Novel Approaches to the Design of Inertial Confinement Fusion Implosions



Jordan Lee

New College

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Trinity 2026

Acknowledgements

My time spent at Oxford during this chapter of my life has been incredible, and I am deeply grateful to everyone who has made it possible and enjoyable. First and foremost, I would like to thank my supervisor, Peter, for his invaluable guidance, support, and encouragement throughout. I am also grateful to him for providing me with the opportunity to work on a range of exciting projects, and for supporting me in pursuing my own research interests. I would also like to thank all the group members that I have had the pleasure of working with for the conversations, collaboration, and general camaraderie that made the experience so rewarding. In particular, Rob, whose early work and guidance helped shape this project, and who was an excellent collaborator throughout, and Sunny, Heath, Rusko, and Robin, who have made office days, group meetings, and conferences especially entertaining. This experience has been made even more enjoyable by the vast number of internal and international conferences, experiments, and... ski trips that I have been fortunate to attend, and I am grateful to all the people that have made these possible. I would also like to thank my examiners for their time and effort reading this very long document (apologies!) - I owe you a drink! Finally, I would like to thank all my friends and family, especially Hannah, for their support and encouragement throughout this process. Of course, none of this would have been possible without the support of my parents, Wendy and John, who throughout my life have always encouraged and facilitated my academic pursuits, and for this I am most thankful.

Abstract

Inertial confinement fusion (ICF) offers a promising pathway toward controlled thermonuclear energy generation. Despite significant progress, achieving robust and efficient ignition remains a major challenge due to the complex, multi-scale physics governing implosion dynamics, the sensitivity of performance to design parameters, and uncertainties in the physical models used to describe material behaviour and energy transport. This thesis addresses these challenges through the development of new approaches to improve both the design and modelling of laser-driven ICF implosions.

First, optimisation frameworks based on Bayesian optimisation and particle swarm methods are developed to efficiently explore the high-dimensional implosion design space. Applications to wetted foam and fast ignition configurations identify improved designs while enabling more efficient and systematic exploration of the parameter space.

Second, a novel auxiliary heating scheme for hotspot-based implosions is investigated. Analytical modelling and kinetic Vlasov-Maxwell simulations are used to characterise beam-plasma instabilities and energy coupling from relativistic electron beams. The applicability of the auxiliary heating mechanism to augment implosion performance is then explored.

Finally, an experimental study measuring the equation of state of low-density silica aerogel foam using laser-driven shock experiments up to 160 GPa is presented. These results provide important constraints for radiation-hydrodynamics simulations and are crucial for accurately modelling their behaviour in ICF implosions.

Together, these contributions address key limitations in current ICF research by combining optimisation techniques, novel physical mechanisms, and improved modelling ability.

Contents

1 Introduction **1**

1.1 A Brief History of Nuclear Fusion 1
1.2 Thesis Focus 6
1.3 Thesis Structure 7
1.4 Author Contributions 9
1.5 Publications 10
1.5.1 First author publications 10
1.5.2 Other publications 10
1.6 Presentations 11
1.7 Grants and Funding 12

2 Theory **13**

2.1 Introduction 14
2.2 Plasma 14
2.2.1 Thermodynamic Equilibrium, Temperature, and the Maxwell-Boltzmann Distribution 15
2.2.2 Plasma Screening, Debye Length, and the Plasma Parameter 16
2.2.3 Plasma Oscillations and Waves 18
2.2.4 Coulomb Collisions 21
2.3 LASERs 25
2.3.1 Stimulated Emission 25
2.3.2 Population Inversion and Practical Lasers 26
2.3.3 Chirp Pulse Amplification and Short Pulse Lasers 26
2.3.4 Electromagnetic Waves in Plasmas 27
2.3.5 Critical Density 28
2.3.6 Absorption Mechanisms 29
2.4 Shocks and the Rankine-Hugoniot Jump Conditions 31
2.4.1 Shock Definition 31
2.4.2 Conservation Laws and the Rankine-Hugoniot Jump Conditions 31
2.4.3 Multiple Shocks and Accessing Off-Hugoniot States 34
2.5 Laser Driven Inertial Confinement Fusion Implosions 35
2.5.1 Capsule Design 35

2.5.2 Laser Driver and Confinement 36
2.5.3 Fusion Fuel 37
2.5.4 Alpha Heating, Ignition and Burn Propagation 39
2.5.5 Fast Ignition 39
2.5.6 Performance Metrics 40
2.6 Instabilities 44
2.6.1 Hydrodynamic Instabilities 45
2.6.2 Parametric Instabilities 47
2.7 Numerical Methods 49
2.7.1 Fluid Description 50
2.7.2 Kinetic Description 53

3 Optimisation Approaches to Inertial Confinement Fusion Implosions 57
3.1 Introduction 57
3.2 General Formulation of Optimisation 59
3.3 Bayesian Optimisation using Gaussian Processes 61
3.3.1 Bayesian Inference (Bayes' Theorem) 62
3.3.2 Gaussian Processes (GPs) 63
3.3.3 Conditioning the GP on Data 71
3.3.4 Acquisition Functions 72
3.3.5 Bayesian Optimisation 73
3.4 Bayesian Package Implementation 75
3.4.1 Computational Complexity 79
3.4.2 Additional Features 81
3.5 Particle Swarm Optimisation (PSO) 90
3.6 PSO Package Implementation 92
3.6.1 Additional Features 93
3.7 Summary 97

4 Optimisation of Inertial Confinement Fusion Implosions 99
4.1 Introduction 100
4.2 Wetted Foam Implosions 101
4.2.1 Wetted Foam Implosions 101
4.2.2 Target Fabrication 102
4.2.3 Convergence Ratio 103
4.2.4 Wetted Foam Experiments 104
4.3 Re-optimisation of Low Convergence Ratio Wetted Foam Implosions 105
4.3.1 Defining the Low Convergence Ratio (Low-CR) Regime . . . 105
4.3.2 Previous Optimisation Approach 106

4.3.3 Re-optimisation of Low-CR implosions . . . 108
4.3.4 Performance Comparison . . . 110
4.3.5 Evolution Through the Parameter Space . . . 112
4.3.6 Bayesian Surrogate Model Accuracy . . . 116
4.3.7 Breakdown of Hydroscaling . . . 118
4.4 Fast Ignition Implosion Optimisation . . . 120
4.4.1 Fast Ignition Implosions . . . 120
4.4.2 Previous Fast Ignition Design Studies . . . 121
4.4.3 Optimisation of Fast Ignition Implosions with Improved Parametric Instability Robustness . . . 121
4.4.4 Optimised Laser Pulse and Compressed Fuel Configurations 123
4.4.5 Modifying the Optimum . . . 125
4.5 Extension of FI Implosions to ArF Drivers . . . 126
4.5.1 Argon Fluoride Laser Technology . . . 126
4.5.2 Optimisation of Fast Ignition Implosions with ArF Drivers . 127
4.6 Summary . . . 128

5 Collisionless Heating of Plasma Using Relativistic Electron Beams 131
5.1 Introduction . . . 132
5.2 System Setup, Parameterisation and Linearisation . . . 134
5.2.1 Multi-Fluid Theory and Governing Equations . . . 135
5.2.2 Linearisation of the Governing Equations . . . 136
5.2.3 Numerical Identification of Modes . . . 139
5.3 Stable Plasma Modes (No Beam) . . . 140
5.3.1 Characterising Electrostatic and electromagnetic Modes . . . 140
5.3.2 Dispersion Relations and Comparison with Theory . . . 142
5.4 Beam-Plasma System and the Onset of Instabilities . . . 143
5.4.1 Beam Modes in Isolation . . . 144
5.4.2 Splitting and Instability of the Beam Mode . . . 144
5.4.3 Two-Stream Instability . . . 147
5.4.4 Weibel Instability . . . 147
5.4.5 Oblique Instabilities and Mixed Modes . . . 148
5.4.6 Comparing Instability Growth Rates . . . 149
5.5 Kinetic Effects and Vlasov-Maxwell Simulations . . . 151
5.5.1 Initial Conditions and Simulation Setup . . . 151
5.5.2 Linear Growth via Bump-on-Tail Instability . . . 153
5.5.3 Saturation from Electron Trapping and Dephasing . . . 156
5.5.4 Non-Linear Mode Coupling . . . 159
5.6 Coupling Efficiency . . . 163

5.6.1 Electron Trapping Theory 164
5.6.2 Electron Momentum Evolution 165
5.6.3 Simulation Parameter Scan 166
5.7 Summary 169

6 Auxiliary Heating Inertial Confinement Fusion Implosions 171
6.1 Introduction 172
6.2 Electron Beam Generation, Transport, and Hotspot Coupling . . . 173
6.2.1 Electron Beam Generation 174
6.2.2 Channel Formation 176
6.2.3 Electron Collisional Scattering and Distribution Evolution . 178
6.2.4 Beam Divergence and Density Evolution 181
6.2.5 Electron Beam Properties at the Hotspot 183
6.2.6 Improving Electron Beam Transport and Divergence 187
6.3 Implementation of Auxiliary Heating in Low-CR Simulations 188
6.3.1 Previous Exploration of AH in ICF Implosions 188
6.3.2 Interface vs Hotspot Energy Deposition 190
6.3.3 Optimum Deposition Time 190
6.3.4 Energy Scan and Implosions Performance 191
6.3.5 Tritium Lean Implosions with Auxiliary Heating 192
6.4 Summary 194

7 Equation of State Measurements of Low-Density Silica Aerogel Foam 197
7.1 Introduction 198
7.2 Experimental Setup 199
7.2.1 Target Design 199
7.2.2 Laser System 200
7.2.3 Deviation From the Plan 201
7.3 VISAR Diagnostic 202
7.4 Determining the Foam Shock State 205
7.4.1 Shock Variables and the Rankine-Hugoniot Relations 205
7.4.2 Imepdance Matching 205
7.4.3 Error Propagation 208
7.5 VISAR Analysis and Results 209
7.5.1 VISAR Analysis 209
7.5.2 Shock Data 211
7.6 Vacuum Gap 213
7.6.1 Vacuum Gap Evidence 213
7.6.2 Characterising the Vacuum Gap 213

7.6.3 Vacuum Gap Simulations . . . 214
7.6.4 Updated Hugoniot . . . 217
7.7 Summary . . . 220

8 Summary and Further Work **221**
8.1 Summary . . . 221
8.2 Further Work . . . 223

Appendices

A Theory Section Derivations **226**
A.1 Maxwell-Boltzmann Distribution . . . 226
A.2 Linearised Governing Equations (Plasma Oscillations) . . . 229
A.3 Coulomb Scattering Angle . . . 229
A.4 Stimulated Emission . . . 230
A.5 ICF Lawson Criterion . . . 231

B Bayesian Optimisation Derivations **233**
B.1 Correlated Multivariate Gaussian Distribution PDF . . . 233
B.2 Cholesky Decomposition . . . 235

C Low-CR Optimisaiton Details **236**
C.1 Low-CR Optimisation Hyper-Parameters . . . 236
C.1.1 Bayesian Hyper-Parameters . . . 236
C.1.2 Particle Swarm Optimisation Hyper-Parameters . . . 236
C.2 Low-CR Optimisation Results . . . 237
C.3 Fast Ignition Optimisation Hyper-Parameters . . . 237

D Beam-Plasma System **242**
D.1 Linearisaing the Lorentz gamma factor . . . 242
D.2 Matrix describing the linearised system, $\mathbf{M}$ . . . 242

References **247**

1 Introduction

Contents

1.1 A Brief History of Nuclear Fusion 1
1.2 Thesis Focus 6
1.3 Thesis Structure 7
1.4 Author Contributions 9
1.5 Publications 10
1.5.1 First author publications 10
1.5.2 Other publications 10
1.6 Presentations 11
1.7 Grants and Funding 12

1.1 A Brief History of Nuclear Fusion

Long before nuclear fusion was understood, the origin of the Sun's energy posed a major scientific problem. In 1862, Lord Kelvin and Helmholtz proposed that the Sun's energy could be explained by gravitational contraction, where the Sun slowly contracts under its own gravity, converting gravitational potential energy into heat. However, this mechanism could only sustain the Sun's luminosity for about 30 million years, which contradicted observations made by paleontologists and Darwin's theory of evolution, suggesting that the Earth and Sun were much

older. This discrepancy highlighted the need for a new energy source to explain the Sun's long-term energy output.

The story of fusion begins around the break of the 20th century. Here, Einstein publishes his theory of special relativity [1] and, with it, the most famous equation known to mankind: $E = mc^2$. This idea, that there lies an equivalence between mass and energy, underpins the entire process of nuclear fusion.

Around this time, Thomson [2] and Aston [3] were working with cathode rays and the use of perpendicular electric and magnetic fields to separate charged particles depending on their charge-to-mass ratio. It was from this research that some of the first mass spectrometers were designed. In 1920, using mass spectrometry, Aston discovered that the mass of four hydrogen nuclei is slightly larger than the mass of a single helium ($^4_2$He) nucleus [4]. This mass deficit provided the first hints of a mechanism by which energy could be released in the fusing of nuclei.

Wasting no time, Eddington [5] speculated that the currently ambiguous source of the Sun's energy could be tied to the large amounts of energy released from the nucleosynthesis of helium. It was further hypothesised that the immense gravitational pressures and temperatures present at the centre of stars could facilitate these nuclear reactions. However, this idea was rebutted by the scientific community as the temperature of the Sun was observed to be too cool for fusion reactions to occur.

The idea was revitalised in 1928 when progress in quantum mechanics and, in particular, the process of quantum tunnelling was introduced by Gamow [6]. Classically, a particle cannot overcome a potential barrier if the particle has insufficient energy. Quantum mechanics, however, concludes that the wavefunction of the particle has non-zero amplitude (exponentially decaying within the potential) on the opposite side of the barrier. This non-zero wavefunction amplitude corresponds to a probability of the particle being found in this location and, hence, the possibility that a particle with classically insufficient energy can 'tunnel' through the potential barrier. When applied to fusion reactions, quantum tunnelling provides a mechanism for fusion reactions to occur below the energy threshold required by classical

mechanics. Atkinson and Houtermans calculated that, when accounting for quantum tunnelling, the temperature of the Sun was, in fact, high enough to produce the required rate of nuclear reactions.

In the early 1930s, huge advancements in nuclear physics were made at the University of Cambridge's Cavendish Laboratory, largely stemming from the leadership and work of Lord Rutherford, of the famous gold foil experiment probing the structure of the atom in 1911 [7]. In building and applying the first ever particle accelerator [8], Cockcroft and Walton fired high-energy protons at a range of targets and, in 1932, split lithium into helium nuclei [9]. This experiment was recognised as the first artificial splitting of the atom and it was in experiments similar to this in which Chadwick discovered the neutron in the same year [10, 11]. Further progressing on this series of breakthroughs, Oliphant began bombarding similar materials with deuterium. It was found that deuterium nuclei could be made to react with one another producing energetic ${}^{3}_{2}\mathrm{He}$ or tritium nuclei [12]. This was the first case of nuclear fusion in the lab and confirmed the large energy output of the reaction. At the end of the decade, Bethe provided more detailed work on the fusion reactions that occur in the Sun, notably the proton–proton chain describing the multi-stage process of fusing ${}^{1}_{1}\mathrm{H}$ to ${}^{4}_{2}\mathrm{He}$ [**ProtonProtonChain**].

During World War II, significant research into the cross section of nuclear reactions was conducted [13, 14]. The cross section essentially refers to the probability of a nuclear reaction occurring as a function of the energy of the interaction and the large cross section of deuterium-tritium reactions at relatively low energy was unearthed. This formed the foundation for the development of the hydrogen bomb, where fission reactions are used to compress, heat and fuse fusion material. Neutrons released from the fusion reactions are then cycled back into the fission material - further stimulating fission reactions and ultimately resulting in an incomprehensible amount of energy release. These bombs were first demonstrated in 1952 and claim their spot in history as the first 'successful' man-made fusion device, albeit with non-containable energy output.

Outside of the use of fusion in bombs, fusion reactors for energy generation using a deuterium-tritium plasma were proposed. Whilst the reactivity of DT reactions peaks around 100 keV, the benefit of plasma confinement is the Maxwellian distribution of particle energies associated with thermal equilibrium. If the plasma can be confined and heated to, say, 10 keV, a small fraction of the particles will still have sufficient energy to fuse. These initial fusion reactions will release and deposit energy back into the plasma, raising the temperature and resulting in further fusion reactions. Collisions within the plasma also act to repopulate the high energy tail of the energy spectrum thereby recycling particles capable of fusion. This type of fusion reaction is know as thermonuclear fusion, in which a high temperature plasma is used to facilitate fusion reactions. From this idea of plasma confinement, magnetic confinement fusion reactors, where external magnetic fields are used to confine the plasma, were initially proposed.

In the early 1950s, Soviet physicists Sakharov and Tamm proposed their idea for the 'tokamak' [15], a new magnetic confinement fusion reactor that improves on previously proposed UK reactor designs [16]. Charged particles interacting with magnetic fields orbit around magnetic field lines, confining the particles along the axis of the magnetic field. It was initially proposed that a solenoid could be used to confine the plasma. A solenoid, however, is open-ended and there will be plasma and energy losses at either side. A simple solution is to bend the solenoid into a torus, connecting the ends and preventing losses. This raises further issues due to the differing strength of the magnetic field at the inner and outer radius of the torus. This radially varying magnetic field strength produces a drift in the particles towards the outside of the chamber. To counteract this drift, a second magnetic field can be used to force the particles into a helical motion. The helical motion, over several cycles, will cause the particles to interact with the magnetic field at the inside and outside of the torus equally and negate the radial drift.

How this second magnetic field causing helical motion is implemented is dependent on the design of the reactor. The tokamak, presented in 1950, essentially treats the reactor as the second arm of a transformer, producing an axial current

and subsequently generating the poloidal magnetic field required for helical motion. To supply the axial current, a constantly increasing magnetic field is required and therefore imposes a maximum confinement time on the tokamak. The stellarator [17], a reactor idea introduced by Spitzer in 1951, removes this pulsed power requirement of the tokamak by using a complex configuration of external magnets. These two designs dominate magnetic confinement fusion research to this day, with reactors being built in a whole host of different countries.

In 1958, Nuckolls and colleagues at the Lawrence Livermore National Laboratory proposed the concept of inertial confinement fusion (ICF) [18, 19]. This process works similarly to that used in the hydrogen bomb, where an external energy source drives the implosion of fusion fuel, compressing it to extremely high temperatures and densities. Under these conditions, fusion reactions occur over very short timescales before the fuel explodes apart. The confinement in this scheme is provided solely by the inertia of the imploding material. In 1960, following the invention of the laser (light amplification by stimulated emission of radiation) [20], Nuckolls proposed the idea of laser-driven fusion, in which laser beams are used to drive the implosion. It is inertial confinement fusion, and in particular laser-driven ICF, which forms another major branch of modern fusion research.

There is a long-standing notorious quip in the fusion community that 'fusion energy is always 20/30 years away' [21]. It is unfortunate that, to this day, this statement still holds true. Whilst the general ideas underpinning fusion reactors remain the same, modern research struggles with efficiently coupling energy to the plasma and counteracting instabilities which act to degrade the performance of the device. Whilst magnetic and laser inertial confinement fusion are not the only proposed ideas, they do dominate modern fusion research. However, pinch-based devices [22] (still falling within the domain of inertial confinement fusion), such as the Z-machine and Scylla [23], the latter of which was the first device to demonstrate controlled nuclear fusion in 1958, deserve notable mention.

1.2 Thesis Focus

The content of this thesis focuses on the laser-driven inertial confinement fusion (ICF) approach, in which high-power lasers are used to compress and heat a capsule containing fusion fuel to conditions suitable for thermonuclear burn. Despite significant progress over the decades, achieving robust and efficient ignition remains a central challenge, limited by the complexity of implosion dynamics, sensitivity to design parameters, and uncertainties in the underlying physical models. This thesis addresses these challenges by developing and applying new approaches to improve both the design and modelling of ICF implosions. The work is organised into three complementary strands:

- **Optimisation of ICF implosion designs.** ICF implosions are characterised by large, high-dimensional parameter spaces describing both target design and laser drive profiles. Exploring these spaces is computationally expensive, as each design evaluation requires radiation-hydrodynamic simulations, making optimisation challenging. Although advanced optimisation techniques have begun to be explored within the ICF community, general-purpose frameworks suitable for expensive ICF optimisation problems are not widely available. This thesis addresses this challenge through the development of complementary Bayesian optimisation and particle swarm optimisation frameworks tailored to ICF design studies. These methods enable more efficient exploration of the design space, identify improved implosion designs, and provide a general optimisation toolset that is accessible to the wider academic ICF community.

- **Auxiliary heating of hotspot-driven implosions.** A novel heating scheme, in which relativistic electron beams are used to collisionlessly deposit energy within a plasma, has recently been proposed for applications in inertial confinement fusion. The aim of this approach is to raise hotspot conditions towards ignition thresholds, or to enhance the gain of already igniting designs. However, the suitability of this heating mechanism under realistic ICF conditions remains uncertain. In this thesis, the underlying

beam-plasma interaction responsible for the auxiliary heating mechanism is investigated, and the conditions required to maximise the energy coupling are identified. Beyond the beam-plasma interaction itself, the generation and transport of relativistic electron beams within realistic ICF implosions are also considered. Together, these studies assess the viability of auxiliary heating for ICF and identify the physical processes that ultimately limit its application.

- **Experimental constraints on equation-of-state models.** The predictive capability of radiation-hydrodynamic simulations depends critically on the accuracy of the equation of state (EOS) models used to describe the response of materials under extreme conditions. As low-density foams receive increasing attention for advanced ICF target designs, the limited experimental data available to constrain their EOS has become an important source of uncertainty in implosion modelling. Improving the predictive capability of these simulations therefore requires new experimental measurements against which EOS models and tables can be benchmarked and refined. This thesis addresses this need through the design and execution of laser-driven shock experiments that measure the principal Hugoniot of 90 mg/cc silica aerogel foam at pressures up to 160 GPa. In addition, a methodology is developed to account for the presence of vacuum gaps in principal Hugoniot experiments, improving the interpretation of these measurements and providing a framework for future EOS studies of similar materials.

Together, these three strands address key limitations in current-day ICF research by combining design optimisation, exploration of novel physical mechanisms, and improved modelling through experimental validation, with the overarching goal of improving the performance and predictive capability of laser-driven fusion implosions.

1.3 Thesis Structure

This thesis is organised to develop the physical background, introduce the methods used, and then present the original contributions across optimisation, auxiliary

heating, and experimental validation of ICF implosions.

Chapter 2 introduces the theoretical framework required for the remainder of the thesis. This includes the fundamental physics of plasmas, laser-plasma interactions, shock dynamics, and inertial confinement fusion, as well as the numerical approaches used to model these systems. This chapter establishes the key concepts and equations that underpin the later analysis.

Chapter 3 presents the optimisation frameworks developed to explore and improve ICF implosion designs. Bayesian optimisation using Gaussian processes and particle swarm optimisation are introduced, alongside their practical implementation. These methods provide the primary tools used in Chapter 4 to navigate the high-dimensional design space of ICF implosions.

Chapter 4 applies these optimisation techniques to the design of inertial confinement fusion implosions. In particular, wetted foam and fast ignition configurations are investigated, and improved designs are identified. This chapter demonstrates how systematic optimisation can enhance implosion performance and provide insight into the underlying parameter space.

Chapter 5 investigates the physics of collisionless heating of plasmas using relativistic electron beams. A combination of analytical theory and kinetic simulations are used to characterise beam-plasma instabilities and energy coupling mechanisms. This provides the physical foundation for the auxiliary heating scheme explored in the following chapter.

Chapter 6 explores the use of the heating mechanism to provide auxiliary heating to inertial confinement fusion implosions. Building on the results of Chapter 5, the integration of electron beam heating into low-convergence-ratio implosions is investigated, with the aim of improving hotspot conditions and overall implosion performance.

Chapter 7 presents an experimental study of the equation of state of low-density silica aerogel foam. Measurements of the principal Hugoniot at pressures up to 160 GPa are obtained using laser-driven shock experiments. These results provide

important constraints for radiation-hydrodynamics simulations and are directly relevant to the modelling and design of wetted foam implosions.

Finally, Chapter 8 summarises the main findings of the thesis and outlines directions for future work.

1.4 Author Contributions

The work presented in this thesis is largely my own and was undertaken under the guidance of my supervisor, Prof Peter Norreys. During the course of the D.Phil, I was day-to-day supervisor of a number of summer and Masters students whose work contributed to aspects of the research. In all cases, I was the primary contributor and led the development and analysis of the work presented in this thesis. The contributions of the supervised students are summarised below:

- Heath Martin (Masters Student): Contributed to investigations of auxiliary heating inertial confinement fusion implosions, particularly multi-pulse approaches and efficiency (related to Chapter 6, but not included in this thesis).
- Rusko Ruskov (Masters Student): Contributed to investigations of auxiliary heating ICF implosions, specifically the nonlinear saturation process (Chapter 5).
- David Coope (Summer and Masters Student):
  - Summer project: Contributed to the development of the Bayesian optimisation package (Chapter 3).
  - Masters project: Explored aspects of the auxiliary heating mechanism, specifically predicting the energy coupling through momentum evolution and investigating auxiliary heating driven by non-idealised electron beam distributions (Chapters 5 and 6).
- Joshua Redfern (Summer Student): Contributed to the development of the particle swarm optimisation package (Chapter 3).

- Ronald Thorpe (Summer Student): Applied the optimisation packages to indirect-drive ICF implosions (work not included in this thesis).
- Julian Mathe (Summer Student): Participated in the equation-of-state experiment conducted at the LULI2000 laser facility (related to Chapter 7, but not included in this thesis).
- Tianchi Liu (Summer Student): Performed kinetic simulations investigating unexpected laser pulse amplification during experiments (work not included in this thesis).
- Zhenshuo Wu (Masters Student): Applied the optimisation frameworks to ArF fast ignition implosions (Section 4.5).

1.5 Publications

1.5.1 First author publications

- J. J. Lee et al. "Toward more robust ignition of inertial fusion targets". *Physics of Plasmas* 30.2 (2023), p. 022702.
- Jordan Lee et al. "Measuring the principle Hugoniot of low-density silica aerogel foam at pressures up to 160 GPa". *Physical Review E* 113 (2026), p. 055210
- J. J. Lee et al. "Bayesian and particle swarm approaches to inertial confinement fusion optimization". *Physics of Plasmas* 33.6 (2026), p. 062702

1.5.2 Other publications

- R. W. Paddock et al. "Energy gain of wetted-foam implosions with auxiliary heating for inertial fusion studies". *Plasma Physics and Controlled Fusion* 66.2 (2024)
- R. J. L. Timmis et al. "Attosecond and nano-Coulomb electron bunches via the Zero Vector Potential mechanism". *Scientific Reports* 14 (2024), p. 10805

- Eduard Atonga et al. “Gravitational waves from high-power twisted light”. *Physical Review D* 110 (2024), p. 044023
- H. Martin et al. “Electrothermal filamentation of igniting plasmas”. *Physical Review E* 110 (2024), p. 035205
- Robin J. L. Timmis et al. “Efficiency-optimized relativistic plasma harmonics for extreme fields”. *Nature* 652.8112 (2026), pp. 1153–1158

1.6 Presentations

- Transfer of Status (2022) - Oxford, UK

  “Auxiliary Heating ICF Hotspots Using Relativistic Electron Beams.”

- Direct Drive and Fast Ignition Workshop (2022) - Madrid, Spain

  “Crossing Relativistic Electron Beams to Heat Dense Plasma.”

- IBM 1 hour Presentation (2023) - Daresbury, UK

  “Multi-Scale Simulation Toolkit for Inertial Confinement Fusion.”

- American Physical Society Division of Plasma Physics (2023) - Denver, Colorado

  “Augmenting Robust Wetted Foam ICF Implosions to Achieve High Gain.”

- NIF/JLF User Group Meeting (2023) - Livermore, California

  “Multi-Scale Simulations of Auxiliary Heated Low-Convergence Ratio Implosions.”

- Matter in Extreme Conditions for Magnatised Plasmas (2024) - Montenevre, France

  “Multi-Scale Simulations of Robust Wetted Foam ICF Implosions.”

- NIF/JLF User Group Meeting (2025) - Livermore, California

  “Bayesian and Particle Swarm Approaches to ICF Optimisaiton.”

- IFE-STAR (2025) - Breckenridge, Colorado

  "Bayesian and Particle Swarm Approaches to ICF Optimisaiton."

- OxCHEDS (2025) - Oxford, UK

  "Bayesian and Particle Swarm Approaches to ICF Optimisation."

- APS-DPP (2025) - Long Beach, California

  "Measuring the Principal Hugoniot of Low-Density Silica Aerogel Foam."

- HPL (2025) - Abingdon, Oxfordshire

  "Measuring the Principal Hugoniot of Low-Density Silica Aerogel Foam."

- **Invited Talk:** PHEDMM - Hirschegg, Austria (2026)

  "Measuring the Principal Hugoniot of Low-Density Silica Aerogel Foam"

- Matter in Extreme Conditions for Magnatised Plasmas (2026) - Montenevre, France

  "Bayesian and Particle Swarm Approaches to ICF Optimisation."

1.7 Grants and Funding

The work was supported by the UKRI-STFC John Adams Institute for Accelerator Science grant ST/V001655/1 as well as the UKRI-EPSRC HEC Plasma Physics EP/X035336/1 grant. The experimental work presented in this thesis was made possible by funding from the Vulcan dark period experiment support grant to conduct the experiment on the GEKKO XII laser facility. Finally, access to the SCARF supercomputer cluster and access the ARCHER2 National Supercomputer Service (https://www.archer2.ac.uk) was provided by the Central Laser Facility and the Scientific Computing Department at the Rutherford Appleton Laboratory.

2 Theory

Contents

2.1 Introduction 14
2.2 Plasma 14
2.2.1 Thermodynamic Equilibrium, Temperature, and the Maxwell-Boltzmann Distribution 15
2.2.2 Plasma Screening, Debye Length, and the Plasma Parameter 16
2.2.3 Plasma Oscillations and Waves 18
2.2.4 Coulomb Collisions 21
2.3 LASERs 25
2.3.1 Stimulated Emission 25
2.3.2 Population Inversion and Practical Lasers 26
2.3.3 Chirp Pulse Amplification and Short Pulse Lasers . . . 26
2.3.4 Electromagnetic Waves in Plasmas 27
2.3.5 Critical Density 28
2.3.6 Absorption Mechanisms 29
2.4 Shocks and the Rankine-Hugoniot Jump Conditions . 31
2.4.1 Shock Definition 31
2.4.2 Conservation Laws and the Rankine-Hugoniot Jump Conditions 31
2.4.3 Multiple Shocks and Accessing Off-Hugoniot States . . . 34
2.5 Laser Driven Inertial Confinement Fusion Implosions 35
2.5.1 Capsule Design 35
2.5.2 Laser Driver and Confinement 36
2.5.3 Fusion Fuel 37
2.5.4 Alpha Heating, Ignition and Burn Propagation 39
2.5.5 Fast Ignition 39
2.5.6 Performance Metrics 40
2.6 Instabilities 44
2.6.1 Hydrodynamic Instabilities 45

2.6.2 Parametric Instabilities 47
2.7 Numerical Methods 49
2.7.1 Fluid Description 50
2.7.2 Kinetic Description 53

2.1 Introduction

The goal of this thesis is to serve as a comprehensive and self-contained document that details the work carried out during the D.Phil project. To achieve this, it is essential to build the physics discussed from the ground up, starting with fundamental principles. This chapter aims to introduce the theoretical foundation that supports the content of the subsequent chapters as intuitively and requiring as little prior knowledge as possible. In the interest of maintaining a self-contained document, effort has been made to include derivations of the mathematical formulas used throughout the text. In many cases, however, these derivations have been moved to Appendices in the interest of space. Key derivations that are especially relevant to the main discussion are presented directly in the text.

2.2 Plasma

The term plasma was first introduced by Irving Langmuir in 1928 to describe a state of matter containing charged particles, such as ions and electrons [32]. This describes an ionised state, in which atoms have acquired sufficient energy for their electrons to overcome the Coulombic attraction of the nucleus, producing free electrons and positively charged ions. Since 1928, the definition of a plasma has become more precise: in addition to being ionised, a plasma must be quasi-neutral and exhibit collective behaviour [33, 34].

Quasi-neutrality is an interesting concept. The "neutral" part is straightforward: the bulk of the plasma must have equal amounts of positive and negative charge, so that, when observed on large scales, the plasma appears electrically neutral.

The "quasi" qualifier indicates that local regions can temporarily deviate from neutrality, allowing for small-scale charge separations.

The second defining property of a plasma is collective behaviour. Because its constituent particles are charged, they rarely interact through direct collisions, but instead via long-range Coulomb forces. Quantitatively, the number of charged particles in a shell of radius r increases as $\propto r^2 dr$, while the Coulomb force from each particle decreases as $\propto r^{-2}$. Thus, the net contribution from each shell is roughly independent of distance, so that distant particles still influence the dynamics significantly. This leads to collective interactions that dominate over short-range collisions, allowing the plasma to respond collectively over lengthscales much larger than the interparticle spacing.

This section on plasmas follows a similar structure to that of the book 'Introduction to Plasma Physics and Controlled Fusion' by Francis F. Chen [33], which is a fantastic introduction to plasma physics, and a rite of passage for all group members! First, the idea of thermodynamic equilibrium is introduced, along with the Maxwell-Boltzmann distribution. Then the response of a plasma to static electric charges and fields is discussed, introducing the concepts of plasma screening, the Debye length, and the plasma parameter. This is then extended to a dynamic system, where plasma oscillations and waves in plasmas are discussed. Finally, the section concludes with a discussion of how energy is primarily redistributed in a plasma through Coulomb collisions.

2.2.1 Thermodynamic Equilibrium, Temperature, and the Maxwell-Boltzmann Distribution

Thermodynamic equilibrium refers to a state in which macroscopic properties of a system, such as temperature, pressure, and density, remain constant over time. This is a fluid perspective of equilibrium, which is concerned with the bulk properties of the system. On a microscopic level, however, the particles that constitute the system are in constant motion, colliding, interacting, and exchanging energy with one another.

From this microscopic perspective, equilibrium corresponds to the condition in which the rate of transitions into and out of each microscopic state are equal [35, 36]. As a consequence, the distribution of particles among available states becomes time-independent. A key result of this principle is the Maxwell-Boltzmann distribution [36, 37], which describes the probability density of finding a particle with momentum $\mathbf{p}$ in equilibrium:

$$f(\mathbf{p}) = N\left(\frac{1}{2\pi m k_B T}\right)^{\frac{d}{2}} \exp\left(-\frac{p^2}{2mk_BT}\right), \tag{2.1}$$

where N is the total number of particles in the system, m is the particle mass, T is the temperature, k_B is the Boltzmann constant, and d is the number of spatial dimensions (usually 3!). The full derivation of the Maxwell-Boltzmann distribution is given in Appendix A.1.

In general, physical systems evolve toward thermal equilibrium, and plasmas are no exception [38]. Interactions between particles continually redistribute energy and momentum, acting to relax deviations from equilibrium. In the non-relativistic limit, this equilibrium is described by the Maxwell-Boltzmann distribution, while in relativistic plasmas the equilibrium distribution generalises to the Maxwell-Jüttner distribution, which incorporates relativistic effects. Thus, on timescales long compared to the characteristic collision or interaction (e.g. through long-range forces) times, the particle distribution relaxes toward these Maxwellian forms.

2.2.2 Plasma Screening, Debye Length, and the Plasma Parameter

Having established what the plasma system looks like in thermal equilibrium, it is a natural next step to consider how the charges in the system respond to electric fields. Consider the introduction of a test charge into an otherwise uniform plasma. In contrast to a neutral gas, the mobile charges in the plasma will rearrange themselves. Electrons, being lighter than the ions, will respond faster, clustering around positive charges and moving away from negative ones.

To describe this more quantitatively, consider a point charge q inserted into the plasma. This introduces an electric potential $\phi(\mathbf{r})$, where $\mathbf{E}(\mathbf{r}) = -\nabla\phi(\mathbf{r})$, which satisfies Poisson's equation

$$-\nabla \cdot \mathbf{E} = \nabla^2\phi(\mathbf{r}) = -\frac{\rho(\mathbf{r})}{\epsilon_0} = -\frac{e}{\epsilon_0}\left(Zn_i - n_e\right), \tag{2.2}$$

where ρ is the total charge density, $n_{i/e}$ are the ion and electron number densities, and Z is the ionisation state. For simplicity, Z will be set to 1 for the rest of this chapter.

As the mass of an ion is significantly larger than the mass of an electron, the electrons will react on timescales much faster than the ions. As such, the ions can be assumed to be stationary as a static neutralising background. The ion density can therefore be taken as the unperturbed electron density infinitely far away from the potential, $n_i = n_0$.

In the presence of the electric field, the electrons will redistribute themselves and settle into a new equilibrium. The electron species therefore follows a new Maxwellian distribution, but with an additional energy term in the form of electrostatic potential, giving

$$f(\mathbf{r}, \mathbf{u}) = A \exp\left[\frac{-\left(\frac{1}{2}mu^2 - e\phi(\mathbf{r})\right)}{k_B T_e}\right]. \tag{2.3}$$

Integrating this distribution over velocity space gives the electron density

$$n_e(\mathbf{r}) = \int f(\mathbf{r}, \mathbf{u})\, \mathrm{d}^3u \tag{2.4}$$

$$= n_0 \exp\left(\frac{e\phi(\mathbf{r})}{k_B T_e}\right), \tag{2.5}$$

where the constant has been chosen such that $n_e \to n_0$ as $\phi \to 0$.

Substituting this into Equation 2.2 yields

$$\nabla^2\phi(\mathbf{r}) = \frac{en_0}{\epsilon_0} \exp\left(\frac{e\phi(\mathbf{r})}{k_B T_e}\right). \tag{2.6}$$

Assuming that the potential energy is small compared to the thermal energy, $e\phi \ll k_B T_e$, the exponential can be linearised using Taylor expansion:

$$\nabla^2\phi(\mathbf{r}) \approx \frac{n_0 e^2}{\epsilon_0 k_B T_e}\, \phi(\mathbf{r}). \tag{2.7}$$

By defining the Debye length as

$$\lambda_D = \left(\frac{\epsilon_0 k_B T_e}{n_0 e^2} \right)^{1/2}, \tag{2.8}$$

the governing equation becomes

$$\nabla^2 \phi(\mathbf{r}) = \frac{1}{\lambda_D^2} \phi(\mathbf{r}). \tag{2.9}$$

Which, in spherical symmetry, $\nabla^2 \phi(r) = r^{-2} \frac{\mathrm{d}}{\mathrm{d}r} \left(r^2 \frac{\mathrm{d}\phi}{\mathrm{d}r} \right)$ has a solution of the form

$$\phi(r) = \frac{q}{4\pi\epsilon_0 r} \exp\left(-\frac{r}{\lambda_D} \right). \tag{2.10}$$

In comparison to the potential of an isolated charge in vacuum, which scales as $\propto 1/r$, the plasma potential acquires an additional exponential decay. This indicates that the free charges in the plasma act to screen the potential, with the Debye length setting the characteristic lengthscale of this screening.

This screening is only effective if there are enough particles in the plasma to provide it. The plasma parameter is therefore defined as the number of particles in a Debye sphere (a sphere of radius λ_D). Collective behaviour, a criterion required for the definition of a plasma, requires that

$$N_D = \frac{4}{3} \pi n_e \lambda_D^3 \gg 1. \tag{2.11}$$

2.2.3 Plasma Oscillations and Waves

Next, let's consider how a plasma responds to a small perturbation in the absence of external magnetic fields. Assuming that, again, the ions remain stationary over the electron response timescales, there are three equations that govern the evolution of this system. The first is the equation of motion of an electron, which is described by the Lorentz force

$$m_e n_e \left[\frac{\partial \mathbf{u}_e}{\partial t} + (\mathbf{u}_e \cdot \nabla) \mathbf{u}_e \right] = -e n_e \left(\mathbf{E} + \mathbf{u}_e \times \mathbf{B} \right), \tag{2.12}$$

with the electric and magnetic fields governed by Maxwell's equations. This form of the Lorentz force neglects collisions and pressure forces. The second is the equation

of continuity, which states that the rate of change of the number of particles in a volume is equal to the net flux of particles across the surface bounding the volume

$$\frac{\partial n_e}{\partial t} + \nabla \cdot (n_e \mathbf{u}) = 0. \tag{2.13}$$

Finally, there is Poisson's equation which was stated earlier (Equation 2.2).

To study the behaviour of small perturbations in a plasma, it is common practice to linearise the governing equations. This allows for the analysis of how infinitesimal disturbances evolve, and whether they grow, decay, oscillate, or propagate as waves.

At the heart of linearisation is Fourier analysis - the greatest gift ever handed to a physicist! The central idea is that any reasonably well-behaved function can be expressed as a continuous superposition of plane waves, each with a specific wavenumber and frequency. Mathematically, this is expressed as a Fourier transform, which decomposes a function into its constituent frequencies:

$$f(x) = \int_{-\infty}^{\infty} \hat{f}(k) e^{ikx} \, \mathrm{d}k, \tag{2.14}$$

or, more generally in the case of travelling waves,

$$f(\mathbf{x}, t) = \int_{\mathbb{R}^n} \int_{-\infty}^{\infty} \hat{f}(\mathbf{k}, \omega) e^{i(\mathbf{k}\cdot\mathbf{x} - \omega t)} \, \mathrm{d}\mathbf{k} \, \mathrm{d}\omega. \tag{2.15}$$

The key insight is that, in a linear system, these plane wave components evolve independently. This means the response of the system can be fully understood by analysing how each wave mode behaves on its own, whether it is amplified (instability), damped, or sustained.

To linearise the governing equations, it is assumed that each variable, such as density, velocity, and electromagnetic field, consist of a uniform equilibrium value and a small perturbation:

$$\mathbf{A}(\mathbf{x}, t) = \mathbf{A}_0 + \mathbf{A_1}(\mathbf{x}, t), \tag{2.16}$$

where $\mathbf{A}_0$ is the unperturbed quantity, assumed constant, and $\mathbf{A}_1$ is a small perturbation such that $|\mathbf{A}_1| \ll |\mathbf{A}_0|$. This assumption of a small perturbation ensures that higher order terms in perturbed quantities can be neglected. By retaining only first order terms the equations are linearised.

Guided by Fourier analysis, the main interest is in the evolution of a plane wave perturbation, which can be expressed as:

$$\mathbf{A_1}(\mathbf{x}, t) \propto \exp\left(i(\mathbf{k} \cdot \mathbf{x} - \omega t)\right). \tag{2.17}$$

Because of this plane wave ansatz, the derivatives of the perturbed quantities can be expressed in terms of the wavenumber, $\mathbf{k}$, and frequency, ω:

$$\frac{\partial \mathbf{A}}{\partial t} = -i\omega \mathbf{A}_1, \quad \nabla \cdot \mathbf{A} = i\mathbf{k} \cdot \mathbf{A}_1, \quad \nabla \times \mathbf{A} = i\mathbf{k} \times \mathbf{A}_1, \tag{2.18}$$

which is equivalent to the substitutions $\nabla \to i\mathbf{k}$ and $\frac{\partial}{\partial t} \to -i\omega$.

The linearised governing equations are

$$i\omega m \mathbf{u}_1 = e(\mathbf{E}_1 + \mathbf{u}_1 \times \mathbf{B}_0), \tag{2.19}$$

$$\omega n_1 = n_0 \mathbf{k} \cdot \mathbf{u}_1, \tag{2.20}$$

$$i\epsilon_0 \mathbf{k} \cdot \mathbf{E}_1 = -e n_1, \tag{2.21}$$

which are derived in Appendix A.2. In the absence of a magnetic field, $\mathbf{B}_0 = 0$, and considering a perturbation along a single axis, these equations can be simplified to

$$i m_e \omega u_1 = e E_1, \tag{2.22}$$

$$\omega n_1 = n_0 k u_1, \tag{2.23}$$

$$ik\epsilon_0 E_1 = -e n_1. \tag{2.24}$$

These are three unique equations involving three perturbed quantities: n_1, E_1, and u_1. These equations can be solved to yield

$$\omega^2 = \frac{n_0 e^2}{m_e \epsilon_0}, \tag{2.25}$$

which describes a system which supports stable ocillations with a frequency of

$$\omega_p = \left(\frac{n_0 e^2}{\epsilon_0 m_e}\right)^{\frac{1}{2}}, \tag{2.26}$$

defined as the plasma frequency.

Intuitively, this oscillatory motion describes particles which have been perturbed from equilibrium and feel a restoring force from the electric field established due to

charge separation. The particle is therefore accelerated back towards equilibrium and accumulates momentum. Due to the non-zero momentum at equilibrium, the particle overshoots, and is accelerated back toward equilibrium on the other side. This process repeats and produces oscillations around the equilibrium at this characteristic plasma frequency.

A similar approach can be taken without neglecting thermal motion. Thermal motion modifies the equation of motion to include a pressure term. Assuming an isotropic, isothermal pressure with a single thermodynamic degree of freedom, and neglecting collisions, viscosity, and thermal conduction, the equation of motion becomes

$$m_e n_e \left[\frac{\partial \mathbf{u}_e}{\partial t} + (\mathbf{u}_e \cdot \nabla)\, \mathbf{u}_e \right] = -e n_e \mathbf{E} - 3 k_B T_e \nabla n_e. \tag{2.27}$$

These approximations are appropriate for ICF plasmas, where the large Reynolds and Péclet numbers imply that viscous and conductive transport are small compared with inertial and pressure-gradient effects over the timescales of interest. Including this pressure term and carrying through the same linearisation procedure results in the Bohm-Gross dispersion relation [39]:

$$\omega^2 = \omega_p^2 + \frac{3}{2} k^2 v_{th}^2, \tag{2.28}$$

where $v_{th}^2 = 2k_B T/m_e$ is the electron thermal velocity.

The key difference is that the frequency now depends on the wavenumber k. Pressure gradients act as an additional restoring force: when the electron density is locally compressed, the pressure drives particles outward, coupling neighbouring regions. This coupling allows disturbances to travel through the plasma. Thus, oscillations that were previously confined to local charge separation become plasma waves capable of propagating and transporting energy across the system.

2.2.4 Coulomb Collisions

Coulomb collisions are the primary mechanism by which energy and momentum are exchanged between particles in a plasma, and are therefore how the plasma evolves toward thermal equilibrium and the Maxwellian distribution. Unlike hard-sphere

collisions in neutral gases, Coulomb collisions between two charged particles are mediated by the long-range Coulomb force

$$\mathbf{F}_{12} = K\frac{\hat{\mathbf{r}}_{21}}{|\mathbf{r}_1 - \mathbf{r}_2|^2}, \tag{2.29}$$

where $K = q_1 q_2/4\pi\epsilon_0$, q_1 and q_2 are the electric charges of particle 1 and 2, and r_1 and r_2 are the particle positions. Considering the case of an electron passing near an ion: the electron is attracted by the Coulomb force, causing its trajectory to deflect. In the case of electron-electron and ion-ion collisions, the particle will be scattered in the opposite direction. The angle by which the particle is deflected, χ, is determined by the impact parameter (the distance of closest approach, b) and the relative speed, v, and is given by

$$\chi(b) \approx \frac{2K}{\mu v^2 b}, \tag{2.30}$$

where $\mu = m_a m_b/(m_a + m_b)$ is the relative mass of the particles. This equation is derived in Appendix A.3 and assumes that the angle of deflection is small and that scattering through large angles is largely a result of cumulative small angle deflections.

Since the angle of deflection of a single Coulomb collision is often very small, it is therefore difficult to actually call it a 'collision'. Additionally, the electron is equally likely to be deflected in all directions, with the average scattering angle being zero. A usful quantity to help define a collision is the change in velocity

$$\Delta v = v\sin\chi(b) \approx v\chi(b) = \frac{2K}{\mu v b}. \tag{2.31}$$

While the average of this velocity change is zero, the quantity of interest is the deviation from this average value, the root-mean-square $\langle(\Delta v)^2\rangle$. The rate at which this $\langle(\Delta v)^2\rangle$ changes is every time a Coulomb interaction occurs. This is given by $n_i\sigma v$, where σ is the cross section of the interaction. For an impact parameter

b, the cross section is given by $2\pi b db$, and so

$$\frac{d}{dt}\langle(\Delta v)^2\rangle = \int_{b_{\min}}^{b_{\max}} 2\pi b n_i v (\Delta v)^2 \, db \tag{2.32}$$

$$= \frac{8\pi n_i K^2}{\mu^2 v} \ln\left(\frac{b_{\max}}{b_{\min}}\right) \tag{2.33}$$

$$= \frac{8\pi n_i K^2}{\mu^2 v} \ln(\Lambda), \tag{2.34}$$

where $\ln(\Lambda) = \ln(b_{\max}/b_{\min})$ is the Coulomb logarithm.

The effectiveness of Coulomb scattering depends on the ratio of the maximum to minimum impact parameter. The largest relevant scale is the Debye length, λ_D (Equation 2.8), beyond which collective screening suppresses the interaction. The minimum impact parameter can be determined two separate ways: classically, by considering the interaction where the electron collides head on with the ion, or quantum mechanically, by describing the electron as a wave with a wavelength given by the de Broglie wavelength. In the classical case, in a head-on collision, the electron will be reflected, and the change in velocity would be $\Delta v = 2v$. The impact parameter of this interaction is classically the closest the particles can be, where the kinetic energy matches the potential energy. This classical minimium impact parameter is:

$$b_{\min}^{\text{Classical}} = \frac{2K}{\mu v^2}. \tag{2.35}$$

Quantum mechanically, the electron is spread out over a range

$$b_{\min}^{\text{Quantum}} = \lambda_{dB} = \frac{\hbar}{\mu v}, \tag{2.36}$$

making an impact parameter smaller than this value impossible to define. The minimum impact parameter is therefore taken as the maximum of either of these values

$$b_{\min} = \max\left(b_{\min}^{\text{Classical}}, b_{\min}^{\text{Quantum}}\right). \tag{2.37}$$

Using the classical definition of the minimum impact parameter, the argument of the Coulomb logarithm for a particle travelling at the thermal velocity is given by

$$\Lambda = \frac{b_{\max}}{b_{\min}} = \frac{3}{Z} N_D, \tag{2.38}$$

which is approximately equal to the plasma parameter.

The collision time, τ, between two species can now be well defined. Defining a collision as the particle being deflected by 90 degrees, $(\Delta v)^2$ is of the order v^2, and the collision time is given by

$$\frac{1}{\tau}v^2 = \frac{8\pi n_b K^2}{\mu^2 v}\ln(\Lambda), \tag{2.39}$$

$$\tau = \frac{2\pi\epsilon_0^2\mu^2 v^3}{nZ^2e^4\ln(\Lambda)}. \tag{2.40}$$

This formulation of the collision time can then be used to compare the impact of collisions between different species. For both electron-electron and electron-ion interactions, $\mu \approx m_e$ and $v \approx v_{e,\mathrm{th}}$. However, ion-ion collisions have $\mu \approx m_i$ and $v \approx v_{i,\mathrm{th}} = v_{e,\mathrm{th}}/\sqrt{m_i}$, resulting in

$$\frac{\tau_{ii}}{\tau_{ee/ei}} \propto \sqrt{\frac{2m_i}{m_e}}. \tag{2.41}$$

Electrons therefore thermalise faster than the ions. Although electron-ion collisions occur frequently, the large ion-to-electron mass ratio means that very little energy is exchanged per collision. As a result, equilibration of ion and electron temperatures is much slower than equilibration within each species. This allows plasmas, particularly in ICF conditions, to sustain distinct electron and ion temperatures for significant periods.

The typical distance a particle travels before significant cumulative deflections is given by the mean free path, $\lambda_{\mathrm{mfp}} = v\tau$. Often, plasmas are referred to as collisional or collisionless. This terminology refers to the comparison between the Coulomb mean free path and the characteristic lengthscale of the system.

Coulomb collisions are relatively weak compared to collective plasma oscillations, especially in high-temperature or low-density plasmas where $\lambda_{mfp} \gg \lambda_D$. Nevertheless, they provide the microscopic mechanism for momentum and energy exchange between particles. Over long timescales, they ensure that deviations from equilibrium relax toward a Maxwellian distribution and that different species (electrons and ions) exchange energy until they reach a common temperature. Thus, while collisionless dynamics govern short-timescale plasma behaviour, Coulomb collisions underpin the transport and equilibration processes required to relax to thermal equilibrium.

2.3 LASERs

LASER was initially used as an acronym for 'Light Amplification by Stimulated Emission of Radiation,' but is now accepted as a word of it is own. As the acronym implies, light can be amplified when interacting with excited atoms by stimulating the emission of more radiation. Key characteristics of laser light include its narrow spectral bandwidth, the ability to remain collimated over long distances, and its high degree of spatial and temporal coherence. These properties enable lasers to deliver extremely high intensities, to be focused to small spots, and to be precisely controlled in both time and space. Lasers play a crucial role in laser driven inertial confinement fusion, by supplying the high-energy-density source required to drive the implosion. Due to their narrow bandwidth and high degree of coherence, laser beams form a highly non-thermal driving source, enabling energy to be concentrated and coupled into the plasma far more efficiently than would be possible using thermal radiation.

2.3.1 Stimulated Emission

The fundamental idea underpinning lasers was first introduced in 1917 by Einstein [40], although lasers were only experimentally realised in 1954 [41]. It was known at that point that bodies in thermal equilibrium emit electromagnetic radiation that follows a Plankian distribution [42] of the form

$$\rho(\omega) = \frac{\hbar\omega^3}{\pi^2 c^3} \frac{1}{\exp\left(\frac{\hbar\omega}{k_B T}\right) - 1}, \tag{2.42}$$

where ρ is the spectral energy density.

Einstein set out to derive this form from the consideration of photons interacting with matter. In order to fully derive the Plankian distribution, it was necessary to introduce a term to the emission process which depends on ρ. The derivation of the stimulated emission term is provided in Appendix A.4. This additional term requires that the rate of de-excitation is dependent on the photon density. Stimulated emission is therefore a strict requirement for the Plankian distribution of a body in thermal equilibrium. This theoretical framework indicates that the rate of photon emission of a specific wavelength can be increased by a seed pulse

of photons into the system. Stimulated emission is not just a mathematical trick, but the physical process that allows for light amplification.

2.3.2 Population Inversion and Practical Lasers

For stimulated emission to dominate over absorption, a non-thermal distribution of atomic populations is required, known as a population inversion ($N_2 > N_1$). In practice, this is achieved by optical or electrical pumping of a gain medium. For example, in inertial confinement fusion facilities, neodymium-doped yttrium aluminum garnet (Nd:YAG) or neodymium-doped phosphate glass (Nd:glass) are excited using flash lamps [43, 44], producing emission around $\lambda = 1053$ nm. Alternatively, excimer (excited dimer) lasers [45, 46] (e.g. ArF and KrF) can be pumped electrically to combine the constituent atoms (e.g. Ar and F into ArF), which can then be separated to produce emission at specific wavelengths (e.g. ArF at $\lambda = 193$ nm, and KrF at $\lambda = 248$ nm).

The laser medium is placed inside an optical resonator, typically a pair of mirrors, ensuring that photons traverse the gain medium multiple times. This feedback mechanism, together with population inversion, produces a highly coherent, collimated laser beam.

2.3.3 Chirp Pulse Amplification and Short Pulse Lasers

There is a limit to the intensity and fluence a laser pulse can reach before the onset of non-linear effects and damage to the optical equipment. For long pulse lasers (e.g. the National Ignition Facility (NIF) [47], with long pulse laser duration of $\mathcal{O}(10$ ns)) the beam can be streched spatially, amplified, and focused by the final lens onto the target. This allows facilities like the NIF to acheive peak powers of 500 TW (split between 192 beams) and delivering up to 2.2 MJ (as of 2026) laser energy onto the target.

However, to reach ultra-high intensities ($> 10^{18}$ W/cm^2), the pulse must be compressed both spatially and temporally. To do this, a seed pulse with some initial spread of frequencies (bandwidth) is first temporally stretched by passing

it through a dispersive optical system, which introduces a frequency-dependent delay (chirp). This increases the pulse duration by several orders of magnitude, significantly reducing its peak intensity and allowing safe amplification. After amplification, the pulse is passed through a second dispersive system that reverses the initial chirp, recompressing the pulse to its original duration while maintaining the increased energy. This results in an ultra-short, high-energy pulse that can be tightly focused to achieve extremely high intensities. This technique is known as chirped pulse amplification (CPA) [48], for which Donna Strickland and Gérard Mourou were awarded the 2018 Nobel Prize in Physics.

2.3.4 Electromagnetic Waves in Plasmas

To determine how lasers propagate in plasmas, the linearisation approach introduced in Subsection 2.2.3 can be used. In the case of electromagnetic waves, there is now a non-zero magnetic field. As such, two more of Maxwell's equations must be satisfied

$$\nabla \times \mathbf{E} = -\frac{\partial \mathbf{B}}{\partial t}, \tag{2.43}$$

$$c^2 \nabla \times \mathbf{B} = \frac{\mathbf{j}}{\epsilon_0} + \frac{\partial \mathbf{E}}{\partial t}, \tag{2.44}$$

which can be linearised and combined by taking the curl of Equation 2.43 or the time derivative of Equation 2.44 to yield

$$\left(\omega^2 - c^2 k^2\right) \mathbf{E}_1 + c^2 \mathbf{k} \left(\mathbf{k} \cdot \mathbf{E}_1\right) = \frac{-i\omega}{\epsilon_0} \mathbf{j}_1. \tag{2.45}$$

For electromagnetic waves, the electric field is perpendicular to the direction of propagation, so $\mathbf{k} \cdot \mathbf{E}_1 = 0$, and the equation simplifies to

$$\left(\omega^2 - c^2 k^2\right) \mathbf{E}_1 = -i\omega \mathbf{j}_1 / \epsilon_0. \tag{2.46}$$

$\mathbf{j}_1$ can then be related to the density and velocity of the electrons via

$$\mathbf{j}_1 = -n_0 e \mathbf{v}_1. \tag{2.47}$$

Using the linearised equation of motion introduced earlier (Equation 2.19), and again assuming the external magnetic field $\mathbf{B}_0 = 0$, the following dispersion

relation can be obtained

$$\omega^2 = \omega_p^2 + c^2 k^2. \tag{2.48}$$

Note that the choice of $\mathbf{k} \cdot \mathbf{E}_1 = 0$ nullifies Equation 2.21 and subsequently Equation 2.20.

This is the dispersion relation of light in a plasma, which describes how electromagnetic waves propagate through a plasma. From this dispersion relation, the phase and group velocities can be written as

$$v_\phi = \frac{\omega}{k} = \frac{c}{\sqrt{1 - \frac{\omega_p^2}{\omega^2}}}, \quad v_g = \frac{\partial \omega}{\partial k} = c\sqrt{1 - \omega_p^2/\omega^2}, \tag{2.49}$$

and the corresponding refractive index is

$$n = \frac{c}{v_\phi} = \sqrt{1 - \frac{\omega_p^2}{\omega^2}}. \tag{2.50}$$

An interesting feature is that the refractive index is always less than unity for propagating waves ($\omega > \omega_p$), an uncommon property in optics. This implies that rays bend away from regions of higher plasma density, opposite to their behaviour in ordinary dielectrics. When combined with the ponderomotive expulsion of plasma from the high-intensity regions of a laser pulse, this effect can lead to self-focusing of light within the plasma. Remarkably, such self-focusing can concentrate light beyond the normal diffraction-limited focusing constraints, enabling extremely high intensities to be reached inside the plasma.

A variety of other wave modes [49] can be derived by relaxing the assumption that $\mathbf{B}_0 = 0$, such as the ordinary and extraordinary modes, and the right (whistler) and left circularly polarised modes.

2.3.5 Critical Density

Equation 2.48 introduces another interesting feature of plasmas: there is a cutoff frequency below which light waves cannot propagate. If $\omega < \omega_p$, there is no real k that satisfies the dispersion relation, and the refractive index becomes imaginary. Waves at these frequencies are therefore unable to propagate and are reflected.

The key insight is to observe that the plasma frequency is dependent only on the plasma density (Equation 2.26). As a result, there is a density above which light of a given frequency is reflected. This is referred to as the critical density, n_c, and is given by

$$n_c = \frac{\epsilon_0 m_e}{e^2}\omega^2. \tag{2.51}$$

This is a crucial property of plasmas, and a difficulty in laser driven ICF, as it limits the depth at which a laser can deposit energy.

For example, for the 351 nm third harmonic of the Nd:glass laser used at the NIF, the critical density is $n_c \sim 9 \times 10^{21}\,\mathrm{cm}^{-3}$. This is much lower than the solid density of the capsule, which is $\mathcal{O}(10^{23}\text{–}10^{24})\,\mathrm{cm}^{-3}$. As a result, the laser cannot penetrate into the dense capsule shell; instead, it is absorbed or reflected at the critical surface in the coronal plasma. The energy is then transported inward by thermal conduction, shocks, and hydrodynamic motion, rather than by direct laser penetration. This limits the efficiency of laser coupling to the fuel and highlights the importance of the critical density in understanding energy coupling.

2.3.6 Absorption Mechanisms

There are a range of mechanisms by which laser light is absorped by an ICF capsule [50]. Initially, the laser pulse interacts with the solid surface of the ICF capsule. Here, multi-photon ionisation [51, 52] liberates some bound electrons from their atoms. These liberated electrons are then accelerated by, and oscillate in, the electric field of the laser pulse and can collide with surrounding atoms, liberating further electrons via collisional ionisation and leading to a cascade of ionisation. This ionisation process leads to the heating and theromodynamic expansion of the plasma from the surface.

Once the material has expanded, the laser interacts with the blown-off ionised coronal plasma. In this regime, the dominant absorption mechanism is inverse bremsstrahlung [53, 54]. Inverse bremsstrahlung can be understood as a three-body process: an electron oscillates in the laser electric field while undergoing Coulomb

collisions with nearby ions. During these collisions, the electron can absorb a photon from the laser field, converting the oscillatory energy into thermal energy. This is the inverse process of bremsstrahlung emission, in which an electron radiates a photon while being decelerated by an ion.

The absorption rate of inverse bremsstrahlung depends strongly on the collision frequency between electrons and ions, and therefore on both plasma density and temperature. The effective absorption coefficient can be written as [55]

$$\alpha_{\mathrm{IB}} \propto \frac{nZ\nu_{ei}}{\omega^2}, \tag{2.52}$$

where ν_{ei} the electron-ion collision frequency. Since $\nu_{ei} \propto nZ/T^{3/2}$, inverse bremsstrahlung absorption is most efficient in relatively cool, dense plasmas, and becomes weaker as the coronal plasma expands and heats.

Eventually, the laser pulse propagates through the low-density coronal plasma and reaches the critical density surface. If the pulse strikes the surface obliquely and the electric field has a component normal to the surface, resonant absorption can occur [50]. In this case, the laser frequency matches the local plasma frequency, allowing the electric field to resonantly drive electron plasma waves. After this interaction, the laser light is reflected at the critical surface and propagates back through the underdense plasma, where it can undergo further absorption via inverse bremsstrahlung.

In addition to inverse bremsstrahlung and resonant absorption, laser energy can be coupled into the plasma through a variety of laser-plasma interactions (LPIs), which occur resonantly near the quarter-critical ($n_e \simeq n_c/4$) and critical surface. These include stimulated Raman and Brillouin scattering, which redirect energy away from the target, as well as two-plasmon decay, which generates hot electrons that preheat the fuel. These processes are generally detrimental to implosion performance and are discussed in more detail in Section 2.6.2.

For very short laser pulses (sub-picosecond), the plasma has insufficient time to expand significantly during the interaction, and the laser couples to the material very close to the initial surface. In this ultrashort-pulse regime, additional mechanisms

dominate. One is vacuum (Brunel) heating [56], in which the strong oscillating fields extract electrons into the vacuum each half-cycle; these electrons are then rapidly accelerated back into the target, depositing their energy efficiently. Another is $\mathbf{J} \times \mathbf{B}$ heating [57], where the oscillatory magnetic field drives transverse electron oscillations into the material, again resulting in effective energy deposition.

2.4 Shocks and the Rankine-Hugoniot Jump Conditions

A laser incident on a target surface results in rapid heating and expansion of the surface, producing a high pressure region that presses on the underlying material. These large pressure gradients drive strong shocks into the material which is compressed and heated as the shock passes. The shocks are a key part of inertial confinement fusion implosions, as they are the primary mechanism by which laser energy is transferred to the capsule to compress and heat the fusion fuel.

2.4.1 Shock Definition

A shock wave [55] is a rapidly moving surface that produces an abrupt, discontinuous change in the physical properties of a material. Ahead of the shock front, the material exists in its initial state, denoted by subscript 0 in Figure 2.1, characterised by properties such as density (ρ_0), particle velocity (u_0), pressure (P_0), and specific internal energy (ϵ_0). As the shock passes through, it compresses and heats the material, transitioning it into a new state behind the shock front, labelled with subscript 1, where these properties - ρ_1, u_1, P_1, and ϵ_1 - have sharply changed.

2.4.2 Conservation Laws and the Rankine-Hugoniot Jump Conditions

The shock transition occurs over a very thin region, and the changes across this interface are governed by the fundamental conservation laws of mass, momentum, and energy. Applying these laws to a control volume enclosing the shock leads

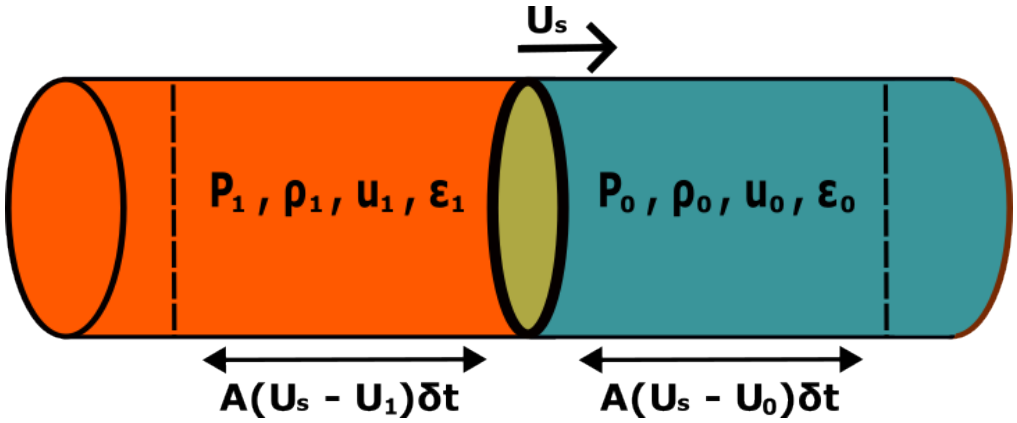


Figure 2.1: Illustration of a shock wave propagating through a material. The shock moves with velocity U_s, and the material ahead of the shock is in an initial state described by subscript 0. As the shock passes through the material, the shock introduces a discontinuous jump in material properties, transitioning material behind the shock to a state described by subscript 1.

to the Rankine-Hugoniot jump conditions, which relate the upstream (unshocked) and downstream (shocked) states.

Considering a shock moving at velocity U_s through a uniform material, the shock will convert a volume $A(U_s - u_0)\delta t$ of unshocked material to a volume of $A(U_s - u_1)\delta t$ of shocked material in a time δt, where u_0 and u_1 are the unshocked and shocked material velocities (particle velocities in the lab frame) respectively. Conserving mass, momentum, and energy across this control volume gives the Rankine-Hugoniot equations [55, 58]:

$$\rho_1(U_s - u_1) = \rho_0(U_s - u_0), \tag{2.53}$$

$$P_1 - P_0 = \rho_1(U_s - u_1)u_1 - \rho_0(U_s - u_0)u_0, \tag{2.54}$$

$$\rho_1(U_s - u_1)\left(\epsilon_1 + \frac{(U_s - u_1)^2}{2} + \frac{P_1}{\rho_1}\right) = \rho_0(U_s - u_0)\left(\epsilon_0 + \frac{(U_s - u_0)^2}{2} + \frac{P_0}{\rho_0}\right). \tag{2.55}$$

In the case where the material ahead of the shock is stationary ($u_0 = 0$),

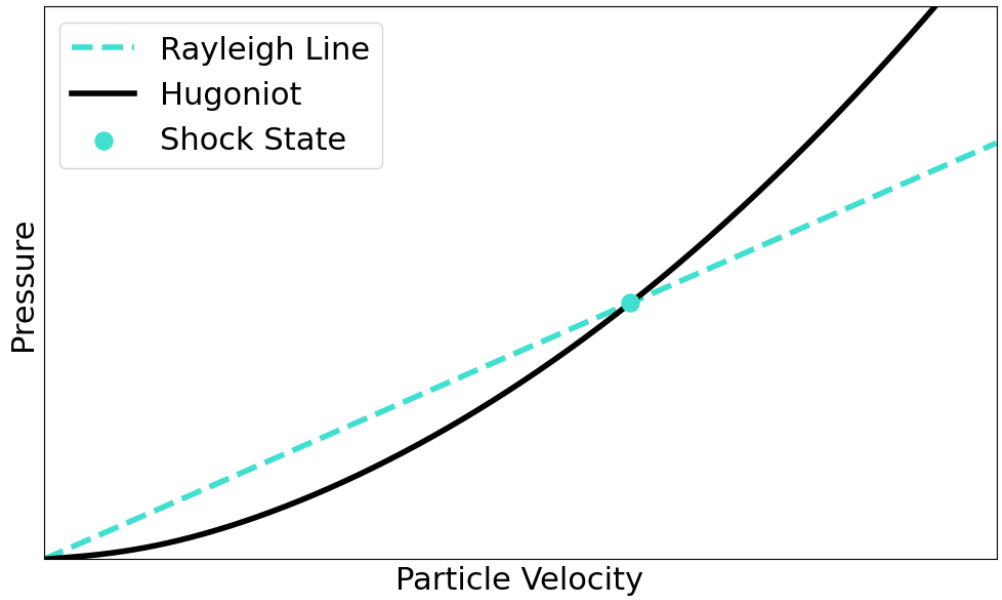


Figure 2.2: Typical Hugoniot curve in the particle velocity (u_p)-pressure (P) plane, together with the corresponding Rayleigh line. Shocked material states must lie on both the Hugoniot and Rayleigh line. The intersection of the two therefore uniquely determines the material shock state.

the equations reduce to:

$$\rho_1 = \frac{\rho_0 U_s}{(U_s - u_1)}, \tag{2.56}$$

$$P_1 - P_0 = \rho_0 U_s u_1, \tag{2.57}$$

$$\epsilon_1 - \epsilon_0 = \frac{1}{2}(P_0 + P_1)\left(\frac{1}{\rho_0} - \frac{1}{\rho_1}\right). \tag{2.58}$$

These relations form a coupled system: knowing any two of the shock variables (U_s, u_1, P_1, ρ_1, and ϵ_1) allows all the shock variables to be calculated, defining the unique shock state.

The possible states that the material can be shocked into from a single shock are constrained by the Rankine-Hugoniot equations. For a given initial state, the locus of all possible end states that can be reached by a single shock is referred to as the Hugoniot. In the case of the upstream material being in a rest state, this is referred to as the principal Hugoniot. The Hugoniot is typically visualised in (P, u_p) or (P, V), where $V = 1/\rho$ is the specific volume, space as shown in Figure 2.2.

Equation 2.57 shows that the change in shock pressure is linearly proportional to the particle velocity. As the Hugoniot is often visualised in the (P, u_p) plane,

for a measured shock velocity this constraint appears as a straight line, known as the Rayleigh line. The slope of the Rayleigh line is proportional to the shock velocity, and the intersection of the Rayleigh line with the Hugoniot defines the shock state. The Rayleigh line is also indicated in Figure 2.2.

It can also be seen from Equation 2.56 that $U_s \propto \rho_1$. As such, the shock velocity increases as the density it is travelling through increases. As a result of this, if the shock is assumed to have some finite width, the rear of the shock will be travelling through higher density material and will therefore be travelling faster than the front of the shock. This leads to the shock steepening and the justification as the why shocks present as discontinuous jumps.

2.4.3 Multiple Shocks and Accessing Off-Hugoniot States

In inertial confinement fusion implosions, a sequence of shocks is launched through the target [59]. The initial shock drives the material to a state on the principal Hugoniot. Each subsequent shock then propagates through material that has already been shocked, and therefore follows a new Hugoniot originating from the previous state. In this way, successive shocks enable access to thermodynamic states that lie off the principal Hugoniot.

This multi-shock approach allows a broader range of material states to be reached, in particular achieving higher densities at a given pressure than would be possible with a single shock. Consequently, the required drive pressure is reduced, lowering the laser energy and intensity needed to compress the fuel to the desired conditions.

The shocks are launched with increasing strength, such that each subsequent shock catches up with the preceding ones, eventually leading to shock coalescence. This staged compression reduces entropy generation compared to a single strong shock, enabling higher final densities to be achieved. The shocks are typically timed to coalesce within the fuel layer, producing a dense shell that compresses the central, lower-density vapour. Precise timing of this coalescence is critical, and mistimed shocks can severely degrade implosion performance.

2.5 Laser Driven Inertial Confinement Fusion Implosions

Inertial confinement fusion (ICF) refers to a method of achieving thermonuclear fusion in which the only confinement mechanism is the nuclei's own inertia. This is acheived by rapidly imploding a target containing fusion fuel. As the target implodes, the fusion fuel is compressed and heated to form a dense, high temperature plasma where fusion reactions occur. As the plasma is heated, it is internal pressure increases. Once the internal pressure exceeds the imploding pressure, the plasma will begin to decelerate and eventually expand outward, leading to loss of confinement, cooling, and the fusion reactions will cease. Therefore, the reactions must occur within a very short time frame.

The implosion process in ICF can be driven by a variety of methods, including large magnetic fields [22], particle beams [60], projectiles [61], and, in the case of nuclear bombs, fission reactions. However, the most commonly explored approach for fusion power generation is laser driven ICF [62], where high power lasers are used to rapidly heat and compress a small spherical capsule containing nuclear fusion fuel.

2.5.1 Capsule Design

Standard laser driven ICF target designs are based on spherical capsules consisting of layered shells. A simplified capsule can be described by three material layers: an outer ablator shell (often made from a CH compound), an inner shell containing the bulk of the fusion fuel, and a central cavity filled with low-density fusion fuel vapour. This standard target design is illustrated in Figure 2.3.

While this capsule description captures the basic structure, real target designs often incorporate additional complexities, such as dopants and high-Z layers. These features primarily mitigate preheat from energetic electrons or high frequency radiation, which would otherwise degrade the performance of the implosion. An additional feature that has been the subject of recent research and is discussed in detail in this thesis is the use of foam layers within the capsule. Foam layers and their purposes are discussed in more detail in Section 4.2.1 and 7.1.

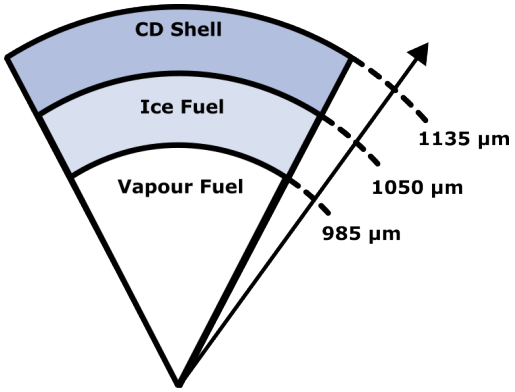


Figure 2.3: Standard laser driven ICF capsule design, consisting of an outer ablator shell, an inner shell containing the bulk of the fusion fuel, and a central cavity filled with low-density fusion fuel vapour. The target dimensions displayed are those used at the NIF to achieve gain greater than 1 for the first time [63].

2.5.2 Laser Driver and Confinement

Even laser driven ICF can be subdivided further into direct [64] and indirect drive [65]. In direct drive, the lasers are focused directly onto the capsule, while indirect drive focuses the laser pulses onto a hohlraum, a high-Z (often gold) cylindrical cavity. In indirect drive (also called X-ray direct drive), the capsule sits in the centre of the cavity, and the laser energy is absorbed by the hohlraum walls, which is then re-emitted as a bath of X-ray radiation which illuminates the capsule. There are trade-offs to consider between the two approaches: direct drive is more efficient, with the laser light directly illuminating the capsule, while indirect drive provides a more uniform bath of radiation, reducing instability growth.

Regardless of which approach is used, the implosion is driven by the ablator layer absorbing energy from incoming radiation (either direct laser light or X-rays), which then heats and expands. This expansion imparts momentum onto the underlying fuel layer, driving it inward in a spherical implosion. The temporal intensity profile can be tailored to drive the required series of shocks through the target.

Once the imploding shell has coalesced, a central region is compressed to high temperatures and densities, forming a 'hotspot'. This hotspot is where fusion reactions begin. As fusion reactions occur, DT nuclei collide and fuse into alpha particles and neutrons. These fusion reactions release energy, which can further

heat the hotspot and drive more fusion reactions. If enough fusion reactions occur, the heating can spread to the surrounding dense fuel layer, igniting it and releasing a large amount of fusion energy.

This all must occur within a very short time frame, typically on the order of 10s of ns, due to the confinement mechanism. As the hotspot is compressed and heats, the internal pressure of the hotspot increases. Once the pressure in the hotspot exceeds the pressure of the imploding shell, the material will stagnate and begin the explode outward, leading to a loss of confinement, cooling of the hotspot, and fusion reactions ceasing.

2.5.3 Fusion Fuel

There are many products that release energy when fused. A non-exhaustive list of the most relevant of these fusion reactions is [66]:

$$\mathrm{D} + \mathrm{D} \rightarrow \mathrm{T}\big(1.01\,\mathrm{MeV}\big) + \mathrm{p}\big(3.02\,\mathrm{MeV}\big) \tag{2.59}$$

$$\mathrm{D} + \mathrm{D} \rightarrow {}^{3}\mathrm{He}\big(0.82\,\mathrm{MeV}\big) + \mathrm{n}\big(2.45\,\mathrm{MeV}\big) \tag{2.60}$$

$$\mathrm{D} + \mathrm{T} \rightarrow \alpha\big(3.5\,\mathrm{MeV}\big) + \mathrm{n}\big(14.1\,\mathrm{MeV}\big) \tag{2.61}$$

$$\mathrm{D} + {}^{3}\mathrm{He} \rightarrow \alpha\big(3.6\,\mathrm{MeV}\big) + \mathrm{p}\big(14.7\,\mathrm{MeV}\big) \tag{2.62}$$

$$\mathrm{T} + \mathrm{T} \rightarrow \alpha + 2\,\mathrm{n} + 11.3\,\mathrm{MeV} \tag{2.63}$$

$$\mathrm{p} + {}^{6}\mathrm{Li} \rightarrow \alpha\big(1.7\,\mathrm{MeV}\big) + {}^{3}\mathrm{He}\big(2.3\,\mathrm{MeV}\big) \tag{2.64}$$

$$\mathrm{p} + {}^{7}\mathrm{Li} \rightarrow 2\,\alpha + 17.3\,\mathrm{MeV} \tag{2.65}$$

$$\mathrm{p} + {}^{11}\mathrm{B} \rightarrow 3\,\alpha + 8.7\,\mathrm{MeV} \tag{2.66}$$

There are a few key parameters that determine the suitability of a fusion reaction for energy generation. The most obvious consideration is the total energy generated per fusion reaction. The more energy released per reaction, the more energy that will eventually be released as heat for electrical energy generation. There are, however, additional physical and safety related considerations.

First, neutrons are often generated in these reactions. Neutrons are neutral and interact very weakly with the plasma, and they therefore escape the plasma with very

high energies. These high energy neutrons then interact with the surrounding reactor materials, activating them and leading to long-term radioactive waste. Additionally, since these neutrons escape the plasma without depositing their energy, they do not contribute to plasma heating and sustaining the high temperatures required for fusion reactions. For sustaining fusion conditions, aneutronic reactions, of which neutrons are not a product, are therefore preferred [67]. However, a lack of neutrons presents challenges for reactor design, as the neutrons are typically used to breed tritium fuel and to generate heat for electricity generation.

Second, some of these fusion reactants are scarce or radioactive. For example, tritium is both [68], appearing in only trace amounts in nature and having a half-life of 12.3 years. As such, tritium must be produced, typically by bombarding lithium with neutrons. It is therefore possible for a reactor to breed it is own fuel by using the neutrons generated from the fusion reactions. However, this would add additional complexity to the reactor design, and again adds safety concerns with the handling of radioactive materials. The radioactive nature of the tritium also causes concerns in target fabrication, with targets needing to be produced and fielded within about a month of each other to ensure the correct fuel mixture is present in the target and the higher Z He dopant, of which the tritium decays into, is minimised. Tritium also causes issues when paired with wetted foam target designs, as the heating from tritium decay could comprimise the foam structure.

However, the most important consideration at the current stage of fusion research is the temperature required to get the products to fuse. Different fusion reactions have different cross-sections – a measure of the probability of a fusion reaction occurring between two nuclei. The cross-section is dependent on the centre-of-mass energy of the ineracting nuclei. In a thermal plasma, this leads to a temperature dependent reactivity $\langle \sigma v \rangle$ and different reactions have different temperature ranges where the reactivity is maximised [69]. Currently, facilities are capable of reaching temperatures around 10 keV, and as such, reactions that can occur at lower temperatures are preferred. This is the primary reason why deuterium and tritium (DT), which has substantially higher reaction rates in the

sub-100 keV range (peaking around 60 keV), is the fuel of choice for current fusion research, along with it is high energy yield.

2.5.4 Alpha Heating, Ignition and Burn Propagation

As previously mentioned, the energy generated from fusion reactions can be re-deposited back into the plasma, heating it further and sustaining the high temperatures required for fusion reactions. However, not all fusion products can contribute to this heating. Charged particles, such as alpha particles, interact strongly with the plasma, and are able to re-deposit the bulk of their energy back into the plasma. Neutrons, on the other hand, and neutral and interact very weakly with the plasma, only through direct collisions with nuclei. As a result, their stopping distance is very long, and they often escape the plasma without depositing their energy. It is therefore important to distinguish between the total fusion energy generated, and the amount of energy generated that can be re-deposited back into the plasma. This redeposited energy is referred to as alpha heating, as it is primarily the alpha particles that contribute to this heating, although any charged fusion products will be included when considering this heating.

There are also many energy loss mechanisms that act to cool the plasma. These include radiative losses through bremstahlung radiation, conduction, and convection. The point at which the alpha heating overcomes these losses is referred to as ignition and is the point at which a thermal runaway instability occurs; as the plasma gets hotter, the fusion reaction rate increases, leading to more alpha heating, which heats the plasma further. This positive feedback loop leads to a rapid increase in the plasma temperature, and a self-sustaining burn wave propagates through the fuel.

2.5.5 Fast Ignition

A more niche area of fusion research is fast ignition (FI) [70]. In this approach, the implosion is separated into two stages. The first stage compresses the fuel isochorically to high densities, but without forming a hotspot, in contrast to the approximately isobaric hotspot ignition scheme desribed in Section 2.5.2. The

second stage then uses a short, high intensity laser pulse to generate a beam of energetic particles, which is focused onto the compressed fuel and collisionally deposits energy, igniting the fuel. Separating compression and heating allows the fuel to be compressed to higher densities without simultaneously satisfying the hotspot ignition conditions. Consequently, the required implosion velocity and compression energy are reduced, potentially enabling higher fusion gains. Furthermore, the compression and ignition phases are separated, and therefore both stages of the implosion can be optimised largely independently. However, this approach introduces additional complexity, requiring efficient generation, transport, and deposition of the energetic particles, as well as additional laser systems and asymmetric heating. Consequently, FI remains less experimentally mature than the conventional hotspot ignition approach.

2.5.6 Performance Metrics

There is a large amount of terminology used to describe the performance of an ICF implosion, with lots of overlap or amiguity in the definitions. This subsection aims to introduce these terms and clarify their definition as used in this thesis.

The Lawson Criterion and the n-T-τ Triple Product

The Lawson criterion [71] is a figure of merit that describes the conditions required for a fusion plasma to become self-sustaining. The criterion is derived by balancing the alpha heating generated by fusion reactions against the energy losses from the plasma. The rate of fusion reactions in a plasma consisting of a 50/50 mix of deuterium and tritium can be written as

$$R = n_D n_T \langle \sigma v \rangle = \frac{n^2}{4} \langle \sigma v \rangle, \tag{2.67}$$

where $n_D = n_T = n/2$ are the deuterium and tritium densities, n is the total ion density, and $\langle \sigma v \rangle$ is the Maxwellian averaged reaction rate of the fuel. The alpha heating power density generated by these reactions is then

$$P_\alpha = \frac{n^2}{4} \langle \sigma v \rangle E_\alpha, \tag{2.68}$$

where $E_\alpha = 3.5$ MeV is the energy of the alpha particles produced in DT reactions.

The internal energy of the plasma is given by the ideal gas law, $W = 3nk_BT$. By defining the energy confinement time, τ_E, as the time taken for the plasma to lose this internal energy, the average power loss from the plasma can be written as

$$P_{\text{loss}} = \frac{W}{\tau_E} = \frac{3nk_BT}{\tau_E}. \tag{2.69}$$

The Lawson criterion can then be expressed in terms of a n-T-τ triple product by balancing the alpha heating against the power losses, $P_\alpha \geq P_{\text{loss}}$, which gives

$$nT\tau_E \geq \frac{12k_BT^2}{E_\alpha \langle \sigma v \rangle}. \tag{2.70}$$

This expression defines the minimum conditions required for a plasma to become self-sustaining.

In ICF, this equation can be written in a more convenient form by expressing the confinement time as the time taken for an ion to traverse the plasma (radius R) at the ion sound speed. This allows the Lawson criterion to be written in terms of the areal density, ρR, as

$$\rho R \gtrsim 0.3 \text{ g/cm}^2, \tag{2.71}$$

which uses DT specific values for the constants and assumes a temperature of $T \approx 10$ keV. The derivation of this equation is given in Appendix A.5.

This value of 0.3 g/cm^2 refers to the minimum hotspot areal density required to confine the alpha particles within the hotspot and sustain self-heating. However, this is not sufficient to ensure efficient burn of the entire fuel assembly. For the thermonuclear burn wave to propagate from the hotspot into the surrounding dense fuel before the target hydrodynamically disassembles, a second criterion is imposed on the areal density of the compressed fuel. This "whole fuel" criterion requires a total fuel areal density of approximately 3 g/cm^2.

Burning Plasma and Ignition

An effective ICF implosion relies on significant alpha heating to drive fusion reactions. The first milestone that must be reached is a burning plasma, where the alpha heating is the dominant heating mechanism in the plasma compared to external drivers. This milestone was first achieved at the NIF in November 2020 [72] where approximately 10 kJ of $P\mathrm{d}V$ work was done on the hotspot, and around 20 kJ of alpha heating energy was deposited. A burning plasma is a necessary, but not sufficient, condition for the plasma to become self-sustaining. In a burning plasma, the plasma may still be cooling if the external heating and alpha heating combined are insufficient to overcome the energy losses.

The next milestone is ignition, which is the point at which the alpha heating alone overcomes the energy losses, and the plasma becomes self-sustaining. Once ignition is reached, the thermal runaway instability occurs, and a burn wave propagates through the fuel, rapidly increasing the temperature and fusion reaction rate. It is at this point that huge jumps in the amount of fusion energy generated occur. Ignition has also been acheived at the NIF, with the first instance occuring in August 2021 [73], with a fusion yield of 1.37 MJ (around $1370/5 = 274$ kJ of alpha heating energy deposited).

Yield and Gain

The fusion yield, Y, is the total fusion energy released in an implosion, while the gain, G, is the ratio of the fusion energy released to the energy input into the implosion. Several "gain" definitions coexist; they differ only in what is counted as input energy. To avoid ambiguity, this thesis will use the following:

- **Scientific (target) gain**, $G_{\mathrm{sc}} \equiv Y/E_{\mathrm{laser}}$: fusion energy out divided by the total laser energy on target.
- **Hotspot (ignition) gain**, $G_{\mathrm{hs}} \equiv Y/E_{\mathrm{hotspot}}$: fusion energy out divided by the energy deposited to form the hotspot.

- **Engineering (plant) gain**, $G_{\text{eng}} \equiv E_{\text{electric,out}}/E_{\text{electric,in}}$: wall-plug electrical gain for an end-to-end power system (includes electrical energy to laser energy efficiency).

In this thesis, if a gain value is given without a modifier, it refers to scientific gain, which is the most common metric used in the literacture.

Breakeven is defined as $G = 1$, where the fusion energy out equals the energy in. Scientific breakeven ($G_{\text{sc}} = 1$) is the most commonly used definition, and it was first achieved at the NIF in December 2022 [74], with a yield of 3.15 MJ from 2.05 MJ of laser energy ($G_{\text{sc}} \approx 1.54$). Since then, repeated ignition with higher G_{sc} have been acheived, with the current record yield being around 8.6 MJ from 2.08 MJ of laser energy ($G_{\text{sc}} \approx 4.13$) in April 2025 [75]. While there is an increasing trend in scientific gain, engineering gain remains far below breakeven due to the present wall-plug efficiency ($\sim 1-2\%$) of high-energy lasers, and the use of the indirect drive approach, which involves an additional laser-to-X-ray conversion step compared to direct drive.

Alternate Metrics

There are additional metrics that are sometimes used to describe the performance of an implosion. One of these is the areal density, ρR, which is a measure of the column density of the fuel. This metric is important as it describes how effectively the alpha particles are confined within the fuel and appears directly in the Lawson criterion (Equation 2.71). If the areal density is too low, the alpha particles escape the burning region without re-depositing their energy, reducing alpha heating and making ignition more difficult. This metric is also particularly important in fast ignition [70], which relies on heating the compressed fuel with a beam of energetic particles. Here, the areal density must be sufficiently high to absorb the energy deposited by the incoming particle beam and initiate burn. Unlike conventional hotspot ignition, however, the particle beam heats the compressed fuel anisotropically, depositing energy as it penetrates the target. Consequently, part of the heated region lies close to the outer surface of the compressed fuel, where it is not confined by dense surrounding material and can expand freely.

This hydrodynamic expansion removes thermal energy from the ignition region, introducing an additional hydrodynamic loss mechanism that is absent in the hotspot ignition scheme. Consequently, fast ignition requires a larger areal density than would be expected from alpha-particle confinement alone in order to compensate for these additional hydrodynamic losses.

A metric that areal density is often related to is the burn up fraction, which refers to the fraction of the fuel that undergoes fusion reactions during the implosion. This metric is usful as it indicates how effectively the burn wave propagates into the dense fuel. The burn up fraction, Φ, can be estimated from the areal density [76] as

$$\Phi \approx \frac{\rho R}{\rho R + H_B}, \tag{2.72}$$

where $H_B \approx 7$ g/cm^2 is the burn parameter for DT fuel.

2.6 Instabilities

Instabilities play a central role in determining the performance of ICF implosions. Small perturbations, whether arising from imperfections in the target or from non-uniformities in the laser drive, can grow exponentially during the compression process. Such growth can seed asymmetries [77], mix ablator material into the fuel [78], preheat the fuel [79], or redirect laser energy away from the capsule [80], all of which reduce the efficiency of the implosion and make ignition more difficult.

Broadly, instabilities in ICF fall into two categories: hydrodynamic and parametric. Hydrodynamic instabilities typically occur at material interfaces within the imploding capsule and are driven by acceleration or shock dynamics. Parametric instabilities, on the other hand, arise in the plasma surrounding the target when the intense laser light couples to plasma waves. The following subsections describe each class of instability, their underlying physical mechanisms, and their impact on implosion performance.

2.6.1 Hydrodynamic Instabilities

Hydrodynamic instabilities [81] arise from the motion of materials with different densities during an implosion. In ICF, they are most prominent at two key interfaces: between the ablator and the fuel during the compression phase, and between the hotspot and the surrounding dense fuel during burn. If a small perturbation exists at either interface, differences in pressure and acceleration can cause it to grow, leading to mixing, distortion of the fuel shell, and reduced confinement.

The most common hydrodynamic instability discussed in the context of ICF is the Rayleigh-Taylor (RT) instability, which occurs when a dense fluid is accelerated by a lighter one. During the acceleration phase of an ICF implosion, the ablator (denser material) is driven inward by the low-density plasma blowoff (lighter material), an RT-unstable configuration. The growth rate of small-amplitude perturbations in the classical incompressible case is given by

$$\gamma = \sqrt{Akg}, \tag{2.73}$$

where $A = (\rho_1 - \rho_2)/(\rho_1 + \rho_2)$ is the Atwood number, $k = 2\pi/\lambda$ is the perturbation wavenumber, and g is the effective acceleration. This expression shows that growth is faster for shorter-wavelength perturbations and for larger density contrasts.

In real ICF capsules, density gradients and the ablation of material from the interface introduces a stabilising effect. The ablative RT growth rate can be approximated by [55, 82]

$$\gamma = \sqrt{\frac{Akg}{1 + kL}} - \beta k V_a, \tag{2.74}$$

where $L = \min\left(\rho\left(\partial\rho/\partial r\right)\right)$ is the density scale length, V_a is the ablation velocity, and $\beta \approx 2$ is an empirical constant. The modification of the first term indicates that a finite scale length smooths the interface and reduces the effective Atwood number, thereby lowering the growth rate. The second term represents the damping caused by mass ablation that carries perturbations away from the interface. These effects therefore suppress short wavelength (large k) modes, though large-scale perturbations can still grow significantly. Thermal conduction and viscosity also

act to suppress RT growth, and can be incorporated into Equation 2.74 to form a more complete description of the RT growth rate.

Another related process is the Richtmyer-Meshkov (RM) instability, which occurs when a shock wave interacts with a perturbed surface. Unlike the continuously driven RT instability, the RM instability grows linearly with time rather than exponentially. In ICF, it can seed perturbations at material interfaces during the sequence of shocks used to compress the capsule, providing initial conditions for subsequent RT growth.

Finally, the Kelvin-Helmholtz (KH) instability occurs when a shear flow intersects a stratified fluid. In ICF, this becomes relevant in the evolution of RT instability bubbles, where fingers develop and create regions of high shear. KH instabilities can further enhance mixing and degrade implosion performance.

Hydrodynamic instabilities are seeded by many sources, including surface roughness on the capsule, imperfections in layer thickness, and non-uniformities in the laser drive.

The impact of these instabilities, particularly the RT instability, can be significant. Initially, instability seeds grow to cause asymmetries in the imploding shell. These asymmetries can lead to non-uniform compression of the fuel, and increase the surface area of the fuel, thereby increasing radiative losses. In severe cases, the instabilities can cause mixing of material from the shell into the fuel, adding cooler, higher-Z particles to the hotspot that again increases radiative losses.

There are many mitigation strategies that are currently being implemented or explored to reduce the growth of these instabilities. The first mitigation strategy is to reduce the initial perturbations within the capsule. Stringent tolerances are placed on the surface roughness of the capsule, with target fabrication now achieving surface perturbations below ~ 10 nm [83]. Next is to reduce the non-uniformity of the laser drive, which can be achieved through beam smoothing techniques such as phase plates [84], and smoothing by spectral dispersion (SSD) [85]. Laser pulses are typically freckled profiles, with high intensity speckles that can seed instabilities. By varying the spatial and temporal phase of the laser pulse, these speckles can be smoothed out, reducing the initial perturbations.

Beyond directly reducing the initial perturbations, the growth rate of the instabilities can also be reduced through design choices such as drive pulse shaping [86, 87], graded dopant layers [88], and the inclusion of low-density foam overcoats to thermally smooth energy transport [89, 90]. The drive pulse can be shaped to launch a decaying shock into the target. This increases the adiabat of the outer fuel and ablator material, while keeping the inner fuel in a low adiabat state. This improves the stability of the outer layers, while keeping the inner fuel compressible. Dopants are often used within the target to prevent preheat from energetic electrons or high frequency radiation. By grading the dopant concentration, sharp material interfaces can be smoothed out, reducing the growth of instabilities. Finally, low-density foam overcoats can be used to thermally smooth energy transport through the ablator, further reducing drive asymmetries.

Even with the implementation of such measures, RT-driven mixing remains one of the dominant limits on achievable convergence and fusion yield in current ICF experiments.

2.6.2 Parametric Instabilities

Parametric instabilities [50] occur when the incident laser light couples to plasma waves in the low-density coronal plasma surrounding the target. These interactions can scatter light away from the target and generate energetic electrons that preheat the fuel, both of which degrade implosion performance. They are driven by the oscillating electromagnetic field of the laser, which can parametrically excite waves in the plasma when specific frequency and wave-vector matching conditions are satisfied.

In general, a parametric instability occurs when an incident electromagnetic wave (frequency ω_0, wavevector $\mathbf{k}_0$) decays into two daughter waves with frequencies and wavevectors $(\omega_1, \mathbf{k}_1)$ and $(\omega_2, \mathbf{k}_2)$ that satisfy the resonance conditions

$$\omega_0 = \omega_1 + \omega_2, \qquad \mathbf{k}_0 = \mathbf{k}_1 + \mathbf{k}_2. \tag{2.75}$$

Depending on which plasma modes are involved, several specific instabilities can arise.

Stimulated Raman Scattering (SRS)

In SRS, the incident laser photon decays into a scattered photon and an electron plasma wave. This process occurs primarily near the quarter-critical surface, where the laser frequency is roughly twice the local plasma frequency. The scattered light propagates either backward toward the laser source (backward SRS) or outward at an angle (side-scatter SRS), while the generated plasma wave can accelerate electrons to tens or hundreds of keV. SRS therefore impacts an implosion in two ways: through the scattering of light away from the target, and through hot electron generation, which can preheat the DT fuel, increasing its adiabat and reducing compressibility.

Stimulated Brillouin Scattering (SBS)

SBS involves the decay of the laser photon into a scattered photon and an ion-acoustic wave. Because ion-acoustic waves are low-frequency modes, the frequency shift between the incident and scattered light is small, allowing SBS to occur up to the critical density surface. SBS primarily redirects laser energy away from the target, reducing the drive efficiency and introducing asymmetries in the drive.

Two-Plasmon Decay (TPD)

TPD becomes important around the quarter-critical surface, where one photon can decay into two electron plasma waves. Like SRS, this process produces energetic electrons that can preheat the fuel. TPD and SRS often coexist near the quarter-critical surface.

Cross-Beam Energy Transfer (CBET)

When multiple laser beams intersect in the coronal plasma, they can exchange energy through SBS. In direct drive implosions, this cross-beam energy transfer [80] redistributes laser power from incoming to outgoing pulses, reducing the laser energy incident on the target and reducing the drive symmetry. CBET has become one of the dominant degredation mechanisms, with the reduction of CBET being a key design goal for future laser facilities. However, in indirect drive, CBET

provides a mechanism to transfer energy between inner and outer beams, which can be used to tune the drive symmetry.

Mitigation Strategies

One of the main strategies for mitigating parametric instabilities on future laser facilities is the use of laser bandwidth broadening [91]. Bandwidth refers to the range of frequencies present in the laser pulse. Increasing the laser bandwidth breaks the strict frequency matching required for resonant coupling, reducing growth rates. By spreading the laser intensity over a range of frequencies, the effective intensity at any given frequency is reduced, mitigating the instability.

The use of laser 'zooming' techniques [92], where the focal spot size evolves during the drive, can also be used to mitigate parametric instabilities, especially CBET. Laser zooming reduces the intensity of the laser pulse in the coronal plasma, which reduces the growth rates of all parametric instabilities, and has the added effect of reducing the volume of plasma where the laser beams overlap, reducing the region where CBET can occur.

The combined effect of parametric instabilities is a reduction in laser-target coupling efficiency, drive symmetry, and fuel pre-heating, all of which directly impact implosion performance. As future laser facilities push to higher energies and intensities, understanding and mitigating these instabilities will be critical for achieving ignition and high gain.

2.7 Numerical Methods

The physics described in the preceding sections is inherently complex and highly non-linear, involving tightly coupled processes that span a wide range of spatial and temporal scales. Experimental investigation of this physics requires large-scale, high-energy laser facilities, which are costly to build and operate, and offer limited shot opportunities [47]. Moreover, even with advanced diagnostics, the amount of information that can be extracted from a single experiment is fundamentally constrained by the extreme conditions, short timescales, and small spatial scales

characteristic of ICF implosions. Many quantities of interest cannot be measured directly and must instead be inferred indirectly [93].

Numerical simulations therefore play a central role in ICF research [94]. Predictive computational models are used to design experiments, interpret diagnostic data, and explore new target and drive concepts prior to experimental validation. In modern ICF research, simulation is not merely supportive but essential to both physical understanding and experimental progress [95].

There are multiple numerical approaches to modelling ICF plasmas, each with distinct assumptions, strengths, and limitations. The appropriate framework depends on the dominant physical processes, the relevant spatial and temporal scales, and the required level of physical detail. Broadly, modelling approaches fall into two categories: fluid descriptions and kinetic descriptions.

2.7.1 Fluid Description

A fluid description of a plasma treats it as a continuous medium characterised by macroscopic quantities such as mass density ρ, bulk velocity $\mathbf{u}$, pressure P, and internal energy per unit mass ϵ. The plasma is described by a set of fluid equations, which are derived from the conservation of mass, momentum, and energy. Neglecting viscosity and conductivity, and assuming the material flow is non-relativistic, the governing equations reduce to the Euler equations:

$$\partial_t \rho + \nabla \cdot (\rho \mathbf{u}) = 0, \tag{2.76}$$

$$\partial_t (\rho \mathbf{u}) + \nabla \cdot (\rho \mathbf{u} \otimes \mathbf{u} + P\mathbf{I}) = 0, \tag{2.77}$$

$$\partial_t E + \nabla \cdot ((E + P)\mathbf{u}) = 0, \tag{2.78}$$

where

$$E = \rho\epsilon + \frac{1}{2}\rho \left|\mathbf{u}\right|^2, \tag{2.79}$$

is the total energy density of the fluid.

The Euler equations contain four unknown thermodynamic variables (ρ, $\mathbf{u}$, P, and ϵ, or equivalently E) but provide only three independent conservation laws. To close the system, an additional thermodynamic relation must be specified in the form

of an equation of state (EOS). The EOS relates the pressure to the thermodynamic state of the material and is commonly written as

$$P = H(\rho, \epsilon). \tag{2.80}$$

Together with the conservation equations, the EOS forms a closed system that uniquely determines the evolution of the material state.

For many problems in high-energy-density physics, and particularly in inertial confinement fusion, it is advantageous to adopt a Lagrangian formulation in which the computational mesh moves with the fluid. Because the mesh deforms with the material, regions undergoing compression are automatically refined, allowing steep gradients and shock fronts to be resolved with high resolution. In this frame, the time derivative follows a material element and is expressed using the material derivative

$$\frac{\mathrm{d}}{\mathrm{d}t} = \frac{\partial}{\partial t} + \mathbf{u} \cdot \nabla. \tag{2.81}$$

The conservation equations can then be rewritten in a form that directly describes the evolution of fluid parcels. In one spatial dimension, as considered in the simulations presented in this thesis, the equations reduce to

$$\frac{\mathrm{d}\rho}{\mathrm{d}t} + \rho \frac{\partial u}{\partial x} = 0, \tag{2.82}$$

$$\frac{\mathrm{d}u}{\mathrm{d}t} + \frac{1}{\rho} \frac{\partial P}{\partial x} = 0, \tag{2.83}$$

$$\frac{\mathrm{d}\epsilon}{\mathrm{d}t} + \frac{P}{\rho} \frac{\partial u}{\partial x} = 0 \tag{2.84}$$

supplemented by the equation of state $P = H(\rho, \epsilon)$. In this representation, the mass of each computational cell remains constant, and compression or expansion is captured through the motion of cell boundaries.

To evolve this system numerically, the spatial domain is discretised into a set of Lagrangian cells whose boundaries move with the fluid. The discretisation and evolution of the Lagrangian equations closely follows the treatment presented in Colvin and Larson's '*Extreme Physics*' [55]. Scalar quantities such as density ρ, internal energy ϵ, and pressure P are defined at cell centres, denoted by half-integer spatial indices $(j + \frac{1}{2})$, while velocities u are defined at cell boundaries,

denoted by integer indices j. A staggered leapfrog scheme is used in which velocities are defined at half-integer timesteps ($n \pm \frac{1}{2}$), while thermodynamic quantities are defined at integer timesteps n. Considering some initial known state where the cell-centred quantities $\rho^n_{j+\frac{1}{2}}$ and $\epsilon^n_{j+\frac{1}{2}}$, together with boundary velocities $u_j^{n-\frac{1}{2}}$, are known. The equation of state is first used to determine the pressure $P^n_{j+\frac{1}{2}}$ corresponding to the current material state.

The boundary velocities are updated by discretising the momentum equation (Equation 2.83), which states that fluid acceleration is driven by pressure gradients, using a finite difference approximation

$$u_j^{n+\frac{1}{2}} = u_j^{n-\frac{1}{2}} - \frac{P^n_{j+\frac{1}{2}} - P^n_{j-\frac{1}{2}}}{\frac{1}{2}\left(\rho^n_{j+\frac{1}{2}} + \rho^n_{j-\frac{1}{2}}\right)} \frac{1}{x^n_{j+\frac{1}{2}} - x^n_{j-\frac{1}{2}}} \Delta t^n. \tag{2.85}$$

Once the updated boundary velocities $u_j^{n+\frac{1}{2}}$ have been determined, the cell boundary positions are advanced in time according to

$$x_j^{n+\frac{1}{2}} = x_j^{n-\frac{1}{2}} + u_j^{n+\frac{1}{2}} \Delta t^{n+\frac{1}{2}}. \tag{2.86}$$

In the Lagrangian formulation, the mass contained within each cell remains constant throughout the simulation. Changes in density therefore arise solely from changes in cell volume. Denoting the constant cell mass by $m_{j+\frac{1}{2}}$, the density at the next timestep is given by

$$\rho^{n+1}_{j+\frac{1}{2}} = \frac{m_{j+\frac{1}{2}}}{x_{j+1}^{n+\frac{1}{2}} - x_j^{n+\frac{1}{2}}}. \tag{2.87}$$

The internal energy is evolved using the Lagrangian energy equation (Equation 2.84), in which changes arise from compressional work

$$\epsilon^{n+1}_{j+\frac{1}{2}} = \epsilon^n_{j+\frac{1}{2}} - \frac{P^n_{j+\frac{1}{2}}}{\rho^n_{j+\frac{1}{2}}} \frac{u_{j+1}^{n+\frac{1}{2}} - u_j^{n+\frac{1}{2}}}{x_{j+1}^{n+\frac{1}{2}} - x_j^{n+\frac{1}{2}}} \Delta t^{n+\frac{1}{2}}. \tag{2.88}$$

This update accounts for the work done by pressure forces during compression or expansion of the cell. The hydrodynamic update cycle for a single timestep can

then be completed by using the updated density and internal energy to determine the pressure at the next timestep using the equation of state

$$P_{j+\frac{1}{2}}^{n+1} = H(\rho_{j+\frac{1}{2}}^{n+1}, \epsilon_{j+\frac{1}{2}}^{n+1}). \tag{2.89}$$

In practical inertial confinement fusion simulations, the idealised Euler system described above is extended to include additional physics relevant to high-energy-density plasmas. Radiation-hydrodynamic codes incorporate source terms representing electron thermal conduction, radiation transport (often in a multi-group diffusion approximation) [59, 96], laser energy deposition, artificial viscosity [97], and collisional energy exchange between electrons and ions. Separate energy equations are typically evolved for electrons and ions, allowing for non-equilibrium temperatures ($T_e \neq T_i$) and energy equilibration through collisional coupling terms. The radiation-hydrodynamic codes used during the DPhil, HYADES [98] and HELIOS [99], employ a one-dimensional Lagrangian framework of this form, with the hydrodynamic update described above forming the core of the timestep cycle.

The equation of state used in these codes is vital to the accurate evolution of the system, and remains an area of active research. Often simulations rely on tabulated equation of state tables, constructed from a combination of theoretical models and experimental measurements. During a simulation, interpolation within these tables is used to determine the material state at each timestep. The experiment discussed in Chapter 7 of this thesis addresses the scarcity of experimental data for low-density foams under ICF relevant conditions and provides new data to constrain the EOS of these materials.

2.7.2 Kinetic Description

A fluid description of a plasma is valid when the mean free path of particles is much smaller than the lengthscales of interest and the plasma is in local thermodynamic equilibrium. However, many plasma physics phenomena occur outside of these limits, such as laser-plasma interactions and instabilities.

To capture these effects, a kinetic description of the plasma is required [100]. In kinetic theory, the state of each species is described not by macroscopic fluid moments, but by a distribution function $f_s(\mathbf{x}, \mathbf{p}, t)$. This function gives the number of particles of species s within a differential phase-space volume $\mathrm{d}^3\mathbf{x}\,\mathrm{d}^3\mathbf{p}$ centred at position $\mathbf{x}$ and momentum $\mathbf{p}$:

$$f_s(\mathbf{x}, \mathbf{p}, t)\,\mathrm{d}^3\mathbf{x}\,\mathrm{d}^3\mathbf{p}. \tag{2.90}$$

Kinetic simulations are substantially more computationally demanding than fluid models. In addition to evolving quantities in phase space rather than configuration space alone, kinetic methods must resolve the fastest microscopic timescales present in the plasma. Explicit electromagnetic kinetic codes must resolve electron plasma oscillations and electromagnetic wave propagation, requiring timesteps smaller than the inverse plasma frequency. Spatial resolution must also be sufficient to resolve the Debye length in order to avoid unphysical numerical heating. As a result, fully kinetic simulations are significantly more expensive than hydrodynamic simulations, and are often restricted to limited spatial domains, shorter timescales, and reduced dimensionality.

While kinetic codes are not tractable over full implosion timescales (~ 10 ns), they are essential for understanding the underlying physics, which occur on much shorter timescales (< 10 ps) and smaller spatial scales (< 10 μm). Kinetic simulations are therefore used to study the initial stages of laser energy deposition, the growth of parametric instabilities, and the generation of hot electrons, all of which have a significant impact on implosion performance. In this thesis, these codes are used to study the non-linear evolution of beam-plasma systems, as discussed in Chpaters 5 and 6.

There are two main approaches to numerically evolving the distribution function: Particle-in-Cell (PIC) and Vlasov-Maxwell simulations.

Particle-in-Cell Simulations

PIC simulations [101] represent the distribution function using a finite number of macro-particles, each representing a large number of physical particles. These macro-particles are evolved in a continuous phase space according to the Lorentz force

$$\frac{\mathrm{d}\mathbf{p}}{\mathrm{d}t} = q_s \left(\mathbf{E} + \frac{\mathbf{p}}{\gamma m_s} \times \mathbf{B} \right). \tag{2.91}$$

The electromagnetic fields $\mathbf{E}$ and $\mathbf{B}$ are defined on a spatial grid and are updated self-consistently using Maxwell's equations, with charge and current densities computed from the macro-particles trajectories.

The macro-particles contribute to electromagnetic fields defined on a spatial grid by depositing their charge and current onto the grid points using a weighting scheme. The fields are then updated using Maxwell's equations, and the updated fields are interpolated back to the macro-particle positions to compute the Lorentz force acting on them for the next time step. This cycle of particle pushing, charge deposition, field solving, and field interpolation is repeated to evolve the system in time.

PIC simulations are widely used due to their flexibility and ability to capture complex and non-linear plasma dynamics. However, they can suffer from numerical noise due to the finite number of macro-particles, which can leave fine-scale features in the distribution function unresolved, especially in low-density regions of phase space, which will be poorly sampled. The simulation codes `EPOCH` [102] and `SMILEI` [103] are PIC codes used during the DPhil, which are both explicit, relativistic, electromagnetic PIC codes.

Vlasov-Maxwell Simulations

Vlasov-Maxwell simulations [104] take a different approach by discretising the distribution function directly on a fixed phase-space mesh. The distribution function is represented on a grid in both position and momentum space, allowing for high-resolution capture of fine-scale structures in phase space without the statistical noise inherent in PIC methods.

The time evolution of the distribution function is governed by the collisionless Vlasov equation, which governs the conservation of particles in phase space:

$$\frac{\mathrm{d}f_s}{\mathrm{d}t} = \frac{\partial f_s}{\partial t} + \frac{\mathbf{p}}{\gamma m_s} \cdot \nabla_{\mathbf{x}} f_s + q_s \left(\mathbf{E} + \frac{\mathbf{p}}{\gamma m_s} \times \mathbf{B} \right) \cdot \nabla_{\mathbf{p}} f_s = 0, \tag{2.92}$$

where q_s and m_s are the charge and mass of a particle of species s, $\gamma = \sqrt{1 + (|\mathbf{p}|/mc)^2}$ is the Lorentz factor, and $\mathbf{E}$ and $\mathbf{B}$ are the electromagnetic fields. The electric and magnetic fields evolve according to Maxwell's equations, coupled self-consistently through the charge and current densities above.

While Vlasov-Maxwell simulations avoid the statistical noise of PIC methods, they are computationally intensive due to the high dimensionality of phase space. The need to resolve both spatial and momentum dimensions requires significant memory and processing power, especially in higher-dimensional simulations.

Together, the Vlasov and Maxwell equations define the Vlasov-Maxwell system, a fully kinetic model of collisionless plasma dynamics. This is the system of equations numerically solved by the one-dimensional Vlasov-Maxwell simulation code `VALIS` [105], which is utilised in this thesis.

3

Optimisation Approaches to Inertial Confinement Fusion Implosions

Contents

3.1 Introduction 57
3.2 General Formulation of Optimisation 59
3.3 Bayesian Optimisation using Gaussian Processes . . . 61
3.3.1 Bayesian Inference (Bayes' Theorem) 62
3.3.2 Gaussian Processes (GPs) 63
3.3.3 Conditioning the GP on Data 71
3.3.4 Acquisition Functions 72
3.3.5 Bayesian Optimisation 73
3.4 Bayesian Package Implementation 75
3.4.1 Computational Complexity 79
3.4.2 Additional Features 81
3.5 Particle Swarm Optimisation (PSO) 90
3.6 PSO Package Implementation 92
3.6.1 Additional Features 93
3.7 Summary 97

3.1 Introduction

Inertial confinement fusion implosions are highly sensitive to the precise configuration of the implosion, including the target design and laser pulse profile. This has been demonstrated experimentally on the NIF, in which large variations

to the performance of an implosion are observed from small shifts in configuration [106]. Optimisation therefore plays a crucial role in the development of high-performance ICF designs.

Optimising ICF target designs is inherently challenging. The relevant design space is typically high-dimensional, encompassing laser pulse timings and powers, material properties, and capsule dimensions, all while satisfying constraint conditions. Moreover, implosion dynamics are strongly nonlinear, and each function evaluation requires performing computationally expensive simulations. Performance metrics such as shock timing, fuel areal density (ρR), and gain are often noisy and can exhibit multiple physically distinct local optima associated with different implosion pathways. Consequently, simple strategies such as random sampling or sequential linear scans are inefficient and often fail to capture strong parameter coupling or identify globally optimal configurations.

In recent years, advanced optimisation methodologies have begun to attract increasing attention within the ICF community. Approaches including population-based algorithms [107], Bayesian methods [108, 109], genetic algorithms [110], and machine-learning-based approaches [111] have been explored as alternatives to traditional scan-based optimisation. These studies demonstrate that modern optimisation frameworks can more effectively navigate high-dimensional, nonlinear design landscapes. However, comparisons between different optimisation approaches remain scarce, especially under experimentally relevant constraint handling and across distinct implosion regimes. As a result, it is not yet clear which classes of optimisation methods are most appropriate for specific ICF design problems.

The central contribution of this chapter is the development of two dedicated optimisation frameworks for inertial confinement fusion design: Gaussian process-based Bayesian optimisation and particle swarm optimisation (PSO). Both are implemented in Python and specifically adapted for high-cost radiation-hydrodynamic simulations. The frameworks are designed to enable efficient global exploration of complex, noisy, and multi-modal design spaces, while accommodating the

parallel execution potential of large-scale simulation campaigns on high-performance computing (HPC) facilities.

This capability is achieved through the integration of enhancements tailored to the demands of ICF optimisation. While the individual components are drawn from the broader optimisation literature, their curation and integration into a unified optimisation package forms a key contribution of this thesis. The resulting packages form general, modular, and extensible optimisation platforms for high-energy-density physics applications.

This chapter presents the theoretical foundations, implementation, and benchmarking of the two optimisation frameworks. Section 3.2 formulates the general optimisation problem and outlines the challenges inherent to high-cost black-box optimisation. Section 3.3 introduces Bayesian optimisation, including Bayesian inference and Gaussian process surrogate modelling, which form the basis of the framework. Section 3.4 then details the practical implementation and the additional algorithmic enhancements incorporated to augment the approach.

Section 3.5 introduces particle swarm optimisation (PSO), presenting the mathematical formulation of the swarm evolution and the underlying search dynamics. Finally, Section 3.6 describes the PSO implementation and the extensions developed to adapt the method to ICF simulation campaigns.

3.2 General Formulation of Optimisation

An optimisation problem can be posed in its most general form as the task of finding an optimum state, $\mathbf{x}^*$, within some domain, $\mathcal{X}$, which maximises a function, $f : \mathcal{X} \mapsto \mathbb{R}$ [112]. This can be more compactly written as:

$$\text{find} \quad \mathbf{x}^* \in \mathcal{X} \quad \text{such that} \quad f(\mathbf{x}^*) = \max_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x}) \,. \tag{3.1}$$

Although Equation 3.1 formulates the optimisation as a maximisation problem, it can be easily converted to a minimisation problem by considering the negative of the function, i.e. $f(\mathbf{x}) \mapsto -f(\mathbf{x})$.

Due to the ubiquity of optimisation, there is no shortage of approaches to these problems. Each approach is often best suited to specific optimisation tasks where certain knowledge or assumptions can be made about the underlying function, $f(\mathbf{x})$, or the domain on which the function is defined, $\mathcal{X}$. The simplest approaches involve randomly sampling or grid searching the domain. These methods assume no prior knowledge of the optimisation task and are not influenced by previous samples. While naive in nature, they can still be useful in certain scenarios, particularly in low-dimensional problems or when performing hyper-parameter tuning, where exhaustive search may be computationally feasible and performance is less sensitive to parameter choice.

A slight improvement over these basic methods can be found in space-filling designs, which aim to cover the domain more evenly. One such method is Latin Hypercube Sampling (LHS) [113], which stratifies the sampling space along each dimension to ensure better uniformity, requiring fewer samples than randomly searching. While this leads to more representative coverage of the domain, LHS still does not incorporate feedback from the objective function and remains inefficient in high-dimensional or complex landscapes.

When the structure of the objective function is known (or assumed), regression-based approaches [114] such as linear regression, polynomial regression, and generalised linear models can be highly effective. These approaches assume the function has a parametric form and estimate the coefficients associated with this form to best fit the observed data. An alternative set of approaches consists of gradient-based methods, such as gradient descent [115], Newton's method [116], and quasi-Newton methods. These approaches leverage the derivatives of the function to better search for the optimum. These methods are often highly effective and can converge quickly, but require the function to satisfy certain criteria, such as smoothness and differentiability, and perform best on convex objective functions. However, they come with significant limitations: they are sensitive to the choice of initial conditions, can become trapped in local optima, and often perform poorly on noisy, non-smooth, or discontinuous functions. Moreover, in high-dimensional spaces, the computational

cost of evaluating gradients can be substantial. Issues such as vanishing or exploding gradients, saddle points, and poor conditioning can further complicate convergence.

Many practical applications typically fall outside of this regime [117]. The function to be optimised is often expensive to evaluate, noisy, and lacks a known parametric form. These functions are commonly referred to as black-box functions [118], where the internal structure is unknown, and optimisation must proceed based solely on input-output evaluations.

In ICF, each evaluation may require hours of HPC time, making sample efficiency critical. The performance landscape is frequently high-dimensional, nonlinear, and multi-modal. Furthermore, ignition-relevant metrics can exhibit threshold behaviour, where small parameter variations produce discontinuous or cliff-edge changes in performance. Numerical noise, resolution effects, and complex parameter coupling further complicate optimisation. These problems are exacerbated in experimental settings, where measurement noise, shot-to-shot variability, and limited experimental runs introduce additional challenges.

These characteristics necessitate the use of optimisation strategies that are both sample-efficient and robust in challenging parameter spaces. Two complementary approaches for tackling black-box optimisation problems are Bayesian Optimisation [112] and Particle Swarm Optimisation (PSO) [119].

3.3 Bayesian Optimisation using Gaussian Processes

Bayesian optimisation using Gaussian processes (GPs) [112, 120] takes a statistical approach to optimisation. An initial GP, which will be described in detail in Section 3.3.2, is used to define a probabilty distribution over an infinite set of possible functions that could represent the true objective function. This initial probability distribution is known as the prior, and each possible function (infinite in this case) is considered a hypothesis, $\mathcal{H} = \{\mathcal{H}_0, \mathcal{H}_1, \dots\}$. As data becomes avaliable, Bayesian inference (detailed in Section 3.3.1) is used to update the probability of each hypothesis being the true function, forming the posterior distribution

(Section 3.3.3). An acquisition function (Section 3.3.4) is then used to determine the next most informative sample to evaluate, balancing exploration and exploitation. This process is iteratively repeated until a stopping criterion is met, such as reaching a maximum number of evaluations or achieving convergence. This Section will explain in detail each of these steps.

3.3.1 Bayesian Inference (Bayes' Theorem)

Bayes theorem is the method by which the probabilities after observing data are updated and is given by [121]

$$P(\mathcal{H}|\mathcal{D}) = \frac{P(\mathcal{D}|\mathcal{H})P(\mathcal{H})}{P(\mathcal{D})}. \tag{3.2}$$

In the context of updating the probability of hypotheses, these terms can be considered in the following way

- $\mathbf{P}(\mathcal{H}|\mathcal{D})$: The probability of hypothesis $\mathcal{H}$, given the dataset $\mathcal{D}$ (**posterior**).
- $\mathbf{P}(\mathcal{D}|\mathcal{H})$: The probability of observing the dataset $\mathcal{D}$, given hypothesis $\mathcal{H}$ (**likelihood**).
- $P(\mathcal{H})$: The probability of hypothesis $\mathcal{H}$ before observing any data (**prior**).
- $\mathbf{P}(\mathcal{D})$: A normalisation term ensuring that the probability of the hypotheses given the data sum (or integrate, if considering a continuous hypothesis space) to 1. This term is often treated as an arbitrary constant in Bayesian inference, as it doesn't depend on the hypothesis but just ensures the result is a valid probability distribution.

$$P(\mathcal{D}) = \sum_{\mathcal{H}} P(\mathcal{D}|\mathcal{H})\, P(\mathcal{H}) \;\text{ or }\; \int P(\mathcal{D}|\mathcal{H})\, P(\mathcal{H})\, \mathrm{d}\mathcal{H}. \tag{3.3}$$

Bayesian inference can be intuitively understood by considering a set of potential functions (hypotheses) over a space. Initially, each function may be equally plausible, but as new data is observed, certain functions become more likely, while others become less so. In this sense, Bayesian inference is a formal, quantitative approach

that updates these beliefs about which functions are more likely to explain the observed data. The process of updating beliefs is driven by the evidence (data) and the prior assumptions (prior distribution), with the result being a refined probability distribution over the functions (posterior) to determine which best fits the data. This continuous hypothesis space motivates the use of GPs, which provide a flexible, non-parametric, and analytically tractable framework for modelling distributions over functions.

3.3.2 Gaussian Processes (GPs)

Univariate Gaussian Distributions

Before introducing Gaussian processes, it is helpful to first build up the idea of Gaussian distributions. A Gaussian probability density function (PDF), also known as the normal distribution, or a bell curve, referencing its shape, has the functional form

$$p(x) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right), \tag{3.4}$$

where the prefactor is a normalisation constant which ensures the total probability integrates to 1.

A Gaussian distribution is fully defined by its mean μ and standard deviation σ. The mean determines the central location, while the standard deviation controls the spread of the distribution. A larger σ results in a wider and flatter distribution, while a smaller σ leads to a more peaked curve. Figure 3.1 (a) demonstrates four example Gaussian distributions for a range of μ and σ values. Since the Gaussian distribution only depends on two values, μ and σ, it is convention to denote a variable x sampled from a Gaussian distribution as $x \sim \mathcal{N}(\mu, \sigma^2)$.

Multivariate Gaussian Distributions

The Gaussian distribution can be extended from a single variable (univariate case) to multiple variables, forming a multivariate Gaussian distribution that describes the joint behaviour of multiple correlated variables. Consider an n-dimensional

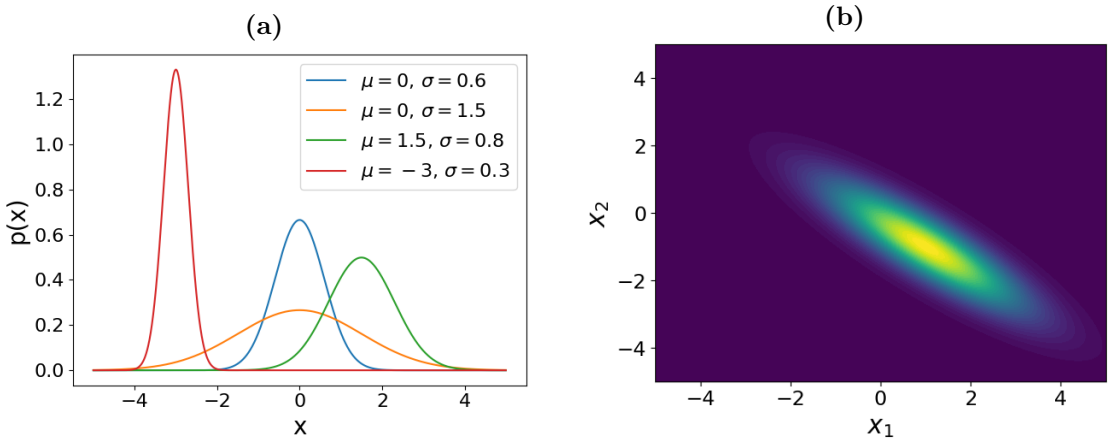


Figure 3.1: (a) Univariate Gaussian distributions for different values of μ and σ. (b) A bivariate (2D) Gaussian distribution, where the shape is influenced by both the individual variances and the correlation between the two variables.

vector, $\mathbf{z} = (z_1, z_2, ..., z_n)$ where each z_i is independently Gaussian-distributed with a mean $\mu = 0$, and unit variance $\sigma^2 = 1$, referred to as a standard normal distribution. Each variable can be expressed as:

$$z_i \sim \mathcal{N}(0, 1). \tag{3.5}$$

Since these are independent variables, the total probability function is the product of each independent probability function

$$\begin{aligned} p(\mathbf{z}) = \prod_{i=1}^{n} p(z_i) = \prod_{i=1}^{n} \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{(z_i)^2}{2}\right) &= \frac{1}{(2\pi)^{n/2}} \exp\left(-\frac{1}{2}\sum_{i=1}^{n} z_i^2\right) \\ &= \frac{1}{(2\pi)^{n/2}} \exp\left(-\frac{1}{2}\mathbf{z}^T\mathbf{z}\right). \end{aligned} \tag{3.6}$$

A more general multivariate Gaussian distribution can be defined by allowing for a nonzero mean vector, $\mu = (\mu_1, \mu_2, ..., \mu_n)$, and a covariance matrix Σ, which captures the variances of each variable along the diagonal ($\Sigma_{ii} = \sigma_i^2$) and the covariances between variables in the off-diagonal terms ($\Sigma_{ij} = \text{Cov}(z_i, z_j)$). A positive off-diagonal entry indicates that two variables tend to increase together, while a negative entry means they are inversely related. Since covariance is symmetric by definition, the covariance matrix Σ is always symmetric, meaning

$\Sigma_{ij} = \Sigma_{ji}$ for all indices i, j

$$\Sigma_{ij} = \begin{pmatrix} \sigma_1^2 & \mathrm{Cov}(z_1, z_2) & \cdots & \mathrm{Cov}(z_1, z_d) \\ \mathrm{Cov}(z_2, z_1) & \sigma_2^2 & \cdots & \mathrm{Cov}(z_2, z_d) \\ \vdots & \vdots & \ddots & \vdots \\ \mathrm{Cov}(z_d, z_1) & \mathrm{Cov}(z_d, z_2) & \cdots & \sigma_d^2 \end{pmatrix}. \tag{3.7}$$

The PDF of the dependent multivariate Gaussian distribution is then given by

$$p(\mathbf{z}) = \frac{1}{(2\pi)^{n/2} |\Sigma|^{1/2}} \exp\left(-\frac{1}{2} (\mathbf{z} - \mu)^T \Sigma^{-1} (\mathbf{z} - \mu)\right), \tag{3.8}$$

which is derived in Appendix B.1 by transforming a standard normal distribution into a general multivarite Gaussian distribution with mean μ and covariance Σ.

An example of a two-dimensional bivariate Gaussian distribution with

$$\mu = \begin{pmatrix} 1 \\ -1 \end{pmatrix} \text{ and } \Sigma = \begin{pmatrix} 2 & -1.5 \\ -1.5 & 1.5 \end{pmatrix}, \tag{3.9}$$

is shown in Figure 3.1 (b). The different variances along each axis ($\Sigma_{11} = 2$, $\Sigma_{22} = 1.5$) cause the distribution to be stretched more along the first dimension and compressed in the second. Additionally, the negative off-diagonal covariance (-1.5) introduces a negative correlation between the two variables, meaning that as one increases, the other tends to decrease. These combined effects lead to an elliptical shape that is tilted downward in the direction of the negative correlation. The principal axes of this ellipse are defined by the eigenvectors of Σ, while the eigenvalues control the lengths of these axes, influencing the overall spread of the distribution.

Gaussian Processes (GPs)

Gaussian Processes (GPs) can be thought of as a continuous extension of a multivariate Gaussian distribution [120]. A multivariate Gaussian describes a probability distribution over function values at a finite set of points. To illustrate this, consider a set of discrete points in the parameter space:

$$\mathbf{X} = (\mathbf{x}_1, \mathbf{x}_2, ..., \mathbf{x}_n), \tag{3.10}$$

where each $\mathbf{x}_i \in \mathbb{R}^d$ is a point in a d-dimensional parameter space. The corresponding function values are assumed to be jointly distributed as a multivariate Gaussian:

$$\mathbf{Y} = (f(\mathbf{x}_1), f(\mathbf{x}_2), ..., f(\mathbf{x}_n)) \sim \mathcal{N}(\mu, \Sigma). \tag{3.11}$$

A Gaussian Process extends this idea to a continuous space, defining a distribution over functions rather than just a finite set of function values. Instead of defining the mean and covariance at a discrete set of points, GPs define them as functions over the entire parameter space. Therefore, the finite mean vector μ is replaced by a mean function $m(\mathbf{x})$, and covariance matrix Σ is generalised by a kernel function $k(\mathbf{x}, \mathbf{x}')$, which captures the correlation between any two input points. Thus, a Gaussian Process is formally written as:

$$f(\mathbf{x}) \sim \mathcal{GP}(m(\mathbf{x}), k(\mathbf{x}, \mathbf{x}')). \tag{3.12}$$

This means that for any finite selection of input points $\mathbf{X}$, the corresponding function values $\mathbf{Y} = \{f(\mathbf{x}_1), f(\mathbf{x}_2), \ldots, f(\mathbf{x}_n)\}$ follow a multivariate Gaussian distribution:

$$\mathbf{Y} \sim \mathcal{N}(\mathbf{m}, K), \tag{3.13}$$

where the mean vector $\mathbf{m}$ and covariance matrix K are obtained by evaluating $m(\mathbf{x})$ and $k(\mathbf{x}, \mathbf{x}')$ at the chosen points, $\mathbf{X}$.

Rather than considering just a single function, a GP represents a probability distribution over an infinite set of possible functions. Each sample drawn from a GP is a single realisation of a stochastic process, where the function values are correlated according to the kernel function.

Kernel/Covariance Function

The kernel function (also known as the covariance function) is central to GPs, as it determines how function values at different points in the parameter space are correlated. Through the kernel, prior assumptions about the function (such as its smoothness, periodicity, or complexity) are encoded into the model. Choosing an appropriate kernel is therefore crucial, as it shapes the GPs behaviour and allows domain knowledge to guide the optimisation process.

Covariance is closely related to correlation, which measures the strength of a linear relationship between two variables and is bounded betweem $[-1, 1]$, where ± 1 represents a perfect linear relationship. The correlation between two points is given by:

$$\text{correlation}(\mathbf{x}_1, \mathbf{x}_2) = \frac{K_{12}}{\sqrt{K_{11}K_{22}}}, \tag{3.14}$$

where K_{12} is the covariance between $\mathbf{x}_1$ and $\mathbf{x}_2$, and K_{11}, K_{22} are the variances at those points.

Different kernels encode different assumptions about the smoothness and periodicity of the function being modelled. A common component that appears in the kernel functions is a standardised distance between two points in the parameter space. This often takes the form

$$d_p(\mathbf{x}_1, \mathbf{x}_2) = \left(\sum_{i=1}^{d} \left| \frac{\mathbf{x}_{1,i} - \mathbf{x}_{2,i}}{l_i} \right|^p \right)^{\frac{1}{p}}, \tag{3.15}$$

where l_i denotes the lengthscale associated with dimension i, which scales an anisotropic parameter space according to the sensitivity of the objective function to each dimension. In the case of $p = 2$, this is the standard Euclidean distance, and $p = 1$ gives the Manhattan distance. Some commonly used kernels are:

- The **radial basis function (RBF)** kernel, which assumes the function is infinitely smooth and is given by:

$$k(\mathbf{x}_1, \mathbf{x}_2) = \exp\left(-\frac{1}{2} d_2(\mathbf{x}_1, \mathbf{x}_2)^2 \right), \tag{3.16}$$

  with the d_2 distance including the lengthscales which control how quickly correlations decay with distance along each axes. A smaller l_i allows the function to change more rapidly, while a larger l_i enforces smoother variations.

- The **Matérn** kernel, which is a generalisation of the RBF kernel and introduces a smoothness parameter ν. The most common choices of ν are 0.5, producing non-differentiable functions, 1.5, producing once-differentiable functions, and 2.5, producing twice-differentiable functions. As $\nu \to \infty$, the Matérn kernel

converges to the RBF kernel, making the function infinitely differentiable. The Matérn kernel is given by:

$$k_\nu(\mathbf{x}_1, \mathbf{x}_2) = \frac{1}{\Gamma(\nu)2^{\nu-1}} \left(\sqrt{2\nu} d_1(\mathbf{x}_1, \mathbf{x}_2)\right)^\nu K_\nu \left(\sqrt{2\nu} d_1(\mathbf{x}_1, \mathbf{x}_2)\right), \tag{3.17}$$

where K_ν is a modified Bessel function and Γ is a continuous extension of the factorial.

- A **periodic kernel** can be used to model periodic functions. A common choice is the decaying periodic kernel, which combines axis-aligned periodicity with an exponential decay

$$k(\mathbf{x}_1, \mathbf{x}_2) = \exp\left(-2\sum_{i=1}^{d} \frac{1}{l_i} \sin^2\left(\frac{\pi\left(\mathbf{x}_{1,i} - \mathbf{x}_{2,i}\right)}{p_i}\right)\right) \cdot \exp\left(-\sum_{i=1}^{d} \frac{|\mathbf{x}_{1,i} - \mathbf{x}_{2,i}|}{\sigma_i}\right), \tag{3.18}$$

where p_i is the period along dimension i, l_i is the lengthscale that controls the width of an individual peak, and σ_i governs the rate of decay in each dimension. The exponential decay term, σ_i, acts as a damping envelope that ensures correlations diminish with distance, preventing overly strong periodic correlations at far-separated points. This kernel is useful in modelling periodic or seasonal patterns, such as daily temperature cycles, economic trends, or periodic oscillations in physical systems.

These kernels are illustrated in figure 3.2 panel (a) in a 1D parameter space, for $x_1 = 0.3$. Each kernel encodes different structural assumptions about the function being modelled, with the RBF (blue) and Matérn (orange) kernels indicating that points closer together in the input space are more strongly correlated, while the decaying periodic kernel (green) enforces that points separated by an integer multiple of the period are correlated, with the strength of this correlation decaying with distance. The choice of kernel depends on domain knowledge and the nature of the function being approximated. In practice, these kernels are often combined additively or multiplicatively to capture more complex behaviour.

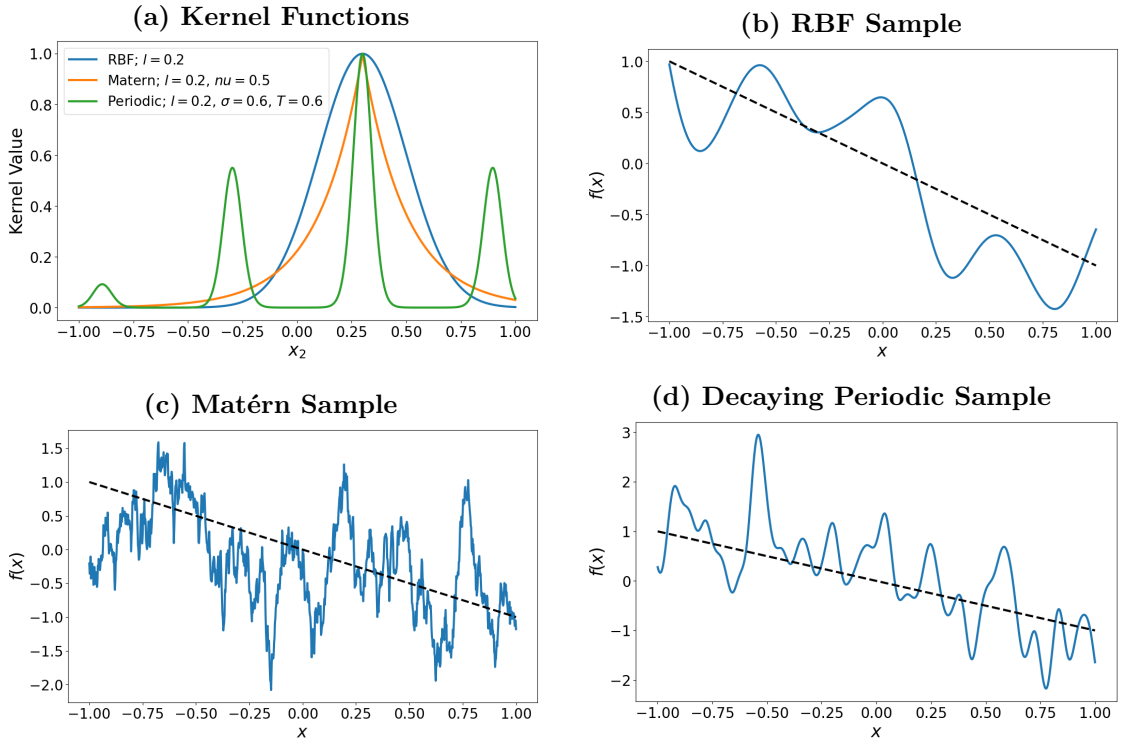


Figure 3.2: (a) Example covariance functions for the Radial Basis Function (RBF), Matérn, and decaying periodic kernels, illustrating how points are correlated with $x_1 = 0.3$. Panels (b)-(d) show representative samples drawn from Gaussian processes with the RBF, Matérn, and decaying periodic kernels, respectively, highlighting the types of functions considered plausible under each prior.

Sampling/Visualising GPs

While the previous sections described the mathematical structure of GPs, it is equally important to build an intuition for how GPs behave. This is best done by visualising sampled functions drawn from a GP. Although a GP defines a distribution over continuous functions, practical sampling involves evaluating the GP at a finite set of input points, yielding a multivariate Gaussian. The result can be plotted to illustrate the kinds of functions the GP considers plausible.

To draw samples from a dependent multivariate Gaussian distribution with a mean function $m(\mathbf{x})$ and a kernel function $k(\mathbf{x}, \mathbf{x}')$, it is most computationally efficient to begin by generating independent samples from a standard multivariate normal distribution. The independent normal variable, $\mathbf{y} = (y_1, y_2, \ldots, y_n) \sim \mathcal{N}(0, I)$,

can then be transformed into the desired correlated variable, $\mathbf{z} \sim \mathcal{N}(\boldsymbol{\mu}, \Sigma)$, via:

$$\mathbf{z} = L\mathbf{y} + \boldsymbol{\mu}, \quad \text{where } LL^T = \Sigma. \tag{3.19}$$

Here, $\boldsymbol{\mu} = m(X)$ is the vector of mean function values, and $\Sigma = k(X, X')$ is the covariance matrix evaluated at the sample locations. The lower triangular matrix L is obtained via Cholesky decomposition, such that $L\mathbf{y} \sim \mathcal{N}(0, \Sigma)$, imposing the desired correlation structure on the samples. A brief illustration of Cholesky decomposition is included in Appendix B.2.

Figure 3.2 (b-d) shows GP samples using different kernels, all with mean $m(\mathbf{x}) = -\mathbf{x}$ (dotted black line). These highlight how each kernel shapes the space of likely functions: the RBF yields smooth curves; the Matérn kernel allows for rougher functions; the decaying periodic kernel captures repeating but fading patterns. Such visualisations reinforce how kernel design encodes structural assumptions into the model, with the GP placing a high probability on functions that align with the properties of the kernel function.

Incorporating Noise into the GP

In practical optimisation problems, observations of the objective function are rarely noise-free. In the context of ICF optimisation, numerical noise may arise from discretisation error, solver tolerances, and stochastic elements within radiation-hydrodynamic simulations. In experimental settings, additional uncertainty arises from measurement noise and shot-to-shot variability. It is therefore essential to explicitly account for observational uncertainty within the Gaussian process framework.

Rather than assuming direct access to the true function value $f(\mathbf{x})$, observations are assumed to be drawn from a distribution centred on the observation. The measured value, y_i is assumed to deviate from the true value by some amount, ϵ_i

$$y_i = f(\mathbf{x}_i) + \epsilon_i. \tag{3.20}$$

The noise term (ϵ_i) is typically modelled as independent and identically distributed Gaussian noise with zero mean and variance σ_n^2:

$$\epsilon_i \sim \mathcal{N}(0, \sigma_n^2). \tag{3.21}$$

Under this assumption, the covariance between observed outputs is modified to include an additional variance term,

$$\mathrm{Cov}(y_i, y_j) = k(\mathbf{x}_i, \mathbf{x}_j) + \sigma_n^2 \delta_{ij}, \tag{3.22}$$

where δ_{ij} is the Kronecker delta. In matrix form, this corresponds to an additional diagonal noise term,

$$K(\mathbf{X}, \mathbf{X}) \to K(\mathbf{X}, \mathbf{X}) + \sigma_n^2 I. \tag{3.23}$$

The practical effect of this noise term is to prevent the GP from being forced to interpolate the observed data exactly. In the absence of noise ($\sigma_n^2 = 0$), the posterior variance at observed data points collapses to zero, and the GP perfectly interpolates the measurements. Introducing observational noise allows for finite posterior uncertainty at training points and permits small deviations from the observed values, reflecting the possibility that the observed values may not perfectly reflect the true function.

3.3.3 Conditioning the GP on Data

A key property of Gaussian processes is that they are closed under conditioning. This means that a prior GP, when conditioned on observed data, yields a posterior GP that remains Gaussian. The posterior incorporates both the structural assumptions encoded in the kernel function and the information contained within the observations.

Consider a set of training inputs $\mathbf{X} = \{\mathbf{x}_i\}_{i=1}^n$ with corresponding observed outputs $\mathbf{y} = (y_1, \ldots, y_n)^T$. Let $\mathbf{X}_* = \{\mathbf{x}_*^{(j)}\}_{j=1}^m$ denote a set of test points at which predictions are sought. Under the Gaussian process prior, the joint distribution of the observed outputs and the observed function values at the test points is given by:

$$\begin{pmatrix} \mathbf{y} \\ \mathbf{f}_* \end{pmatrix} \sim \mathcal{N}\left(\begin{pmatrix} m(\mathbf{X}) \\ m(\mathbf{X}_*) \end{pmatrix}, \begin{pmatrix} K(\mathbf{X}, \mathbf{X}) + \sigma_n^2 I & K(\mathbf{X}, \mathbf{X}_*) \\ K(\mathbf{X}_*, \mathbf{X}) & K(\mathbf{X}_*, \mathbf{X}_*) \end{pmatrix} \right), \tag{3.24}$$

where $K(\mathbf{X}, \mathbf{X})$ is the covariance matrix evaluated at the training points, $K(\mathbf{X}, \mathbf{X}_*)$ is the cross-covariance between training and test points, and $K(\mathbf{X}_*, \mathbf{X}_*)$ is the covariance matrix between the test points.

Using standard results for conditioning multivariate Gaussian distributions, the posterior distribution over the function values at the test points is also Gaussian,

$$\mathbf{f}_* \mid \mathbf{X}, \mathbf{y}, \mathbf{X}_* \sim \mathcal{N}(\boldsymbol{\mu}_*, \boldsymbol{\Sigma}_*), \tag{3.25}$$

with posterior mean

$$\boldsymbol{\mu}_* = m(\mathbf{X}_*) + K(\mathbf{X}_*, \mathbf{X}) \left[K(\mathbf{X}, \mathbf{X}) + \sigma_n^2 I\right]^{-1} (\mathbf{y} - m(\mathbf{X})), \tag{3.26}$$

and posterior covariance

$$\boldsymbol{\Sigma}_* = K(\mathbf{X}_*, \mathbf{X}_*) - K(\mathbf{X}_*, \mathbf{X}) \left[K(\mathbf{X}, \mathbf{X}) + \sigma_n^2 I\right]^{-1} K(\mathbf{X}, \mathbf{X}_*). \tag{3.27}$$

The posterior mean can be interpreted as the prior mean corrected by a data-dependent adjustment term. This correction is weighted by the covariance between test and training points and scaled by the inverse covariance of the observations. Points that are strongly correlated with well-constrained training locations receive larger updates.

The posterior covariance reflects the reduction in uncertainty due to observed data. The subtraction term reduces prior uncertainty in regions that are strongly correlated with the training inputs, while uncertainty remains large in poorly sampled regions of the parameter space.

3.3.4 Acquisition Functions

After conditioning on observed data, the Gaussian process posterior provides both a predictive mean and a predictive variance across the parameter space. A defining feature of Bayesian optimisation is that this predictive uncertainty can be exploited to guide the sampling strategy.

An effective optimisation strategy must balance two competing objectives: exploration, which favours sampling in regions of high uncertainty to improve the global model; and exploitation, which prioritises regions with high predicted objective values. This exploration-exploitation trade-off is formalised through the acquisition function [122]. The acquisition function is defined over the whole

parameter space, and the next evaluation point is chosen by maximising this function. Unlike the objective function, the acquisition function is inexpensive to evaluate and optimise.

A commonly used acqusition function is the upper confidence bound (UCB), defined as:

$$A_{UCB}(x) = m_*(x) + \kappa\sigma_*(x), \tag{3.28}$$

where $m_*(x)$ and $\sigma_*(x)$ are the posterior mean and standard deviation at location x, respectively. The equation introduces a hyper-parameter, κ, which determines the relative weighting of the uncertainty. Larger values of κ encourage exploratory behaviour by favouring regions of high predictive variance, while smaller values bias the optimisation towards exploitation of regions with high predicted objective values.

Alternative acquisition functions [112] include Expected Improvement (EI) and Probability of Improvement (PI), which instead quantify the expected gain relative to the current best observation. While differing in formulation, these functions similarly balance exploration and exploitation using the posterior predictive distribution.

3.3.5 Bayesian Optimisation

The previous subsection outlined all the pieces of the puzzle underlying Bayesian optimisation. Bayesian optimisation proceeds as an iterative loop comprised of three stages: specification of a prior model, conditioning on observed data to obtain a posterior, and selection of a new evaluation point via an acquisition function. This is illustrated in Figure 3.3, which shows how each stage builds upon the last. The GP is illustrated in this figure through 'sausage plots' which depict the mean (black line) and 95% confidence interval, $\pm 1.96\,\sigma(x)$ (grey shaded region). The true function is shown by the dotted black line. The four panels in this figure can be broken down as follows:

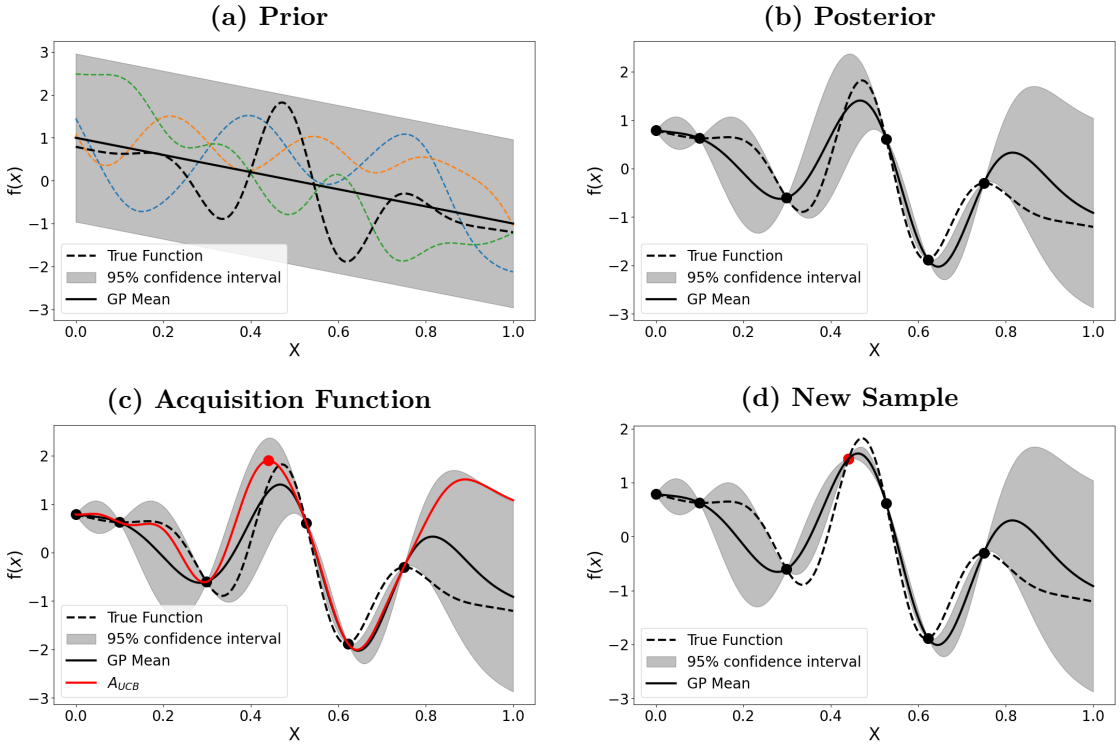


Figure 3.3: One-dimensional example of a Bayesian optimisation iteration. (a) GP prior, defined by a mean (solid black line) and kernel function (RBF). Dashed coloured lines indicate functions the prior deems plausible. The true function is indicated by the dotted black line. (b) GP posterior after conditioning on 6 observed data points. (c) UCB Acquisition function (red line) of the GP posterior. The maximum of this function is indicated and marks the location of the next sample. (d) Updated posterior after sampling the selected point and incorporating the new observation.

(a) GP Prior: The process begins with a GP prior, defined by a mean function and kernel, that encodes prior assumptions about the unknown objective function. In this case, it is assumed that there is a negative linear correlation between the objective function and X ($m(X) = 1 - 2X$), and an RBF kernel is used to capture smooth variations. This captures the general trend of the function, but fails to account for the oscillations in the function. The coloured lines indicate samples from the GP, which represent different plausible functions that are consistent with the prior assumptions.

(b) GP Posterior: The GP is then conditioned on some avaliable data to form a posterior distribution. The posterior represents the model's updated beliefs

about the function by modifying the mean and reducing uncertainty in regions correlated with observed points.

(c) Acquisition Function: The acquisition function leverages both the predicted mean and uncertainty of the posterior to guide the search for the optimum. In this example, the upper confidence bound (UCB) acquisition function (Equation 3.28) is used, and is shown by the solid red line. It balances exploration and exploitation by favouring points with high mean or high uncertainty. The next point to evaluate is chosen as the maximum of this acquisition function, marked by the red dot. A κ value of 2 is used here, which provides a moderate balance between exploration and exploitation.

(d) New Sample: The location identified by the acqusition function is then sampled, and the function value at this point is observed. The process then repeats, with the posterior being updated, a new acquisition function being computed, and the next sample location being selected. As more data is accumulated through this iterative process, the model's understanding of the objective function improves.

A significant advantage of Bayesian optimisation is that it produces a probabilistic surrogate model of the objective function. This surrogate model can therefore provide insight into the function's behaviour, such as sensitivity to different parameters, the presence of local optima, and the overall shape of the function. This is particularly useful when trying to understand the underlying physics of a system or extracting a more general understanding of the function being optimised.

3.4 Bayesian Package Implementation

Having established the theoretical formulation of Bayesian optimisation, this Section describes the practical implementation within the developed optimisation package. The implementation follows the iterative structure outlined in the previous Section, translating the probabilistic model and acquisition strategy into a computational workflow.

Consider a d-dimensional parameter space $\mathcal{X}$ over which the objective function $f(x)$ is defined. Prior to sampling, a mean function $m(x)$ and kernel function $k(x, x')$ are specified to define the GP prior, thereby encoding structural assumptions about the objective function.

If no initial dataset is available, an initial set is generated by evaluating the objective function at a small number of randomly selected points within the parameter bounds. These samples provide the GP with preliminary information about large-scale trends in the objective function, allowing it to begin resolving macro-scale structure before focusing on finer local features. If a dataset already exists, this step is skipped and the posterior is formed directly.

After n function evaluations, the dataset is represented as $\mathcal{D} = \{X, y\}$, where $X \in \mathbb{R}^{n \times d}$ contains the sampled locations (row-wise) and $y \in \mathbb{R}^n$ contains the corresponding objective values.

To ensure consistent numerical behaviour across objective functions with differing magnitudes, the outputs are normalised to the interval $[-1, 1]$ according to

$$y_{\text{norm}} = \frac{2\,(y - y_{\min})}{(y_{\max} - y_{\min})} - 1. \tag{3.29}$$

This transformation is well-defined provided $y_{\max} \neq y_{\min}$. In the degenerate case where all observed values are identical, the normalised outputs are set to zero, reflecting the absence of observable structure in the current dataset.

Normalisation ensures that the GP operates on a controlled output scale, preventing the magnitude of the objective function from implicitly influencing the relative weighting of predictive mean and variance within the acquisition function. This removes the need for problem-specific rescaling of acquisition hyper-parameters and improves robustness across optimisation tasks with differing objective magnitudes.

To determine the next evaluation point, a set of a candidate locations X_*, where $X_* \in \mathbb{R}^{a \times d}$, is generated by uniform random sampling within the parameter bounds. The posterior mean and covariance at the candidate locations are computed using

the standard GP update equations. Three covariance matrices are constructed:

$$K = k(X, X) + \sigma_n^2 I, \quad K_* = k(X, X_*), \quad K_{**} = k(X_*, X_*). \tag{3.30}$$

The posterior mean and variance are then given by

$$\mu(X_*) = m(X_*) + K_*^T K^{-1}(y_{\text{norm}} - m(X)), \tag{3.31}$$

$$\Sigma = K_{**} - K_*^T K^{-1} K_*, \tag{3.32}$$

$$\sigma^2(X_*) = \text{diag}(\Sigma). \tag{3.33}$$

In practice, the inverse of K is never computed explicitly. Instead, a Cholesky decomposition of K ($K = LL^T$) is performed, and linear algebra involving L is used to calculate the mean and variance. Importantly, this is a more computationally stable approach compared to directly inverting K.

The posterior distribution is then passed to the acquisition function $A(\mu(X_*), \sigma(X_*))$, and the candidate location that maximises this function is selected,

$$x_{\text{new}} = \arg\max_{x \in X_*} A(x). \tag{3.34}$$

The objective function is evaluated at this location to obtain y_{new}, and the dataset is augmented accordingly. The procedure is then repeated until a predefined stopping criterion is satisfied.

Restricting the next possible sample location to a randomly generated candidate set avoids the need for a secondary optimisation routine over the acquisition function itself. This approach is computationally efficient and, as discussed later, supports the extension to batch sampling by naturally promoting diversity among proposed evaluation points.

Pseudocode of this optimisation workflow is given in Algorithm 1. This baseline implementation represents the fundamental Bayesian optimisation loop. When applied to high-cost radiation-hydrodynamic simulations, however, several practical limitations arise. Subsection 3.4.2 details the algorithmic enhancements introduced to address these limitations.

Algorithm 1 Base Bayesian Optimisation

Inputs:

m	▷ prior mean function
k	▷ kernel/covariance function
A	▷ acquisition function
$D = \{x_i, y_i\}$	▷ initial (possibly empty) dataset
a	▷ number of candidate points per iteration
n_{rand}	▷ initial random samples if $D = \varnothing$
σ_n	▷ observation noise (or jitter) added to diagonal
Bounds $\in \mathbb{R}^{d\times 2}$	▷ row-wise [lower, upper]

Execution

if $D = \emptyset$ **then**
  $X = \mathcal{U}(\text{Bounds}, n_{\text{rand}})$ ▷ n_{rand} random samples
  $Y = \phi(X)$ ▷ evaluate objective
  $D = \{X, Y\}$
end if

While stopping criterion not met:

$X = D\{x_i\}$
$Y = D\{y_i\}$

$Y_{\text{norm}} = \frac{2(Y - \min(Y))}{\max(Y) - \min(Y)} - 1$ ▷ Guard: if $\max(Y) = \min(Y)$, set $Y_{\text{norm}} = 0$

$X_* \sim \mathcal{U}(\text{Bounds}, a)$ ▷ generate a candidate points

$K = k(X, X) + \sigma_n^2 I$
$K_* = k(X, X_*)$
$K_{**} = k(X_*, X_*)$

$L = \text{Cholesky}(K)$ ▷ $K = LL^T$
Solve $L\alpha = Y_{\text{norm}} - m(X)$
Solve $L^T\beta = \alpha$
Solve $L\gamma = K_*$

$\mu(X_*) = m(X_*) + {K_*}^T\beta$
$\sigma^2(X_*) = \textbf{diag}(K_{**} - \gamma^T\gamma)$

$x_{\text{new}} = \arg\max\left[A(\mu(X_*), \sigma(X_*))\right]$
$y_{\text{new}} = \phi(x_{\text{new}})$

$D = D \cup \{(x_{\text{new}}, y_{\text{new}})\}$

Step	Computation	Complexity
Compute K	$K = k(X, X)$	$\mathcal{O}(n^2 d)$
Compute K_*	$K_* = k(X, X_*)$	$\mathcal{O}(nad)$
Compute K_{**}	$K_{**} = k(X_*, X_*)$	$\mathcal{O}(a^2 d)$
Cholesky Decomposition	$LL^T = K$	$\mathcal{O}(n^3)$
Compute α	$L\alpha = Y_{\text{norm}} - m(X)$	$\mathcal{O}(n^2)$
Compute β	$L^T\beta = \alpha$	$\mathcal{O}(n^2)$
Compute γ	$L\gamma = K_*$	$\mathcal{O}(n^2 a)$
Posterior Mean	$m(X_*) + K_*^T\beta$	$\mathcal{O}(na)$
Posterior Variance	$\text{diag}\,(K_{**}) - \text{diag}\left(\gamma^T\gamma\right)$	$\mathcal{O}(na)$

Table 3.1: The computational scaling of key calcualtions in conditioning the GP. n is the number of training samples, a the number of acquisition candidates, and d the dimensionality of the parameter space.The $O(n^3)$ Cholesky decomposition dominates for large n, while the $O(a^2)$ computation of the K_{**} kernel matrix dominates for large a.

3.4.1 Computational Complexity

Bayesian optimisation provides an efficient framework for optimising expensive black-box functions, particularly when each function evaluation is computationally costly. However, the method is not without limitations. The principal constraint arises from the computational complexity associated with GP regression, which can become prohibitive as the scale of the problem increases.

The primary computational cost arises during conditioning of the GP, i.e. the transition from the prior to the posterior distribution after incorporating observed data. Assuming that the evaluation of the kernel between two points scales linearly with the dimensionality of the parameter space, $\mathcal{O}(d)$, the complexity of the calculations involved in updating the mean and covariance functions are given in Table 3.1.

This table highlights the strong dependence of the computational cost on the number of training points, n. The dominant term is the $\mathcal{O}(n^3)$ scaling associated with the Cholesky decomposition of the covariance matrix, K. This cubic scaling rapidly becomes the principal bottleneck as the dataset grows, and becomes prohibitive as n exceeds 10^5, placing practical limits on the applicability of standard GP-based Bayesian optimisation for large n.

For the random acqusition sampling approach used in this package, it is typically the case that $n \ll a$. In this regime, the dominant computational cost is the $\mathcal{O}(a^2)$ scaling of the K_{**} kernel computation. To mitigate this cost, the candidate set can be partitioned into smaller subsets. The acquisition function is then evaluated independently on each subset, and the optimal candidate across all subsets is selected.

The benefit of this approach is that K, L, α, and β depend only on the existing dataset, X, and are computed once per iteration and reused. Only the candidate-dependent terms require repeated evaluation. This reduces the $\mathcal{O}(a^2)$ scaling in the K_{**} computation, at the expense of repeated $\mathcal{O}(n^2)$ operations associated with the γ computation. In practice, when $n \ll a$, this trade-off reduces overall computational cost. However, in regimes where $n \gtrsim a$, partitioning the candidate set may increase computational overhead. The balance between these costs must therefore be considered when selecting an appropriate candidate batching strategy.

Reducing the computational dependence on a is particularly important when scaling Bayesian optimisation to higher-dimensional parameter spaces (e.g. $d \sim 10$). To illustrate this, consider a one-dimensional parameter space. If too few candidate points are sampled, the acquisition function may not be adequately resolved, and the optimiser may only select points of limited utility, effectively reverting to near-random sampling. In one dimension, resolving the acquisition landscape may require on the order of five candidate samples. Extending this argument to two dimensions, maintaining the same sampling density requires 5^2 candidates. In general, for a d-dimensional parameter space, this scales as 5^d. In ten dimensions, this corresponds to $\sim 10^7$ candidate samples. This issue can be mitigated through alternate acquisition function optimisation approaches, such as gradient-based methods; however, such approaches reduce the diversity random candidate sampling provides when coupled with batch sampling (discussed later). Improving acquisition function optimisation in high-dimensional spaces remains an area for future development within the package.

To further accelerate these computations, the package supports the use of both CPUs (central processing units) and GPUs (Graphics processing units). GPUs are

particularly well suited to the dense linear algebra operations performed in the GP regression. The implementation is hardware-agnostic and automatically detects available computational resources, allowing users to leverage GPU acceleration where available without modification to the code. This capability is especially beneficial for large-scale optimisation studies, where kernel evaluations and matrix decompositions dominate computational cost.

3.4.2 Additional Features

While the standard Bayesian optimisation framework discussed previously and outlined in Algorithm 1 provides an efficient global optimisation strategy, several practical challenges arise when applying it to ICF simulations. In particular, ICF optimisation is characterised by high computational cost per objective evaluation, the requirement to exploit parallel simulation workflows on HPC systems, and moderately high-dimensional parameter spaces.

To address these constraints, a set of physics-motivated extensions to the Bayesian optimisation procedure has been implemented within the developed package. These modifications are designed to reduce time-to-solution, improve sampling efficiency in parallel environments, and enhance robustness and scalability in practical ICF optimisation tasks. The extensions are modular, allowing the optimisation strategy to be tailored to the specific requirements of the given study. The additional features are outlined in the following subsections.

Batch Sampling

By default, Bayesian optimisation is a sequential iterative process in which maximising the acquisition function determines a single point to sample next. In the case where sampling the object function takes considerable time, such as carrying out and analysing experimental data, or executing long simulation workflows, it can take a sizeable amount of time to build a sufficiently large dataset to accurately model the objective function. Therefore, the first extension added to the package is batch sampling [123], allowing for multiple points to be sampled per iteration. This

aims to reduce the time-to-solution and, in the case of simulation based optimisation processes, allows for the efficient use of HPC clusters.

In this framework, batch sampling is accomplished by varying the set of hyper-parameters used in the acquisition function. In the case of the UCB acquisition function (Equation 3.28), this simplifies to varying a single hyper-parameter, κ, which controls the exploration-exploitation trade-off. A set of κ values between minimum and maximum bounds, κ_{min} and κ_{max}, are selected, with the number of κ values equal to the number of desired samples in the batch. For each κ value, a set of candidate points, X_*, are randomly sampled and the acquisition function is evaluated to determine the most suitable location to sample. This procedure is repeated for each κ value until a full batch of samples has been constructed.

The process is illustrated in Figure 3.4 for a one-dimensional parameter space using the UCB acquisition function. Panel (a) shows the true objective function (dotted black line) and the GP posterior (mean shown as a solid black line with variance indicated by the shaded grey region) after conditioning on three data points. Panel (b) overlays the UCB acquisition function for a small value of $\kappa = 0.5$, leading to exploitative sampling near the predicted optimum. This is repeated in panels (c) and (d) for $\kappa = 1.72$ and $\kappa = 6$, demonstrating the progression toward increasingly exploratory behaviour. The selected points are then combined to form a batch of three samples. Panel (e) shows the updated GP posterior after conditioning on the new batch.

In Figure 3.4, a small stochastic perturbation ϵ is added to the variance term of the acquisition function,

$$A_{UCB}\left(m_*(x), \sigma_*(x)\right) = m_*(x) + \kappa\left(\sigma_*(x) + \epsilon\right), \tag{3.35}$$

to avoid deterministic behaviour in regions of uniformly high uncertainty. In such regions, multiple points are equivalently valuable from an exploratory perspective, and strict maximisation of the acquisition function would repeatedly select the same location. The perturbation acts as a tie-breaking mechanism, ensuring that once

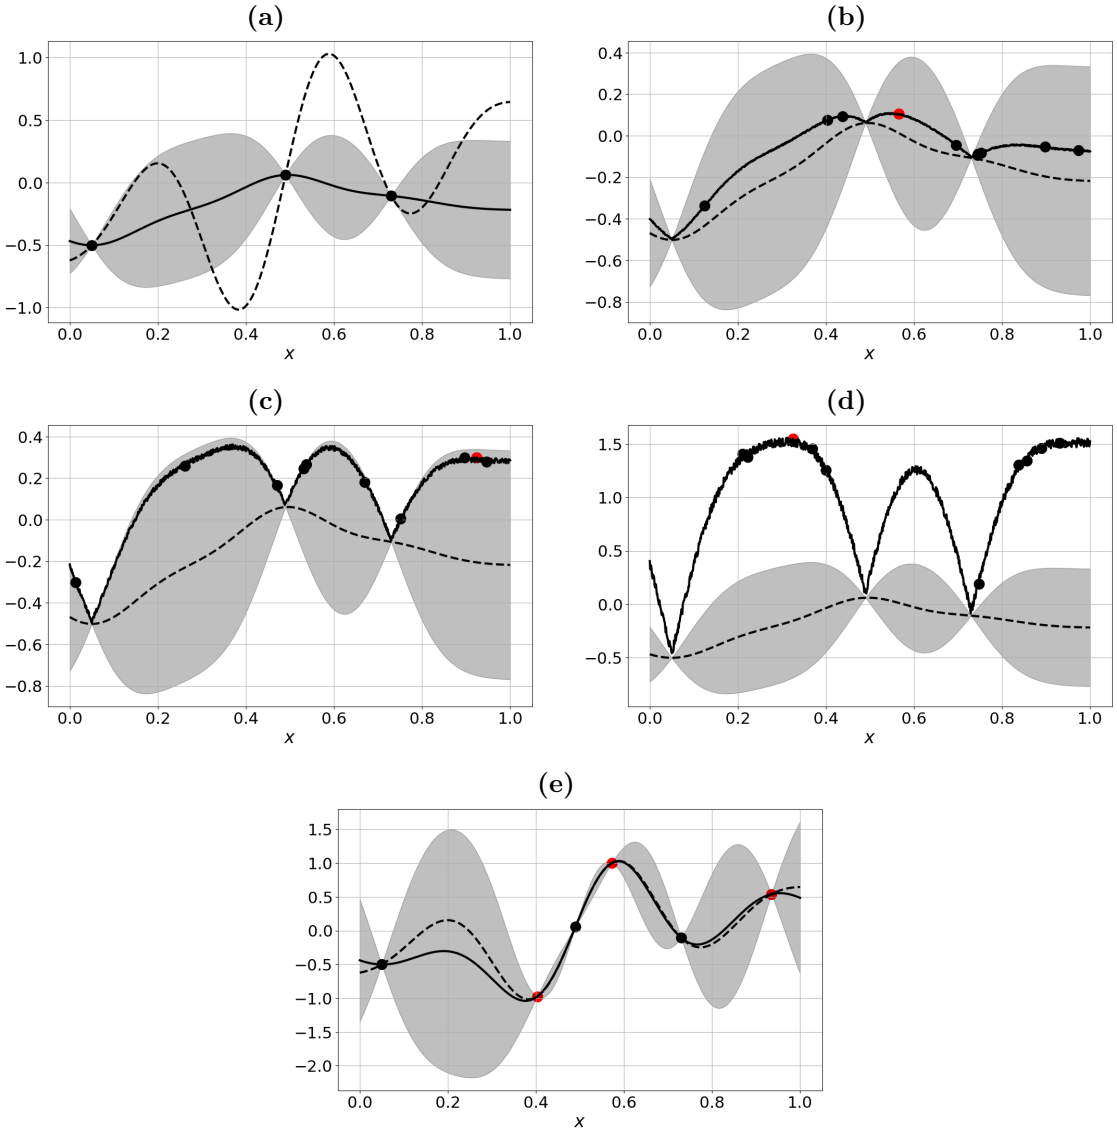


Figure 3.4: Illustration of batch sampling. (a) GP posterior after conditioning on three data points. (b)-(d) UCB acquisition functions (solid black line) constructed with increasing values of the exploration parameter κ, demonstrating the transition from exploitative to exploratory sampling behaviour. Ten randomly selected candidate points are evaluated (dots), and the candidate with the largest acquisition value is added to the batch. (e) Updated posterior after conditioning on the full batch of new observations.

differences in uncertainty fall below a threshold, point selection becomes effectively stochastic, promoting diversity within the batch.

Additional diversity is encouraged by restricting candidate sample locations to a randomly generated subset of the parameter space. Provided the candidate set is sufficiently large, regions of high acquisition values remain likely to be

selected, while the randomly selected candidates provides an additional level of stochasticity to improve diversity.

By explicitly leveraging the influence of κ on sampling behaviour, the package generates batches that span the full spectrum from exploitative to exploratory sampling. This approach not only fosters a richer sampling strategy but also mitigates the sensitivity of the optimisation process to a specific choice of hyper-parameters.

Hallucination

While batch sampling reduces the time-to-solution by selecting multiple points per iteration, it can still lead to inefficiencies when multiple points in a batch cluster around the same region of the parameter space. This occurs because samples within a batch are selected using the same posterior model and therefore have no knowledge of one another during selection. The effect becomes increasingly pronounced as the batch size grows. To mitigate this redundancy, hallucination (or sub-batch sampling) is implemented [124].

A key observation enabling hallucination is that the posterior variance of a Gaussian process does not depend on the observed value at a training point, but only on its location. Regardless of whether the sampled value is large or small, conditioning on that location reduces uncertainty in correlated regions of the parameter space. Only the posterior mean depends on the magnitude of the observation. This property allows the model to incorporate information about sample locations without requiring evaluation of the true objective function.

The procedure begins by dividing each batch into smaller sub-batches. The first sub-batch is selected using the standard batch sampling approach, where the hyper-parameter set is varied and each set corresponds to a new sample. The GP is then temporarily conditioned on these locations using the predicted mean at each point. This reduces the posterior variance in regions correlated with the selected locations while leaving the mean prediction unchanged. As the acquisition function increases in regions of high uncertainty, subsequent samples are discouraged from clustering near previously selected points.

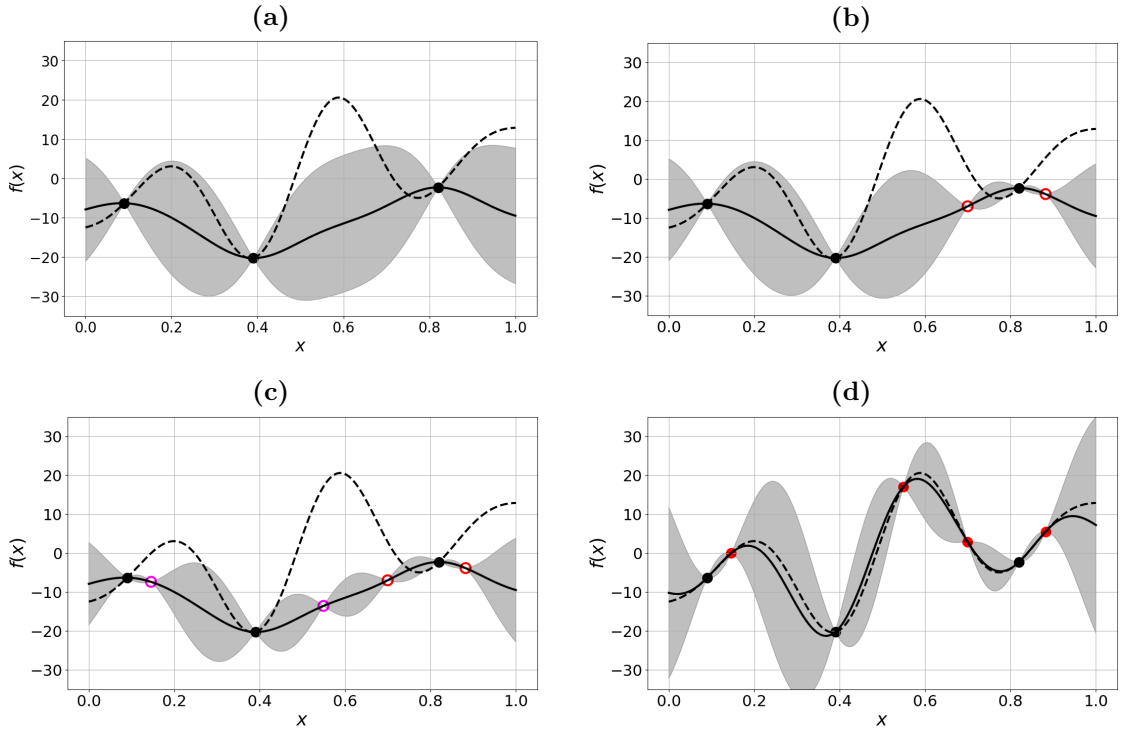


Figure 3.5: Illustration of sub-batch sampling using hallucination. (a) GP posterior after conditioning on 3 data points. (b) Posterior temporarily conditioned on the first sub-batch of 2 samples using the predicted mean values (red circles). The mean of the GP posterior remains the same, while the uncertainty is reduced in correlated regions. (c) Determination of the second sub-batch samples (purple circles) using the posterior in panel (b). (d) Final posterior after evaluating the true objective function at the batch sample locations. Increased diversity in the batch leads to more efficient sampling.

The next sub-batch is then determined using the same acquisition procedure on the temporarily conditioned model. This process repeats until all sub-batches have been selected. Only once the full batch has been constructed is the objective function evaluated at all selected locations, after which the Gaussian process is updated using the true observed values, refining both the posterior mean and variance.

This procedure is illustrated in Figure 3.5 for a batch of four samples constructed from two sub-batches of size two. Panel (a) shows the GP posterior after conditioning on three data points. Panel (b) shows the model after the first sub-batch has been selected and temporarily conditioned using the predicted mean values. Panel (c) shows the second sub-batch selected from the updated model. Finally, panel (d) shows the posterior after evaluating the objective function at all four batch

locations and conditioning on the true observed values. The figure demonstrates the effectiveness of hallucination at generating a diverse set of samples within a batch.

Dynamic Bounds Reduction

High-dimensional parameter spaces present a significant challenge for most optimisation approaches, and particularly for Bayesian optimisation. As shown in Table 3.1, the computational cost scales not only with the dimensionality (d) of the parameter space, but also with the number of acquisition samples (a) and data points (n). In practice, the number of candidate acquisition samples must increase rapidly with dimensionality to avoid the optimisation degenerating into effectively random sampling. This results in substantial computational expense for problems above $d \sim 12$.

To mitigate this behaviour, the package incorporates Dynamic Bounds Reduction [125], a strategy that progressively refines the parameter space during the optimisation process. Initially, the full parameter space is explored, allowing the optimiser to identify promising regions. As the optimisation progresses and improvements stagnate, the bounds are reduced by "zooming in" around high-performing regions while discarding less promising areas. This concentrates computational effort where improvements are most likely, effectively increasing sampling density within relevant regions of the space. By iteratively refining the domain in this manner, the search volume is reduced without sacrificing early-stage global exploration, improving convergence behaviour in moderately high-dimensional problems.

Practically, this procedure is governed by three parameters: a grace period (the number of initial iterations before bounds reduction is permitted), a stagnation threshold (the number of iterations without improvement), and a reduction factor controlling how much the bounds are contracted. After the grace period, if no improvement has been observed for the specified number of iterations, the bounds are reduced by centring them around the current best solution and shrinking their width by the chosen factor. If the contracted bounds exceed the original limits,

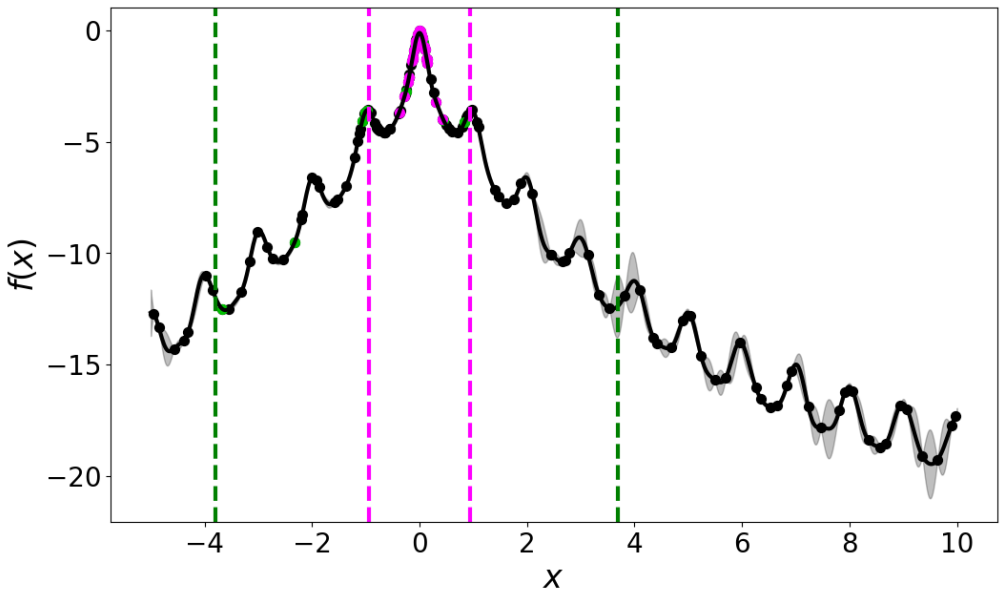


Figure 3.6: Illustration of dyanmic bounds reduction. After an inital grace period, if no improvement is observed for a specified number of iterations, the parameter bounds are contracted around the current best-performing solution. The dashed vertical lines indicate successive reductions of the search domain, concentrating sampling within increasingly localised regions of the parameter space.

they are shifted accordingly to remain within the initial parameter space while preserving the specified width.

An example of this process is shown in Figure 3.6. A grace period of 10 iterations is used, after which the bounds become eligible for reduction if no improvement is observed for 3 consecutive iterations. When this criterion is met, the bounds are reduced by a factor of 0.5, as indicated by the green and magenta dashed vertical lines at iterations 13 and 17. After each reduction, subsequent samples are restricted to the updated domain.

In low-dimensional examples, the effectiveness of this approach is difficult to demonstrate, as the optimiser can often identify the optimum before contraction occurs. However, in higher-dimensional parameter spaces, where acquisition sampling becomes increasingly diffuse, dynamic bounds reduction provides a significant improvement in practical performance.

Kernel Function Optimisation

The kernel function defines the covariance between any two points of the parameter space and thus determines key characteristics of the functions considered by the GP. However, the kernel function typically depends on a set of hyper-parameters, and choosing appropriate values for these hyper-parameters before sampling the objective function can be incredibly difficult. This task often requires extensive domain knowledge, previous experience optimising similar functions, or a lot of luck.

One of the major benifits of Bayesian optimisation is its ability to evaluate how likely the observed data is given the prior hypothesis. This makes it possible to vary the prior hypothesis (e.g. changing the hyper-parameters used in the kernel function) to determine which hypothesis is most likely to explain the observed data. As a result, the model can learn and update its hyper-parameter values throughout the optimisation process.

This is done by conditioning the GP on the observed data and then calculating the marginal likelihood of the data given the hyper-parameters [126]. The marginal likelihood is given by:

$$p(y|X,\theta) = \mathcal{N}(y|m(X), K(X,X;\theta)), \tag{3.36}$$

where θ are the hyper-parameters of the kernel function. The marginal likelihood is then maximised with respect to the hyper-parameters, θ, to determine the most likely values. This is done by optimising the log marginal likelihood, which is given by:

$$\log p(y|X,\theta) = -\frac{1}{2}y^T K(X,X;\theta)^{-1}y - \frac{1}{2}\log|K(X,X;\theta)| - \frac{n}{2}\log 2\pi, \tag{3.37}$$

where n is the number of data points. There are two competing terms in the log marginal likelihood: the first term ($-\frac{1}{2}y^T K(X,X;\theta)^{-1}y$) describes how well the model fits the data, while the second term ($-\frac{1}{2}\log|K(X,X;\theta)|$) penalises overly complex models (short lengthscales). The final term ($-\frac{n}{2}\log 2\pi$) is a normalisation term, which ensures the expression is a true PDF. This balance helps to prevent both overfitting and underfitting. The log marginal likelihood is optimised numerically using gradient-based methods.

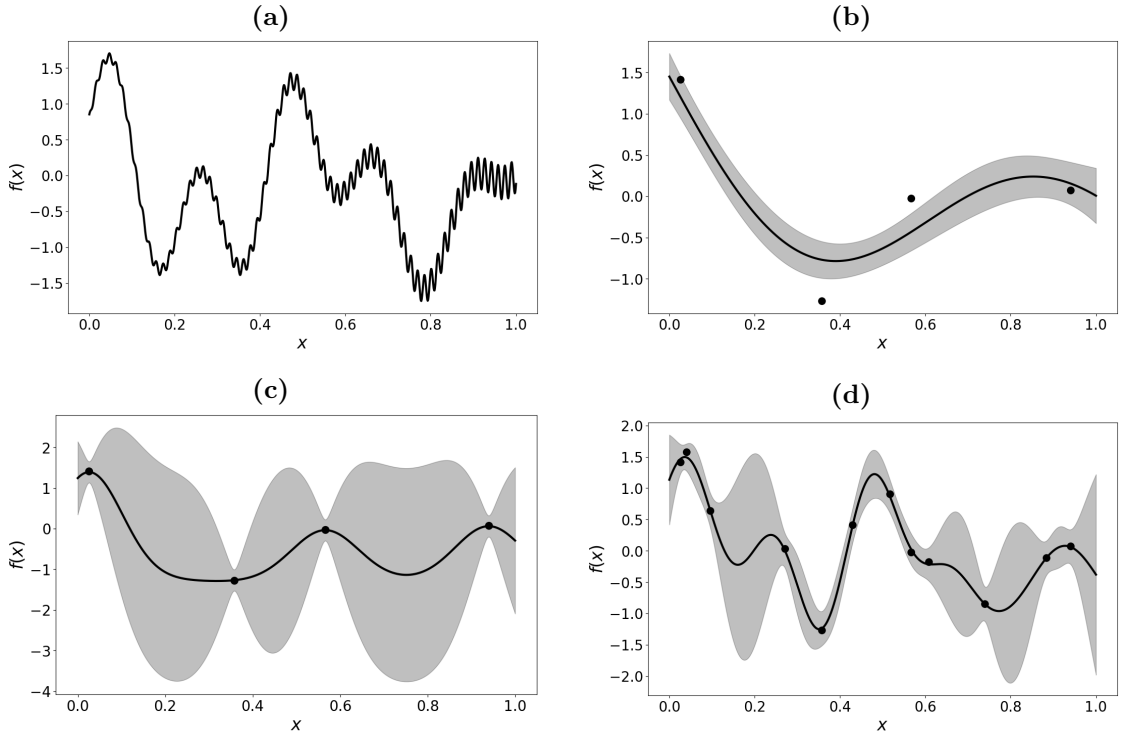


Figure 3.7: Kernel function hyper-parameter optimisation procedure. (a) Example noisy objective function. (b) GP posterior using an initial RBF kernel with an excessively large lengthscale, leading to underfitting. (c) The same GP posterior after optimising the kernel lengthscale. (d) Further lengthscale optimisation after many additional observations. The optimisation process yields a lengthscale which interpolates the data in the least expressive way possible. The lengthscale does not overfit to noise, as a noise term is explicitly included.

In the optimisation studies presented in Section 4, this hyper-parameter optimisation is used to identify suitable lengthscales for the RBF kernel function. An example of this is shown in Figure 3.7, where a noisy 1D objective function as shown in panel (a) is modelled. An RBF kerel function is used with an initial lengthscale of 0.5, which is far too large to accurately model the objective function. Panel (b) shows the GP posterior after being conditioned on 4 data points. The initial lengthscale is too large for the function to interpolate between the data points, and the model is unable to capture any properties of the true function. By scanning candidate lengthscales and maximising the log marginal likelihood, the most likely value given the data is found to be approximately 0.077. Panel (c) shows the Bayesian model after updating the kernel function with this new lengthscale. The model is now able to accurately interpolate between the data points. As more

data points are added, as shown in panel (d), the model is able to further refine its predictions to a lengthscale of 0.063 and accurately model the true function.

The role of the noise tolerance is also illustrated. A noise tolerance of 10^{-2} is used in this example, allowing the model to account for the low amplitude, high frequency oscillations in the objective function without over-fitting to these features. If the noise tolerence is reduced to 10^{-5} and the lengthscale is re-optimised, the model instead selects a lengthscale of 0.016, as it now must account for the high frequency oscillations in the data. This demonstrates how marginal likelihood optimisation selects the simplest model consistent with the observed data, favouring the largest lengthscale that remains supported.

By adaptively refining the kernel hyper-parameters throughout the optimisation process, the surrogate model remains appropriately flexible while avoiding overfitting. This balance arises naturally from the structure of the marginal likelihood, which rewards accurate interpolation of the data while penalising unnecessary model complexity through the log-determinant term. As a result, the optimisation selects the largest lengthscale that remains consistent with the observed data, favouring the smoothest function capable of explaining the measurements (Occam's razor).

3.5 Particle Swarm Optimisation (PSO)

Particle Swarm Optimisation (PSO) is a population-based optimisation algorithm first introduced in 1997 [119] to describe social behaviour in swarms of fish and birds. In contrast to surrogate-based approaches such as Bayesian optimisation, PSO is a model-free method that evolves a population of candidate solutions directly within the parameter space. This makes it particularly well-suited to high-dimensional optimisation problems where surrogate construction becomes computationally expensive or unreliable.

In PSO, a 'swarm' of candidate solutions (referred to as particles) moves through the parameter space in search of an optimum [127]. Each particle represents a point in the parameter space and is associated with both a position and a velocity. During the optimisation process, particles are influenced by their own best-known

position ($\mathbf{L}_p^n$), as well as the best-known position identified by the swarm as a whole ($\mathbf{G}^n$). The velocity, $\mathbf{v}_p$, and position, $\mathbf{x}_p$, of each particle are then updated as

$$\mathbf{v}_p^{n+1} = h_i \cdot \mathbf{v}_p^n + r_l^n h_l \cdot \left(\mathbf{L}_p^n - \mathbf{x}_p^n\right) + r_g^n h_g \cdot \left(\mathbf{G}^n - \mathbf{x}_p^n\right), \tag{3.38}$$

$$\mathbf{x}_p^{n+1} = \mathbf{x}_p^n + \mathbf{v}_p^{n+1}, \tag{3.39}$$

where subscript p denotes particle specific values and superscript denotes the iteration.

Equation 3.38 is comprised of three terms, each introducing a hyper-parameter, h_i, h_l and h_g. In order, these terms can be described as an inertia term, which determines how much the particle remains on its original trajectory, a local attraction, which determines how strongly the particle is attracted to its own best-known position, and a global attraction, which determines how strongly the particle is pulled toward the swarms best found location. In this description, the inertia weight controls the extent to which particles explore new regions of the parameter space, while the local and global hyper-parameters pull particles toward promising areas already explored by themselves or by the swarm. The interplay between these terms governs the balance between exploration and exploitation within the swarm.

The local and global attraction terms also include random numbers, r_l^n and r_g^n, uniformly distributed in $[0, 1]$. The use of these random numbers adds a level of stochasticity to the particle motion which helps to maintain diversity in the swarm. As a result, particles approach their local and global best positions with differing magnitudes and directions, which helps to preserve swarm diversity and mitigates premature convergence to suboptimal solutions. Importantly, while individual particle updates are stochastic, the expected value of the attraction terms remains biased toward the best-known positions, ensuring the particles are encouraged toward promising regions of the parameter space.

Of course, to ensure particle trajectories are bounded and the swarm converges, these hyper-parameters cannot take any value [128]. First, the three hyper-parameters should all be positive, $h_i, h_l, h_g > 0$. Negitive h_l and h_g hyper-parameter values convert the attration points into repulsers, with the strength of the repulsion

increasing with distance. As the particle is pushed away, it gets pushed harder, resulting in instability. A negative inertial term inverts the velocity of the particle every iteration, leading to odd oscillatory behaviour. Also, for stable dynamics, the inertial term $h_i < 1$. This ensure particle motion is damped in the case without any attractors, or in the case when the particle is close to the attractor and the attraction force is negligible. Finally, the sum of $h_l + h_g < 4$. When this is not satisfied, the system can enter oscillatory or divergent regimes.

The PSO algorithm is particularly well-suited for high-dimensional parameter spaces as the algorithm is naturally parallelisable, as each particle can be evaluated independently, and requires only a small amount of stored information to power the search. This property aligns well with HPC-based simulation campaigns.

The PSO algorithm is also robust to noise in the objective function, as it uses a population of particles to explore the parameter space. This means that even if some particles are misled by noise, the swarm as a whole can still converge to a good solution. However, PSO can struggle with premature convergence, particularly in multi-modal landscapes, where particles may become trapped in local optima before adequately sampling the parameter space. Extensions to address these challenges are discussed in Section 3.6

3.6 PSO Package Implementation

The PSO algorithm is comparatively simple to implement, requiring only a small set of hyper-parameters and minimal bookkeeping per particle. Each particle stores only its position, velocity and its personal best-found position and value; the global best-found position and value are maintained once for the swarm. This lightweight representation makes PSO well-suited to large-scale simulation campaigns, where objective evaluations dominate computational cost.

The implementation begins by initialising particle positions randomly within the parameter bounds and assigning velocities with a specified initial magnitude and random direction. The personal best value for each particle is initialised to a placeholder value of $-\infty$ (implemented in Python as `-np.inf`), as is the global best

value. Each particle then evaluates the objective function at its current position. If the sampled value exceeds the particle's stored best value, its personal best is updated accordingly. Similarly, if any particle achieves a value exceeding the current global best, the swarm best is updated.

After evaluation and best-value updates, particle velocities and positions are advanced using Equations 3.38 and 3.39. After each position update, boundary conditions are enforced to ensure particles remain within the domain. In this implementation, reflective boundary conditions are applied: if a particle crosses a parameter bound, its position is reflected back into the domain and the corresponding velocity component is reversed. This approach preserves the particle's velocity magnitude while preventing exploration outside the parameter space. The cycle of evaluation, memory update, and position update is repeated until a stopping criterion is met, such as a maximum number of iterations, computational limits, or a convergence condition. Algorithm 2 summarises this procedure in pseudocode form.

3.6.1 Additional Features

While the PSO algorithm is simple to implement, there are a number of common challenges the algorithm faces. In particular, the optimisation is sensitive to the hyper-parameter selection, and striking the correct balance between exploration and exploitation can be difficult. Furthermore, in high dimensional parameter spaces it is easy for the swarm to converge to a local optimum before adequately sampling the domain. To address these challenges and better adapt PSO to the demands of ICF optimisation, several physics-motivated extensions to the standard PSO framework are implemented in the developed package.

Multi-Species Swarm

In standard particle swarm optimisation, all particles share a common set of hyper-parameters, resulting in a uniform search behaviour across the swarm. However, this uniform behaviour is highly sensitive to the exact choice of hyper-parameters. To address this, a multi-species formulation of PSO is introduced, in which the

Algorithm 2 Particle Swarm

inputs (hyper-parameters):

N ▷ number of particles
h_1 ▷ particle intertia
h_2 ▷ local optimum weighting
h_3 ▷ global optimum weighting
v_0 ▷ initial particle velocity
Bounds $\in \mathbb{R}^{d\times 2}$ ▷ row-wise [lower, upper]

parameters:

$X = [x_1, ..., x_N]$ ▷ N particle positions
$V = [v_1, ..., v_N]$ ▷ N particle velocities
$P = [p_1, ..., p_N]$ ▷ Particles best found value (local)
$P_x = [p_{x1}, ..., p_{xN}]$ ▷ Particles best found location (local)
G ▷ Swarms best found value (global)
G_x ▷ Swarms best found location (global)

execution:

Initialise X with random locations
Initialise V with velocity magnitude v_0 in a random direction
Initialise P and G with $-\infty$

While stopping criterion not met:
for $i = 1$ to N:
$\phi(x_i)$ ▷ sample function at x_i
if $\phi(x_i) > p_i$:
$p_i = \phi(x_i)$ ▷ update particle best found value
$p_{xi} = x_i$ ▷ update particle best found position
if $\phi(x_i) > G$:
$G = \phi(x_i)$ ▷ update global best found value
$G_x = x_i$ ▷ update global best found location

for $i = 1$ to N:
$r_1, r_2 = \mathcal{U}([0, 1], 2)$ ▷ random samples
$v_i = h_1 v_i + h_2 r_1 (p_{xi} - x_i) + h_3 r_2 (G_x - x_i)$ ▷ update particle velocity
$x_i = x_i + v_i$ ▷ update particle position

swarm is composed of multiple particle species, each defined by a distinct set of hyper-parameters governing inertia, local attraction, and global attraction [129, 130]. Particles are assigned to species according to specified probabilities, allowing multiple search behaviours to coexist simultaneously within a single optimisation run. This enables the swarm to maintain global exploration through more exploratory species, while concurrently performing local refinement through more exploitative species.

By allowing different species to operate concurrently, the optimiser can more effectively refine promising configurations while reducing the risk of premature convergence to sub-optimal configurations. The use of multiple species also reduces sensitivity to any single hyper-parameter choice, improving robustness across a range of optimisation problems.

Exploration Term

While the introduction of multiple particle species enables a range of search behaviours to coexist within the swarm, standard PSO dynamics still rely exclusively on attraction toward previously identified promising locations. As a result, even exploratory species can become biased toward regions already sampled by the swarm, particularly once a strong global best has emerged. This behaviour can limit the ability of the swarm to systematically probe poorly explored regions of the parameter space.

To address this, an explicit exploration term [131] is introduced into the velocity update equation (Equation 3.38), providing a mechanism by which particles can be actively encouraged to sample regions poorly explored. This term introduces an additional attraction toward locations in the parameter space that are distant from previously evaluated points, thereby promoting exploration independently of the current best-known solutions. The velocity update equation now reads

$$\mathbf{v}_p^{n+1} = h_i \cdot \mathbf{v}_p^n + r_l^n h_l \cdot \left(\mathbf{L}_p^n - \mathbf{x}_p^n\right) + r_g^n h_g \cdot \left(\mathbf{G}^n - \mathbf{x}_p^n\right) + h_E \left(\mathbf{E}_p^n - \mathbf{x}_p^n\right). \quad (3.40)$$

where h_E is the weight of the exploration term and $\mathbf{E}_p^n$ is the region of low sample density identified for particle p at iteration n.

The regions of low sample density are identified by randomly sampling across the parameter space and determining the sample furthest from any previous samples. In practice, a dedicated, purely exploratory subset of the swarm flits between unsampled regions. By explicitly encoding exploration in this manner, the optimiser is able to maintain coverage of the broader parameter space even in the presence of strong local or global attractors.

In the context of ICF optimisation, such a mechanism may be valuable for identifying alternative implosion pathways that may not be accessible through incremental refinement of existing designs. When combined with the multi-species framework, the explicit exploration term enables a subset of particles to prioritise global search while others focus on local optimisation, improving robustness in complex, multi-modal landscapes. Although this explicit exploration term is now implemented within the PSO framework, it was introduced after completion of the optimisation studies presented in this thesis. Consequently, it was not activated in the optimisation campaigns reported in Chapter 4. The results presented therefore correspond to the multi-species PSO formulation without explicit exploration forcing.

Velocity Boosting

Despite the introduction of multi-species behaviour and an explicit exploration term, particle swarm optimisation can still experience stagnation when particles converge into narrow regions of the parameter space and velocities decay toward zero. In such cases, the swarm may remain trapped near sub-optimal configurations, particularly in noisy or weakly informative optimisation landscapes where small gradients provide little guidance.

To mitigate this behaviour, a velocity boosting mechanism is introduced, designed to actively dislodge particles that have stagnated for an extended period [132]. When a particle's velocity falls below a specified threshold for a certain number of iterations, its velocity is boosted, effectively redistributing the particle within the local region of the parameter space. This controlled perturbation allows particles

to escape shallow basins of attraction or noise-induced traps while preserving the overall swarm dynamics.

In the context of ICF optimisation, velocity boosting is particularly useful for navigating rugged objective landscapes where multiple implosion configurations exhibit comparable performance. By intermittently re-energising particles that have stalled, the optimiser is better able to identify alternative pathways and avoid premature convergence. When combined with multi-species behaviour and explicit exploration, velocity boosting provides a robust mechanism for maintaining long-term swarm activity without compromising convergence toward high-performing solutions.

3.7 Summary

This chapter describes the development of two complementary optimisation frameworks tailored to the demands of ICF design: GP-based Bayesian optimisation and particle swarm optimisation. Both methods are implemented as flexible Python packages and adapted specifically for high-cost, simulation-driven optimisation problems typical of radiation-hydrodynamic ICF studies.

Bayesian optimisation provides a sample-efficient strategy well suited to expensive objective functions, constructing a probabilistic surrogate model that guides evaluation through quantified uncertainty. Its strength lies in minimising the number of costly simulations required to locate high-performing designs. However, this efficiency comes at the expense of increased computational overhead associated with GP regression, which limits scalability as dataset size and dimensionality grow.

In contrast, PSO operates directly in parameter space without constructing an explicit surrogate model. Its lightweight structure and natural parallelism make it well suited to higher-dimensional problems and large-scale HPC campaigns. While typically less sample-efficient than Bayesian optimisation, PSO offers robustness and scalability in complex, multi-modal design landscapes.

Both frameworks are extended with physics-motivated enhancements designed to address the practical constraints of ICF optimisation. For Bayesian optimisation, batch sampling, hallucination, dynamic bounds reduction, and kernel hyper-

parameter optimisation are introduced to improve efficiency, parallelisation, and scalability. For PSO, multi-species swarm dynamics, explicit exploration forcing, and velocity boosting are incorporated to mitigate premature convergence and improve robustness in noisy, high-dimensional spaces.

These developments provide optimisation toolkits capable of addressing the diverse challenges posed by ICF design. The complementary characteristics of the two approaches enable exploration of a wide range of design problems, and providing the methodological foundation for the optimisation studies presented in the following chapter.

4

Optimisation of Inertial Confinement Fusion Implosions

Contents

4.1 Introduction 100
4.2 Wetted Foam Implosions 101
4.2.1 Wetted Foam Implosions 101
4.2.2 Target Fabrication 102
4.2.3 Convergence Ratio 103
4.2.4 Wetted Foam Experiments 104
4.3 Re-optimisation of Low Convergence Ratio Wetted Foam Implosions 105
4.3.1 Defining the Low Convergence Ratio (Low-CR) Regime 105
4.3.2 Previous Optimisation Approach 106
4.3.3 Re-optimisation of Low-CR implosions 108
4.3.4 Performance Comparison 110
4.3.5 Evolution Through the Parameter Space 112
4.3.6 Bayesian Surrogate Model Accuracy 116
4.3.7 Breakdown of Hydroscaling 118
4.4 Fast Ignition Implosion Optimisation 120
4.4.1 Fast Ignition Implosions 120
4.4.2 Previous Fast Ignition Design Studies 121
4.4.3 Optimisation of Fast Ignition Implosions with Improved Parametric Instability Robustness 121
4.4.4 Optimised Laser Pulse and Compressed Fuel Configurations 123
4.4.5 Modifying the Optimum 125
4.5 Extension of FI Implosions to ArF Drivers 126
4.5.1 Argon Fluoride Laser Technology 126
4.5.2 Optimisation of Fast Ignition Implosions with ArF Drivers 127

4.6 Summary 128

4.1 Introduction

This chapter presents the application of the optimisation frameworks developed in Chapter 3 to the design of ICF implosions. Having established the theoretical basis of the optimisation methodologies and their numerical implementation, the focus now shifts to their use in physically relevant ICF design problems. These applications serve both to benchmark the optimisation approaches against established results and to explore regimes in which conventional scan-based strategies are inefficient or insufficient.

Section 4.2 introduces wetted foam implosions, which utilise a foam structure to suspend liquid DT fuel. These targets have attracted significant interest due to the additional control they provide over implosion hydrodynamics [133, 134], their ability to enable hotspot formation at reduced compression [135], and, consequently, their potential to access ignition regimes with improved stability against the growth of hydrodynamic instabilities (see Section 2.6.1) [136].

Wetted foam implosions have been the subject of a previous optimisation study using sequential linear scans of the parameter space [137]. This work defines a “Low Convergence Ratio" regime, seleted based on experimental observations of implosion performance being well described by 1D hydrodynamic simulations [133]. While a sequential linear scan can be effective in identifying viable configurations, such approaches do not capture parameter coupling and inefficiently explore high-dimensional design landscapes. Section 4.3 presents the reoptimisation of these implosions, using the previously published results as a benchmark against which the optimisation strategies developed in this thesis can be evaluated.

An alternative ICF approach, known as fast ignition (FI) [70, 138–140], has also emerged as a promising route to achieving high fusion gain. In FI, the fuel compression and ignition phases are decoupled, with a separate, ultra-intense laser

pulse used to ignite a pre-compressed fuel assembly [141]. This separation allows the fuel to be compressed to high areal densities at relatively low temperatures, forming an isochoric assembly that is subsequently ignited. This enables the compression and ignition stages to be optimised largely independently. If successfully implemented, such high fuel-mass assemblies offer the potential for significantly enhanced fusion gain [76]. Section 4.4 presents the optimisation of the compression stage of FI implosions using the particle swarm approach. Fast ignition presents a different optimisation landscape, characterised by distinct objective functions and constraints compared to conventional central-hot-spot implosions. Its inclusion therefore demonstrates that advanced optimisation frameworks are applicable across a broad range of design challenges. Owing to the increased dimensionality of the FI parameter space, Bayesian optimisation is prohibitively expensive, and is not applied to this task.

Section 4.5 then extends the FI optimisation to using an Argon Fluoride (ArF) laser driver. The benefits of ArF lasers for ICF applications are discussed, and the impact of this change on the optimised implosion performance is evaluated. While ArF lasers are not currently available at the energy scales required for ICF, they are a promising candidate for future high-gain direct-drive facilities [142].

4.2 Wetted Foam Implosions

4.2.1 Wetted Foam Implosions

Wetted foam targets replace the DT ice layer used in conventional target designs with a low-density (< 30 mg/cm^3) foam which is saturated with liquid DT [143]. This substitution allows foam targets to be fielded at higher temperatures, between 20 K and 26 K, compared to the sub-20 K of the ice DT capsules [144, 145]. These higher fielding temperatures correspond to a higher density vapour in the central cavity between 0.6 and 4 mg/cm^3. The increased vapour temperature and density raise the internal pressure within the capsule, which opposes the imploding shell. As a result, the achievable implosion velocity and peak compression are reduced. By varying the fielding temperature, the amount of implosion resistance can be

controlled [133]. This provides a direct means of tuning the hydrodynamic properties of the implosion, including its velocity, shell thickness, and total convergence.

As discussed in Section 2.6.1, hydrodynamic instabilities are a key degradation mechanism in inertial confinement fusion. The growth rate, and therefore impact, of hydrodynamic instabilities depends on parameters such as the implosion velocity, the shell thickness, and the total convergence [146]. In tuning these properties, implosions can be made more robust to hydrodynamic instabilities.

In addition to the hydrodynamic control wetted foams provide, the higher vapour density increases the fuel mass contained as DT vapour within the central cavity. For conventional DT ice targets fielded at the NIF, the vapour mass is typically ~ 0.75 μg, insufficient to produce a hotspot with $\rho R \gtrsim 0.3$ g/cm^2 without substantial compression. Instead, contributions from the inner surface of the DT ice layer are required, although high convergence is still necessary to achieve ignition-relevant conditions. In contrast, the increased vapour mass in wetted foam designs enables hotspot formation from material originating in the vapour region at comparatively lower convergence (convergence ratios ~ 16 compared to ~ 30 for conventional targets) [135]. This shift in hotspot formation mechanism allows access to ignition-relevant conditions within a lower-convergence regime, potentially improving hydrodynamic robustness while maintaining suitable ignition conditions.

4.2.2 Target Fabrication

Foams offer major fabrication advantages over conventional DT ice targets. Ice-layered targets require precise cryogenic layering by slowly cooling the capsule to below 20 K [145]. The solidification of the DT layer must then stem from a single seed crystal to avoid defects at grain boundaries which cannot be corrected once solidified [147]. Consequently, only around one in three ice-layering attempts produces a layer of acceptable quality.

In contrast, wetted foam targets do not require cryogenic cooling, as the liquid DT permeates a pre-fabricated foam scaffold. The defects are determined by the uniformity of the foam itself, and are therefore largely deterministic and reproducible,

rather than stochastic consequences of crystal growth [144]. Overall, fabrication time for a wetted foam capsule is currently approximately one-sixth that of a comparable ice-layered target [144]. This drastic speedup is a critical advantage of foams for eventual large-scale inertial fusion energy (IFE) applications, where high target throughput will be essential.

Furthermore, the formation of an ice layer requires the presence of tritium in the fuel, as the β-decay smooths surface roughness over time [148]. This necessity for tritium complicates research and development, as tritium handling requires specialised facilities and safety protocals. Wetted foam targets, however, can be fabricated and fielded using a wide range of fuel mixtures, including pure deutrium. This allows the testing of capsule performance without the added complexity of tritium handling and lowers the overall fusion yield, reducing optical and diagnostic damage during test implosions.

4.2.3 Convergence Ratio

The convergence ratio (CR) of an implosion is a measure of how much the capsule is compressed during the implosion. In this thesis it is defined as the ratio of the initial inner edge of the fuel layer to the minimum radius of the hotspot

$$CR = \frac{R_i}{R_{\mathrm{HS}}}. \tag{4.1}$$

where R_i is the initial inner radius of the fuel layer and R_{HS} is the minimum hotspot radius.

For the 1D simulations presented in this section, the code enforces perfect spherical symmetry. The hotspot is therefore isotropic, and the minimum hotspot radius is uniquely defined. In practise, real implosions are unlikely to be isotropic due to asymmetric drive conditions and hydrodynamic isntabilities. In such cases, an effective hotspot radius is inferred from the hotspot volume via

$$R_{\mathrm{HS}} = \left(\frac{3V}{4\pi}\right)^{1/3}, \tag{4.2}$$

which corresponds to the radius of a sphere with equivalent volume.

Intuitively, a larger CR, corresponding to a greater degree of compression, should lead to higher temperautres, densities, and hence fusion yields. In an ideal case, there is indeed a strong positive correlation between convergence ratio and fusion yield. However, high CR implosions are more susceptible to hydrodynamic instabilities, which can lead to loss of symmertry and the mixing of cold, dense material into the hotspot. It is therefore a balancing act to choose a CR that is high enough to acheive the required hotspot conditions, but low enough to avoid significant instability growth and degredation.

4.2.4 Wetted Foam Experiments

Targets using wetted foams have been experimentally fielded at the NIF since 2016 [149, 150]. These initial experiments demonstrated that wetted foam targets could be successfully fabricated, fielded, and driven on the NIF, and that measurable neutron yield could be generated. Subsequent experiments [133] further demonstrated the control that wetted foam targets provide over the hydrodynamic properties of the implosion through adjustment of the fielding temperature.

These studies showed that, when the compression of the capsule was limited, key implosion metrics (neutron yield, hotspot temperature, and burn width, etc.) were well reproduced by one-dimensional radiation-hydrodynamic simulations. However, as the level of compression increased, experimental results increasingly deviated from one-dimensional predictions, indicating the growing influence of multi-dimensional effects such as hydrodynamic instability growth and asymmetry.

Work toward implementing direct-drive wetted foam implosions on the NIF culminated in an experiment in 2025 [151], with further work ongoing [136, 152]. These experiments aim to demonstrate significant gain while remaining robust to instabilities; however, they are conducted using the NIF's polar direct drive (PDD) configuration [153]. This geometry is not optimised for direct-drive symmetry, and drive asymmetries may remain a significant performance-limiting factor.

4.3 Re-optimisation of Low Convergence Ratio Wetted Foam Implosions

A simulation-based study of wetted foam performance in a regime that is anticipated to be robust to hydrodynamic instabilities has been conducted by a previous member of the group, (the now Dr) Robert Paddock [137]. In this work, the so-called "Low Convergence Ratio" (Low-CR) regime is defined, which enforces hydrodynamic- and driver-based constraints on the implosion. Under these conditions, five different capsule sizes, each corresponding to a different energy scale, were explored. Each capsule size was optimised for both a 3-pulse and 4-pulse laser sequence, resulting in a total of 10 optimised target configurations.

The hydrodynamic simulation code used to perform these optimisations was `HYADES` [98]. This is a one-dimensional Lagrangian radiation-hydrodynamics code. 1D codes are often inaccurate in modelling full ICF implosions, however is appropriate for the Low-CR regime, where implosions are anticipated to be robust to hydrodynamic instabilities and therefore well modelled by 1D simulations.

4.3.1 Defining the Low Convergence Ratio (Low-CR) Regime

As previously defined, the convergence ratio refers to the amount of compression that the target undergoes during implosion (Equation 4.1). Experimental studies have shown that implosions with low convergence ratios can be well modelled by 1D radiation hydrodynamic simulations [133]. In the Low-CR regime, a threshold value of 16 is imposed on the CR of the implosion.

In addition to the CR, constraints are also imposed on other properties. First, the thickness of the shell plays a crucial role in determining how hydrodynamically stable the implosion is. If the shell is very thin then hydrodynamic instabilities are readily able to puncture through the shell, causing loss of symmetry and mixing. A measure of the shell thickness is the aspect ratio, which compares the shell's outer radius to its thickness. In particular, the aspect ratio is evaluated when the outer shell radius is 2/3 of the initial vapour/fuel boundary radius. The aspect ratio evaluated at this time is referred to as the in-flight aspect ratio (IFAR), and is limited to below 30.

The final hydrodynamic constraint is applied to the velocity of the implosion. This is closely linked to the acceleration that is experienced by the shell, which is explicitly included in the hydrodynamic instability growth rates (e.g. Equation 2.74). The implosion velocity is a metric that is often reported in the literature, making it an easier metric to compare against. In the Low-CR regime, the implosion velocity is limited to 400 km/s.

Outside of hydrodynamic behaviour, a limit is also applied to the product of laser intensity and wavelength, $I\lambda^2 < 10^{14}$ Wμm^2/cm^2. This limit ensures that parametric instabilities (Section 2.6.2) are reduced, which can scatter laser energy out of the implosion and generate hot electrons that preheat the fuel, both of which degrade performance.

The threshold values of these constraints have been chosen based on experimental and computational comparisons. In satisfying these constraints implosions are anticipated to be robust to both hydrodynamic and parametric instabilities, where performance is expected to be accurately modelled by 1D hydrodynamic simulations. A summary of these constraints is provided in Table 4.1

Constraint	**Constraint**
Convergence ratio (CR)	CR < 16
In-flight aspect ratio (IFAR)	IFAR < 30
Mass-averaged implosion velocity	$v_{\text{shell}} < 400$ km/s
Laser intensity–wavelength product	$I\lambda^2 < 10^{14}$ Wμm^2/cm^2

Table 4.1: Summary of constraints defining the low convergence ratio implosion regime.

4.3.2 Previous Optimisation Approach

In the previous optimisation attempt in the Low-CR regime, five different capsule sizes, each corresponding to a different energy scale, were explored. Each capsule size was optimised for both a 3-pulse and 4-pulse laser sequence, resulting in a total of 10 optimised target configurations. These multi-pulse laser profiles are similar to those used in experiments, where the changes in laser intensity are used to launch successive shocks into the target. Increasing the number of pulses allows

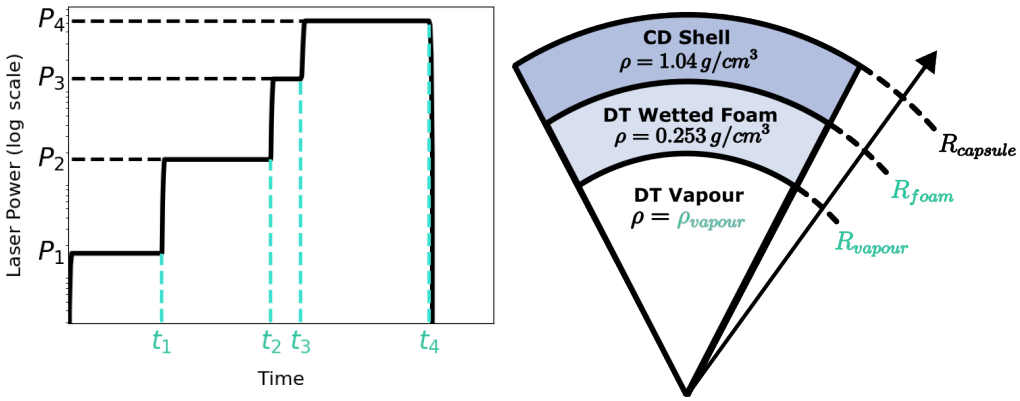


Figure 4.1: 4-pulse laser profile and capsule design construction used in the previous optimisation [137]. Parameters written in turquoise are free parameters to be optimised.

the compression to be distributed over a larger number of shocks, providing greater control over the fuel adiabat and improving compressibility. However, increasing the number of pulses also increases the dimensionality of the parameter space. The construction of the capsule designs and laser pulse profiles are shown in Figure 4.1 with the free parameters to be optimised (laser pulse timings, shell radii, vapour density) indicated in turquoise. Consequently, the 3-pulse optimisations spanned a 6-dimensional (6D) parameter space, and the 4-pulse optimisations were 7D.

The optimisation approach employed began from a 'reasonable' configuration, followed by a sequential linear scan of the parameter space. Simulations in which the Low-CR constraints were not satisfied were disregarded, and the performance metric used for the remaining simulations was the gain (yield/E_{Laser}) of the implosion. The sequential nature of this approach prevents exploration of the interaction between parameters, potentially leading to suboptimal results. Moreover, the laser power of each pulse was fixed and not treated as an optimisation parameter. This severely limits the flexibility of the design space, biasing the optimisation toward a narrow class of pulse shapes and excluding potentially high-performing configurations, such as those associated with shock-augmented ignition [154].

Here a 'reasonable' configuration is determined through hydrodynamic scaling [155]. Hydrodynamic scaling is a method by which the performance of an implosion at one scale can be estimated at a different scale. The Euler equations (Equations 2.76–2.78) are invariant under the transformations $R \to \alpha R$ and $t \to \alpha t$,

with the characteristic implosion velocity u remaining unchanged. Consequently, the imploding mass scales as $M \rightarrow \alpha^3 M$, implying that the laser energy required to drive the implosion scales as $E_{\text{Laser}} \rightarrow \alpha^3 E_{\text{Laser}}$. Implosions of different scales therefore follow similar hydrodynamic behaviour, allowing an optimised configuration at one scale to be used as a starting point for optimisation at a different scale.

Hydrodynamic scaling is commonly used to extrapolate implosion performance to larger energy scales, such as OMEGA implosions at NIF scale energies to predict the potential gain of high-energy direct drive facilities. In the case presented here, a capsule was fully optimised at one scale, requiring many iterations of sequential linear scans. The first optimised configuration was then hydrodynamically scaled to a new scale and used as a starting point for the next optimisation procedure. This required only minor alterations at the new scale to obtain an optimised capsule.

4.3.3 Re-optimisation of Low-CR implosions

The re-optimisation of these configurations introduces several key differences from the previous work. First, the optimised capsule designs from the previous study are retained, and the laser powers are now treated as free parameters, allowing the optimisation process to determine their optimal values. This new optimisation parameter space is shown in Figure 4.2 for the 3- and 4-pulse laser profiles, with the parameters to be optimised again indicated in turquoise. While capsule parameters are no longer being optimised, the dimensionality of the parameter space remains similar, with 6 dimensions in the 3-pulse case, and 8 dimensions for 4-pulses.

In addition to the changes in the optimisation parameter space, several changes have been made in the analysis of the implosions. The implosion velocity has been redefined to be the more physically meaningful mass-averaged velocity of the imploding shell, rather than the volume-averaged velocity used previously. Furthermore, the definition of gain was updated from the neutron fusion energy released (approximately 80% of the total fusion energy for DT reactions) normalised by the laser energy, to the total fusion energy released divided by the laser energy. This is the more common definition of gain found in the literature and remains

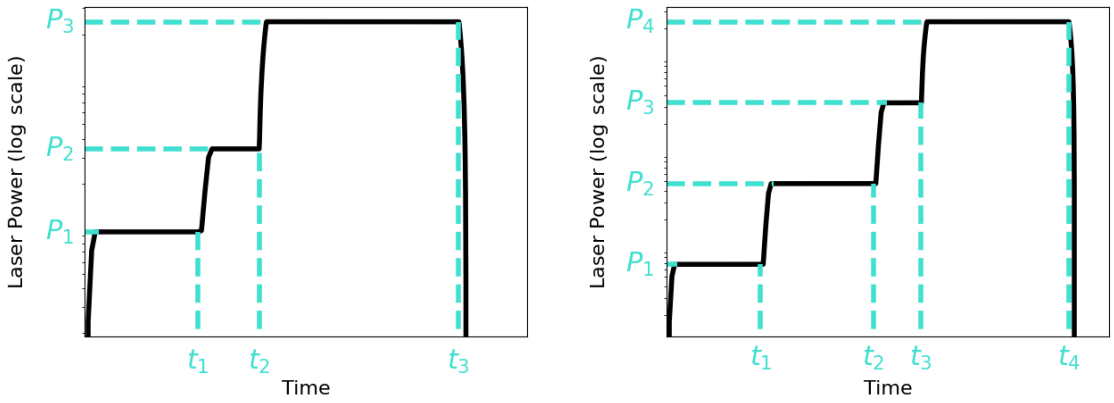


Figure 4.2: General laser pulse profiles to be optimised for the 3- (left) and 4- (right) pulse laser profiles. Parameters to be optimised are indicated in turquoise.

consistent when accounting for fusion reactions beyond DT (such as DD), although these are significantly less frequent, or when using alternate fusion fuel.

It should also be noted that, since the optimisation scanned a large parameter space, for a reasonable number of implosions the CR was infinite, corresponding to cases where the minimum inner shell boundary reached the centre of the implosion. Instead of returning invalid values for these implosions, which can confuse the optimiser, the CR was instead defined to use the outer shell boundary, which is always well-defined, with a threshold value of 13 which was selected empirically. This alternative definition preserves the physical intent of the Low-CR constraint while remaining well defined across the entire parameter space.

Finally, rather than enforcing hard cutoff values for the Low-CR regime parameters, a penalty-based system has been introduced. This allows for smoother exploration of the parameter space, penalising solutions more heavily as they deviate from the Low-CR constraints. This penalty based method is particularly useful due to the ignition instability associated with ICF. Within the Low-CR regime, the implosion may be on the cusp of ignition yet still report little gain. Our approach enables testing the boundaries of the Low-CR regime, potentially surpassing the threshold of ignition slightly and resulting in capsules with significant gain.

The penalty imposed for exceeding the Low-CR thresholds are calculated using

the functional form

$$M = \begin{cases} 1, & X \leq 0 \\ \frac{1}{a}\left[1 - \frac{1}{1+\exp(-(\alpha b X - 3.3))}\right], & X > 0 \end{cases} \tag{4.3}$$

where

$$X = \frac{x - x_{\text{cutoff}}}{x_{\text{cutoff}}}, \tag{4.4}$$

is the deviation of the constraint metric from the threshold value normalised to the threshold value. The constant $a \approx 0.9644$ is used to ensure continuity at $X = 0$, and $b = 52$ was chosen empirically to ensure a suitible decay lengthscale. The α parameter can be used to alter the decay lengthscale, with a larger α resulting in shorter exponential lengthscale and harsher penalty. The value of α is set to 1 for all but the implosion velocity modifier, for which it is set to 2, reflecting the strong influence implosion velocity has on hydrodynamic instability growth. Figure 4.3 illustrates the penalty imposed for exceeding the Low-CR thresholds, using the IFAR as an example.

Taking these penalties into account, the objective function to be optimised is given by

$$\text{Result} = \text{gain} \cdot \text{CR}_M \cdot \text{IFAR}_M \cdot v_M \cdot \left(I\lambda^2\right)_M, \tag{4.5}$$

where subscript M refers to the penalty multiplier corresponding to each constraint.

4.3.4 Performance Comparison

Using the optimisation frameworks described in Chapter 3, all five capsule sizes were re-optimised using both 3- and 4-pulse laser profiles under the revised parameterisation and constraint handling outlined above. Owing to these modifications, the re-optimisation results are not strictly directly comparable to those of the original sequential scan study. The benchmark should therefore be interpreted as a qualitative reference point rather than a controlled one-to-one comparison. Nevertheless, it provides a useful indication that the optimisation frameworks identify plausible optimal implosion configurations within the revised design space.

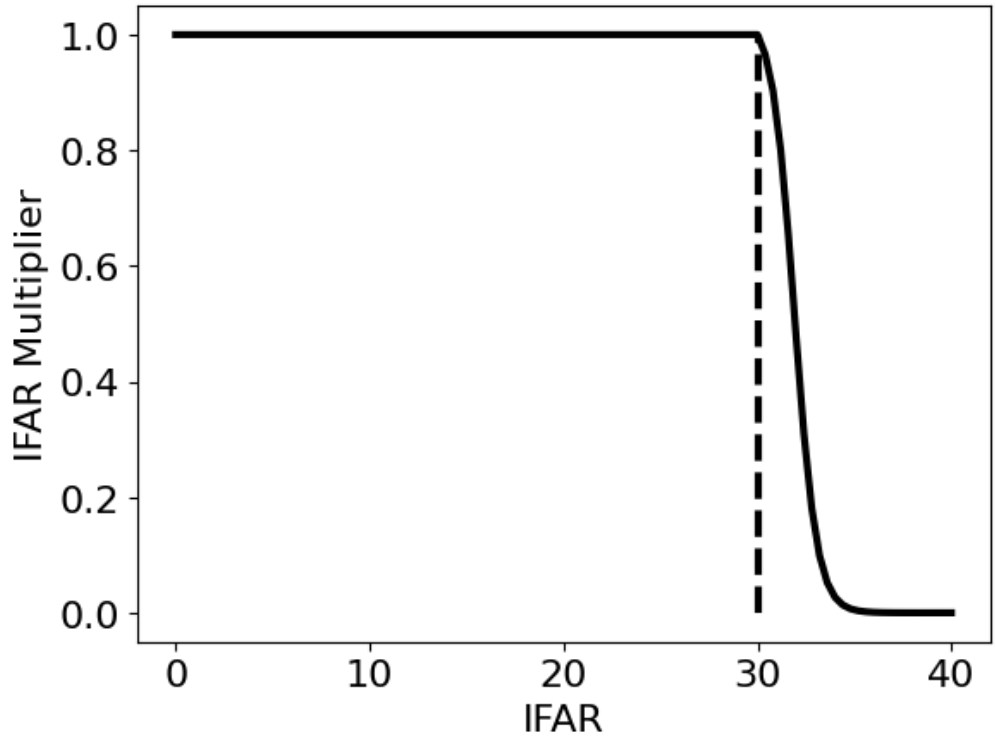


Figure 4.3: Penalty imposed on the gain of the implosion due to exceeding the low-CR IFAR constraint of 30.

The resulting optimal configurations are shown in Figure 4.4, where they are displayed alongside the sequential linear scan results (Equation 4.5) for reference. In all cases, both the Bayesian and particle swarm optimisation identify optima that exceed the results obtained using the sequential scan approach, with comparable final performance between the two advanced methods.

The largest performance improvements are observed for the ~ 2 and ~ 4 MJ implosions, where the achieved gain is improved by a factor of two. These capsules lie close to the ignition threshold, where fusion burn is highly sensitive to small changes in implosion configuration. Ignition is a thermal instability process which, once started, drastically improves the performance of an implosion. There is therefore a stark transition from non-igniting, low-gain implosions to high-gain, igniting configurations. This narrow transition region is referred to as the cliff-edge of ignition, and the ~ 2 and ~ 4 MJ implosions sit within the range of this cliff. It is therefore expected that minor configuration adjustments can lead to drastic

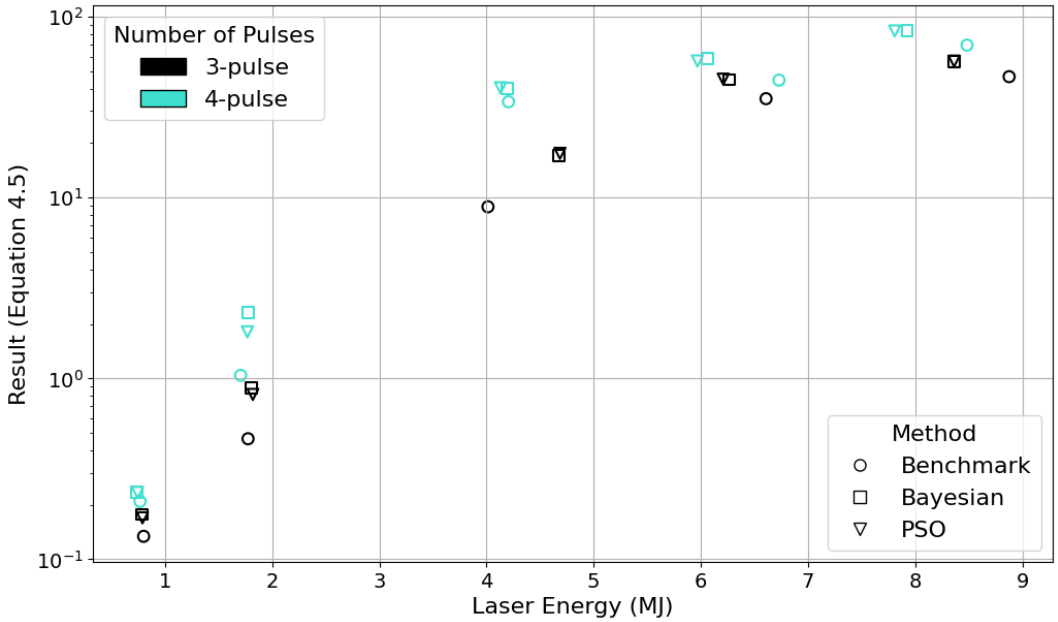


Figure 4.4: Comparison of final results from the sequential linear scan, particle swarm optimisation, and Bayesian optimisation approaches for all five capsule sizes using both 3-pulse and 4-pulse laser profiles. Both advanced optimisation approaches identify improved optimum compared to the sequential linear scan, with comparable performance between the two advanced optimisation techniques.

variations in implosion performance, and capsules operating within this regime are therefore highly sensitive to optimisation improvements.

By contrast, the lowest energy capsule ($\sim$ 0.7 MJ) remains too far from ignition for optimisation to yield substantial gains. Ignition is not possible in the parameter space, and deviations in this space result in minimal change in performance. Similarly, the highest-energy capsules ($\sim$ 6.5 and $\sim$ 8.5 MJ) are already robustly igniting with large burn-up fractions, rendering their performance relatively insensitive to further refinement of the design space.

4.3.5 Evolution Through the Parameter Space

A representative example of the optimisation trajectory for the $\sim$ 2 MJ capsule using a 3-pulse laser profile is shown in Figure 4.5, which displays both the sampled objective values and their locations in the parameter space. The optimised result (Equation 4.5) of the benchmark is indicated by the dashed black line. Distinct behaviours are evident for the two optimisation approaches. In particular, the

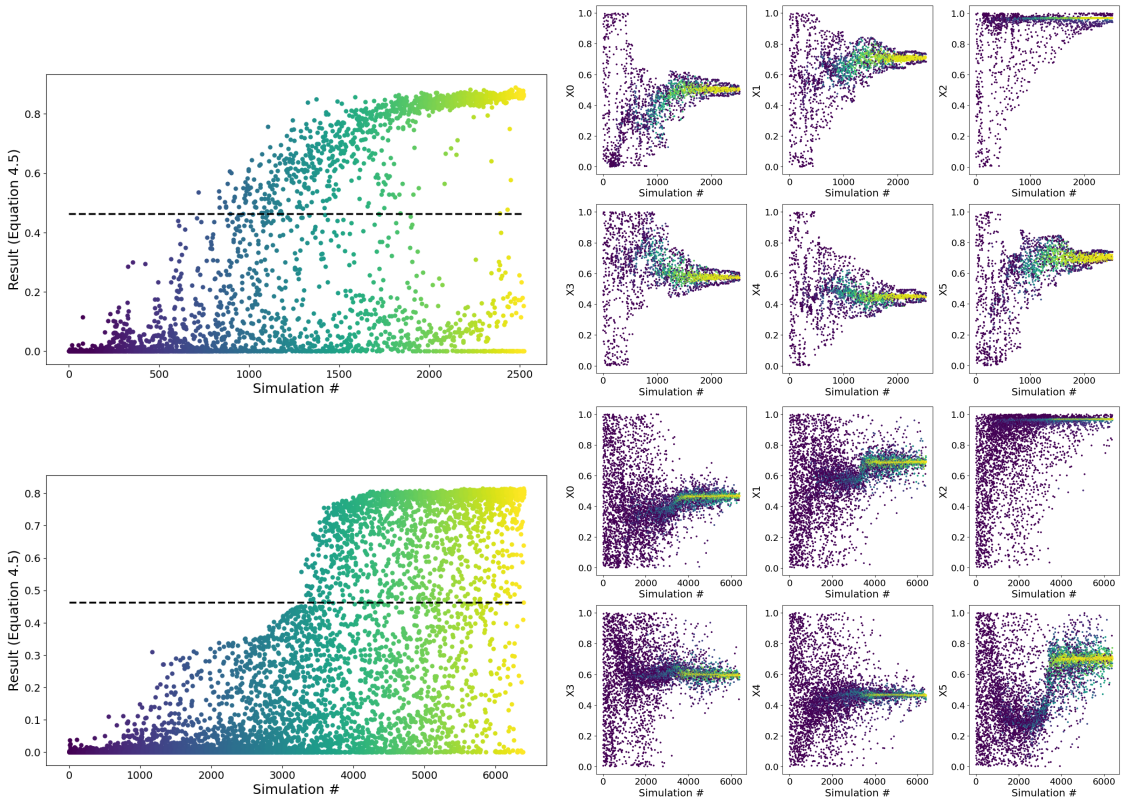


Figure 4.5: Evolution of the optimisation approaches through the parameter space (right) and the associated evaluation values (left). Bayesian (top) exhibits more gradual improvement, while PSO (bottom) converges rapidly on identified optima. The horizontal dashed line indicates the benchmark optimised value.

Bayesian optimiser typically exhibits a slower, more incremental improvement pathway. This behaviour arises from the use of a surrogate model, which typically promotes exploitative sampling through small, local adjustments around the currently identified optimum. The application of dynamic bounds reduction further reinforces this behaviour, as the effective search region contracts smoothly around the best-performing configuration.

The PSO approach contrasts sharply with this behaviour, with the swarm sampling broadly across the parameter space during the early iterations before converging rapidly once a high-performing region is identified. When a distinct region yields improved performance, the swarm quickly reconfigures and coalesces around the newly discovered optimum, as illustrated in Figure 4.5 after $\sim$ 3,000 simulations.

Two key metrics for assessing the performance of these optimisation approaches are the time-to-solution and the computational cost. These two metrics are related,

but have important distinctions. Both approaches utilise parallel processes, where a number of simulations are run concurrently to form a batch of simulations per iteration. The time to solution therefore depends on the number of iterations required to locate the optimum, whereas the computational cost is determined by the total number of simulations performed.

The two approaches differ considerably in this regard. Bayesian optimisation typically prefers smaller batch sizes, with the most efficient approach being single sample batches. This maximises the amount of information available before sampling, leading to a more accurate surrogate model from which samples are determined. PSO sits at the other end of the spectrum, preferring larger numbers of particles in the swarm. This allows the swarm to explore the parameter space thoroughly before settling into an optimum. As the swarm size decreases, the particles become more likely to settle into local optima due to insufficient exploration. As a result, the two metrics, time-to-solution and computational cost, are biased toward a specific approach. PSO often converges to an optimum in fewer iterations, however, requires more simulations to reach this optimum.

In this work, the optimiser is deemed to have identified the optimum once the objective value first reaches 90 % of the final optimum achieved during the run. As can be seen in Figure 4.5, beyond this threshold the sampled points occupy only a very small region of the parameter space, indicating that the optimisation has effectively converged. Differences between configurations beyond this threshold correspond to minor parameter variations and would be indistinguishable within typical shot-to-shot experimental variability. We have verified that adopting alternative thresholds (e.g. 80 % or 95 %) produces negligible changes in the inferred time-to-solution, confirming that the conclusions are not sensitive to the precise numerical choice.

Using these definitions, the number of iterations (time-to-solution) and the total number of simulations (computational cost) for each optimisation approach are summarised in Table 4.2. Each simulation required approximately 3 hours of wall-clock time on a single processor, as the code is not parallelised. With an

average of $2,000$ simulations performed before convergence per optimisation, this corresponds to a total computational cost of $\sim 6,000$ processor-hours per capsule configuration, or $\sim 60,000$ processor-hours for all ten Low-CR re-optimisations. As batches were run concurrently, 35 to 50 iterations corresponds to roughly one working week of wall-clock time per optimisation. Given the substantial duration of each optimisation campaign, the full optimisation process was performed once per capsule configuration. As a result, repeated stochastic restarts and extensive hyper-parameter sensitivity scans were not performed for each capsule configuration. However, both optimisation strategies independently converged to comparable optima across all energy scales, and the optimisation trajectories exhibited stable late-stage convergence behaviour, with sampling concentrated in narrow regions of parameter space once high-performing configurations were identified. Together, these observations suggest that the reported optima are robust within the explored design space.

For comparison, in the benchmark sequential linear scan, approximately 20,000 simulations were required to optimise all 10 implosion designs, corresponding to an average of $\sim$ 2,000 simulations per design. A batch size of 100 simulations was used per iteration, resulting in a total of ~ 20 iterations. It should be noted that the sequential scan was initialised from a reasonable starting configuration and excluded large regions of poorly performing parameter space as a consequence of this choice. In comparison, the optimisation processes presented here started with no prior bias or knowledge, making direct comparison difficult.

Comparisons between the two optimisation strategies indicate that Bayesian optimisation generally requires significantly fewer simulations, while PSO typically converges in fewer iterations, making it the faster method in terms of wall-clock time. Consequently, the preferred approach depends on the optimisation context: Bayesian optimisation is better suited to simulation-limited studies where sample efficiency is critical, whereas PSO is advantageous in scenarios where rapid convergence is prioritised.

		BO		PSO	
Pulse	**Capsule Size**	Iterations	Simulations	Iterations	Simulations
3-pulse	0.5	45	1078	17	897
	0.65	52	1205	56	3067
	0.85	50	1177	13	736
	1.0	35	952	20	821
	1.1	26	713	20	1795
4-pulse	0.5	53	1328	61	5687
	0.65	–	–	66	6162
	0.85	47	1123	29	2482
	1.0	89	2028	41	3441
	1.1	47	1116	35	2904
	Average	49.3	1190	35.8	2800

Table 4.2: Comparison of Bayesian optimisation and PSO convergence performance. The Bayesian approach converges more efficiently, on average in less than half the number of function evaluations, while the PSO approach converges in fewer iterations.

4.3.6 Bayesian Surrogate Model Accuracy

It should be noted that GP-based Bayesian optimisation incurs additional computational overhead associated with surrogate model training and kernel hyper-parameter optimisation. While this overhead is relatively low for the 6- and 8-dimensional parameter spaces considered here, it increases rapidly with dimensionality and becomes a limiting factor beyond $\gtrsim 12$ dimensions. In such higher-dimensional regimes, Gaussian process-based Bayesian optimisation becomes increasingly inefficient, and particle swarm optimisation is therefore the more practical choice. The GP training time during these optimisations scales $\mathcal{O}(n^3)$ with the number of samples, n, and ranged from less than a second in early iterations to approximately 3 minutes in 6D cases and up to ~ 30 minutes in 8D cases after thousands of samples, representing a small fraction of the overall wall-clock time dominated by simulation execution.

The training of a surrogate model does, however, have advantages. By providing an estimate of the objective function across the full parameter space, the surrogate can be used to quantify the stability of the identified optimum, identify alternative optima, flag unphysical solutions, and assess parameter dominance. As an illustrative

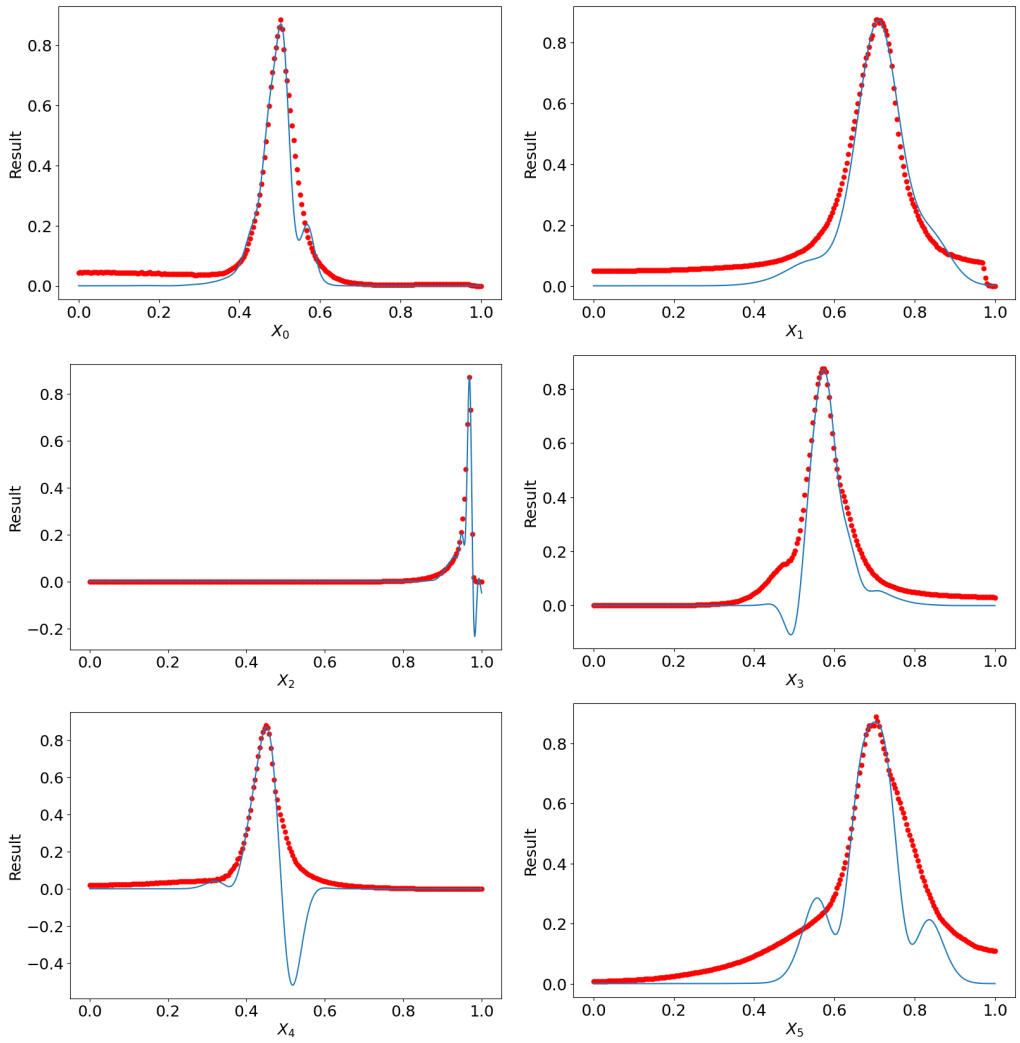


Figure 4.6: Stability analysis of the Bayesian-optimised $\sim$ 0.7 MJ, 3-pulse wetted-foam implosion. The objective function is sampled (red points) along each optimisation parameter axis through the identified optimum, while the corresponding Bayesian surrogate model predictions are shown by the solid blue curves. The width of the optimum along each axis indicates a stable optimum, with the most sensitive parameter (smallest width) given by parameter x_2, corresponding to the power of the final pulse.

example, a stability analysis was performed for the $\sim$ 0.7 MJ, 3-pulse optimum. This analysis consists of sampling the objective function along each parameter axis through the optimal point and comparing these evaluations with the predictions of the Bayesian surrogate model, thereby assessing its local accuracy. The results of this analysis are shown in Figure 4.6.

The surrogate reproduces the objective function well in the vicinity of the optimum, reflecting the high density of samples, and associated high confidence,

accumulated in this region during optimisation. Given the large volume of the parameter space, the surrogate is not expected to provide accurate predictions across the entire domain; rather, its primary utility lies in resolving the local structure around the optimum, where sample density is greatest and optimisation decisions are made. Crucially, the analysis demonstrates that the optimum exhibits significant width along all parameter directions, indicating that the implosion performance is tolerant to realistic experimental variation. The dominant parameter, defined as the parameter with the smallest characteristic lengthscale, is x_2, corresponding to the power of the final laser pulse. This behaviour is physically intuitive, as this parameter controls the strength of the final and strongest shock driven through the target, making the implosion performance particularly sensitive to its value.

Minor discrepancies between the surrogate and sampled objective values are present, most notably along parameters x_4 and x_5. Along x_4, the GP prediction falls below zero. As the objective function is strictly non-negative by construction, this is a result of the unconstrained nature of the GP. Interpreting negative predictions as zero largely removes this discrepancy. Along x_5, the deviation between the GP prediction and sampled values is more pronounced; however, the location and width of the maximum are captured accurately. Since the purpose of this analysis is to assess local stability and parameter sensitivity rather than exact agreement across the entire parameter space, the surrogate model is considered sufficiently accurate in the vicinity of the optimum.

4.3.7 Breakdown of Hydroscaling

A comprehensive summary of the hyper-parameters and results for each optimisation strategy is provided in Appendix Appendix C.1 and C.2. This includes the final laser pulse profiles, Low-CR constraint metrics, and gain. These tables enable direct, quantitative comparison between the different optimisation approaches and clearly illustrate the improvements achieved using Bayesian and particle swarm optimisation.

A key observation that can be made from these results is that the highest-energy implosions ($\sim$ 6.5 MJ and $\sim$ 8.5 MJ) deviate from expectations based on simple hydrodynamic scaling (hydroscaling) [155] of the smaller capsules. Hydroscaling provides a set of scaling relations that allow a capsule and laser drive configuration to be translated between energy scales, and is commonly used to estimate performance on larger laser facilities. However, the optimised pulse shapes for the largest capsules require substantially lower peak laser power (and, to a lesser extent, total laser energy) than would be predicted by directly scaling the smaller configurations.

To investigate this behaviour, test simulations were performed where the 4-pulse, $\sim$ 1.8 MJ configuration identified by the Bayesian optimiser (the highest-performing implosion at this energy scale) was hydroscaled to match the outer radius of the $\sim$ 8 MJ capsule. While the resulting implosion remains viable, achieving an objective function value of 57 compared to 84 for the directly optimised configuration, it is clearly suboptimal. The procedure was then reversed by hydroscaling the optimised $\sim$ 8 MJ configuration down to the radius of the $\sim$ 1.8 MJ capsule. In this case, the degradation is more pronounced, with the objective function decreasing from 2.13 to 0.14. These results indicate that, although hydroscaling can generate reasonable configurations, an optimised solution at one scale cannot be hydroscaled to form an optimised solution at another scale.

A number of factors lead to this deviation from hydroscaling. First, the optimisation objective function incorporates stability constraints and performance metrics that do not scale trivially with the configuration scale. Second, while the hydrodynamic equations can be expressed in a scale-invariant form, additional physics, such as alpha particle deposition, laser energy absorption, and transport processes introduce scale-dependent behaviour. As a result, these observations suggest that while hydroscaling remains a useful guide for estimating performance trends, independent optimisation at each energy scale is necessary to identify truly optimal configurations.

4.4 Fast Ignition Implosion Optimisation

To demonstrate the versatility and broader applicability of advanced optimisation techniques to ICF design, the particle swarm optimisation (PSO) framework was applied to the optimisation of fast ignition (FI) implosions. In contrast to the Low-CR regime, FI is characterised by different objective functions and constraints, providing a distinct optimisation landscape. The dimensionality of the design space, discussed below, however, renders GP-based Bayesian optimisation computationally implausible for this problem, inhibiting its use for this task.

4.4.1 Fast Ignition Implosions

FI is an alternative ICF approach in which the fuel compression and ignition phases are decoupled, offering the potential for higher gain and reduced driver energy requirements. In this scheme, a primary driver compresses the fuel to high areal densities (ρR) after which a secondary, ultra-intense laser pulse generates a beam of energetic particles to rapidly heat the compressed fuel to sufficiently high temperatures to initiate fusion burn.

Achieving high areal density during the compression phase of a fast ignition implosion is critical for both ignition and subsequent burn. Large ρR values ensure that alpha particles produced in fusion reactions redeposit their energy within the fuel, enabling the self-heating required to sustain burn. This behaviour underpins the burn-up fraction, defined as the fraction of fuel undergoing fusion reactions, which can be approximated as [76]

$$\Phi = \frac{\rho R}{\rho R + H_b}, \tag{4.6}$$

where ρR is in units of g/cm^2 and $H_b \approx 7$ g/cm^2. During the ignition phase, high ρR also enhances the collisionality of the fuel, allowing energetic particle beams to deposit energy more efficiently to ignite the fuel. High ρR is therefore essential for both initiating fusion and maximising energy yield in fast ignition.

While high ρR is a necessary condition for effective FI performance, it is not sufficient on its own. Additional constraints must be placed on the density of the

compressed fuel. If the density is too high, energetic particle beams are unable to penetrate the fuel, depositing energy predominantly at the periphery of the compressed region. Conversely, if the density is too low, the energy deposition volume becomes too large, preventing the formation of temperatures sufficient for ignition. Consequently, the compressed fuel density is typically constrained to the range 300-450 g/cm^3 [107, 156–159].

The optimal density within this window depends on the available short-pulse energy. For fixed ρR, increasing the compressed fuel density reduces the characteristic size of the assembly and therefore decreases the total compressed fuel mass, which scales as $M \propto \rho^{-2}$. Since the energy required for ignition is expected to scale approximately with the mass of fuel that must be heated to ignition conditions, the ignition energy requirement similarly decreases with increasing density ($E \propto \rho^{-1.85}$ [160]). The optimal compressed density is therefore facility-dependent and reflects a trade-off between achievable compression, ignition energy requirements, and beam penetration considerations.

4.4.2 Previous Fast Ignition Design Studies

A number of promising fast ignition implosion designs have been explored previously, including the self-similar configurations proposed by Clark and Tabak [161]. While these results are compelling, the optimised laser pulse shapes reported exceed the parametric limit of $I\lambda^2 \leq 10^{14}$ Wμm^2/cm^2. As discussed in the Low-CR context, exceeding this limit increases susceptibility to parametric instabilities and reduces the reliability of one-dimensional radiation-hydrodynamic simulations. Similar issues apply to the optimisation study presented in [107], where a particle swarm approach was used to refine self-similar designs within a highly restricted parameter space.

4.4.3 Optimisation of Fast Ignition Implosions with Improved Parametric Instability Robustness

To address these limitations, the PSO framework presented here is applied to the fast ignition design landscape. The capsule configuration is fixed to that proposed

by Clark and Tabak [161], while the laser pulse profile is parameterised using 8 time-power pairs between which the pulse is linearly interpolated. This parameterisation permits the optimiser to explore a vast volume in configuration space, and does not bias the optimisation process to any preferred pulse shape. An upper bound on laser intensity is set by the constraint $I\lambda^2 \leq 10^{14}$ Wμm^2/cm^2, while the lower bound is chosen to be three orders of magnitude smaller. In addition to the laser intensity constraint, a constraint is also applied to the total laser energy. If the laser energy of a proposed pulse is above 485 kJ, then the whole pulse power profile is reduced so that the pulse energy is 485 kJ. This energy limit fits anticipated FI laser system sizes of around 1 MJ energy [162] being split between both long pulse compression and short pulse ignition laser pulses.

The objective function is comprised of two components: the compressed fuel areal density, ρR, and a density suitability metric designed to penalise density values outside of the 300-450 g/cm^3 range. To identify the compressed fuel region, a shock-finding metric defined in [107] is employed to locate the outgoing shock, which coincides with the outer boundary of the compressed fuel. A shock is characterised by a discontinuous change in material properties such as pressure and density, and its location can be robustly identified by combining the local pressure and density gradients. Evaluating these gradients on a logarithmic scale further enhances the sensitivity to the shock boundary. The shock metric used to identify the index of the outgoing shock, and hence the outer radius of the compressed fuel region, is defined as

$$M = \max\left(\left|\frac{\ln(P_{i+1}) - \ln(P_i)}{r_{i+1} - r_i}\right| \cdot \left|\frac{\ln(\rho_{i+1}) - \ln(\rho_i)}{r_{i+1} - r_i}\right| \cdot m_i\right), \tag{4.7}$$

where P is the pressure, ρ is the density, m is the mass, and subscript i indicates the cell index. The inclusion of the cell mass biases the metric toward larger radii, ensuring the outgoing shock is identified rather than internal shocks. Using this shock index, the compressed fuel ρR can be calculated at each timestep.

A density suitability multiplier is also evaluated to assess whether the compressed fuel lies within the desired density window. For each cell within the compressed region, a target density $\rho_{*,i}$ is assigned. If the cell density lies within the ideal range

of 300-450 g/cm^3, the target density is set equal to the cell density and no penalty is incurred. For densities below 300 g/cm^3 or above 450 g/cm^3, the target density is set to the nearest bound. The suitability penalty, S, is then calculated using a target-density- and cell-width-normalised mean square error:

$$\mathrm{MSE}(M) = \sum_{i=0}^{M} \left(\frac{\rho_i - \rho_{*,i}}{\rho_{*,i}} \right)^2 \frac{(r_{i+1} - r_i)}{r_M}, \qquad S = 1 - \sqrt{\mathrm{MSE}(M)}, \tag{4.8}$$

where r_M denotes the outer radius of the compressed region. The suitability metric is bounded between 0 and 1 (assuming $\rho \leq 900$ g/cm^3), with $S = 1$ corresponding to all compressed fuel lying within the target density range. The identification of the compressed fuel radius, r_M (vertical dotted line), and the target density window (horizontal dotted lines) are illustrated in Figure 4.7(a).

The final objective function combines the compressed fuel areal density with the density suitability metric. To prioritise physically viable compressed density profiles over artificially high ρR values, the suitability metric is squared, biasing the optimisation toward configurations with favourable density distributions:

$$\mathrm{Result} = \rho R \cdot S^2. \tag{4.9}$$

Without this weighting, highly unsuitable density profiles can appear deceptively attractive due to excessively large ρR values alone.

4.4.4 Optimised Laser Pulse and Compressed Fuel Configurations

A swarm consisting of 999 particles were used to explore the 16D parameter space defined by the 8 time-power pairs describing the laser pulse. This is the largest swarm size permitted on the HPC cluster, and was chosen to ensure broad coverage of the high-dimensional design space. The full swarm configuration and hyper-parameter choices are provided in Appendix C.3.

The optimised laser pulse profile and the resulting compressed fuel density distribution are shown in Figure 4.7. In the density profile, the outer shock radius, $r_M = 63.7$ μm, identified using the outgoing shock metric (Equation 4.7), is

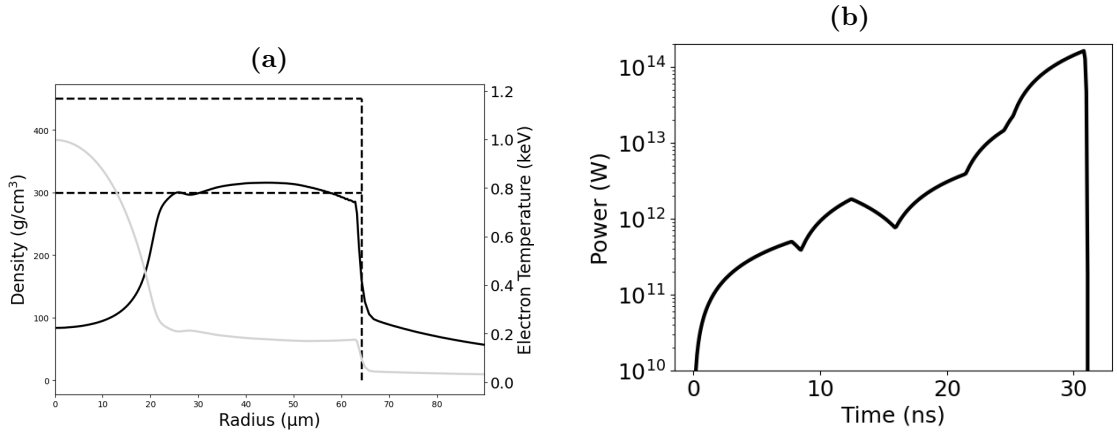


Figure 4.7: (a) Optimised FI density (solid black) and temperature (light grey) of the compressed fuel. (b) Optimised laser pulse profile. The density distribution exhibits a $\sim$ 20 μm hotspot followed by a density profile within the suitable density range of 300-450 g/cm^3, achieving $\rho R = 1.53$ g/cm^2.

indicated by the vertical dotted line. The target density window for fast ignition (300-450 g/cm^3) is indicated by the horizontal dotted lines.

The innermost $\sim$ 20 μm of the compressed fuel lies below the target density threshold of 300 g/cm^3. This behaviour is difficult to avoid under the imposed constraint on the maximum allowable laser intensity. As the capsule is compressed, the central pressure resisting the implosion increases, and once this pressure exceeds the applied drive pressure, the implosion stagnates and subsequently rebounds. The imposed limit on laser intensity therefore constrains the maximum achievable implosion pressure, making it challenging to further increase the central fuel density.

Beyond the central $\sim$ 20 μm, the remaining compressed fuel lies predominantly within the desired density range, with densities clustered around 300 g/cm^3. This region constitutes the majority of the compressed fuel mass, accounting for 98.6 % of the total compressed fuel mass. The resulting compressed areal density is $\rho R = 1.53$ g/cm^2, approaching the 2 g/cm^2 threshold commonly cited [163, 164] as a requirement for high-gain fast ignition designs. If successfully ignited, Equation 4.6 implies that this compressed areal density, together with a total compressed mass of 3.2×10^{-4} g, would correspond to a burn-up fraction of $\sim$ 18 % and a fusion energy release of $\sim$ 19.6 MJ. Assuming the compressed assembly can be ignited using a total laser energy budget of 1 MJ (includes both long and short pulse lasers),

this would imply a gain of ~ 20. It should be emphasised that the optimisation presented here concerns the compression phase only. The generation, transport, divergence, and coupling efficiency of the short-pulse ignition beam are not modelled. This gain estimate therefore represents an implied performance assuming successful ignition is possible with 500 kJ of short pulse energy. Promising routes to further improve this result would be to include the capsule design into the optimisation parameter space, increase the available energy budget, or shifting to a higher frequency laser pulse, such as ArF [142] or KrF [165]. This latter approach is performed and discussed in Section 4.5.

4.4.5 Modifying the Optimum

This laser pulse profile differs reasonably from other presented FI pulse profiles, such as the self-similar profile presented by Clark and Tabak [161]. The optimised laser power profile used to drive the implosion is generally an increasing function, reaching a maximum laser intensity of 1.6×10^{14} W/cm^2, corresponding to $I\lambda^2 = 8 \times 10^{13}$ W μm^2/cm^2, and utilises the full available energy budget of 485 kJ. While the pulse shape is necessarily coarse due to the parameterisation using only 8 time-power pairs, several physically meaningful features can be identified. For much of the drive, the laser power increases gradually, corresponding to ramped compression in which the applied pressure is steadily increased. At ~ 13 ns, the pulse deviates from this trend with a negative power gradient, allowing the plasma to partially relax before being driven again. Since the drive pressure is closely related to the power gradient, this relaxation enables a larger pressure to be applied than would otherwise be possible. This behaviour is qualitatively similar to shock-augmented ignition schemes, in which intentional reductions in laser power are used to facilitate the formation of stronger shocks at later times.

Two additional simulations were performed to analyse the physical importance of the precise structure of the optimised laser pulse. The laser pulses used in these simulations are shown in Figure 4.8. In the first case, the reduction in laser power associated with releasing the plasma was removed, resulting in a

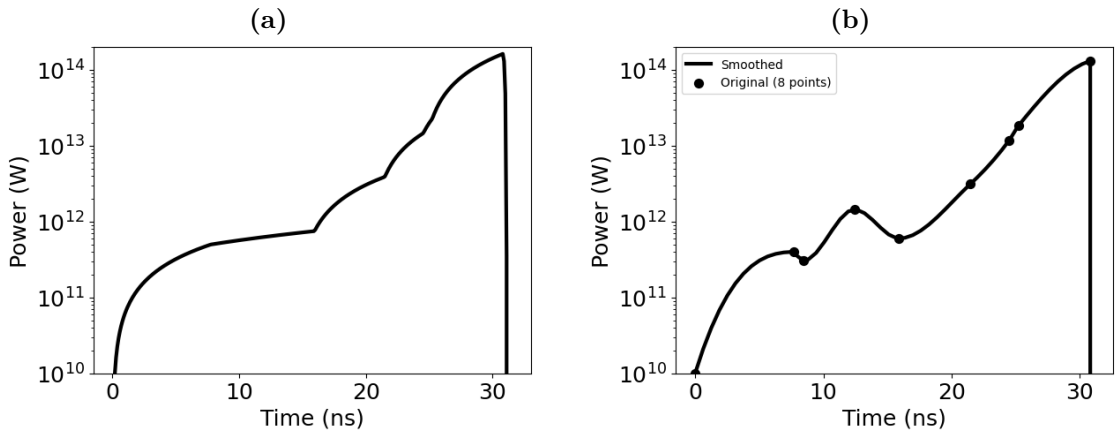


Figure 4.8: Modified optimised laser pulse, with (a) the relaxation feature removed, and (b) a smoothed laser pulse. Both modifications lead to reduced performance.

monotonically increasing drive. In the second case, the pulse was smoothed using spline interpolation through the original time-power pairs, which preserves the structure of the original pulse.

In both cases, implosion performance was degraded relative to the optimised pulse. This sensitivity demonstrates that the relaxation feature identified by the optimiser plays a crucial role in the implosion dynamics. In removing or smoothing this feature the compression of the fuel is hindered. The fact that this feature is not biased or imposed a priori, but instead emerges from the optimisation, highlights the ability of the optimisation process to identify non-intuitive yet physically meaningful drive strategies.

4.5 Extension of FI Implosions to ArF Drivers

4.5.1 Argon Fluoride Laser Technology

Argon fluoride (ArF) lasers are excimer lasers that generate ultraviolet radiation at a wavelength of $\lambda = 193$ nm. In these systems, an electrical discharge excites a mixture of argon and fluorine gases to form a short-lived excited state (ArF*), which subsequently decays and emits coherent radiation. The use of ArF lasers for ICF applications has been heavily considered, especially from the Denver-based fusion start-up Xcimer.

The use of shorter wavelength laser drivers offers several potential advantages for ICF implosions. Most notably, the critical density of a plasma scales inversely with the square of the laser wavelength (Equation 2.51). Consequently, shorter wavelength light can propagate deeper into the plasma corona before reaching the critical surface. This allows laser energy to be deposited closer to the ablation front.

The deeper penetration of the laser light has several important consequences. First, the energy deposited in the corona is more efficiently coupled to the ablation surface, reducing the distance over which electron thermal conduction must transport energy to the shock front. This improves overall laser coupling efficiency and increases the achievable ablation pressure for a given laser intensity. Second, the higher ablation pressure associated with shorter wavelengths enhances ablation stabilisation of hydrodynamic instabilities such as the Rayleigh-Taylor instability, as discussed in Section 2.6.1.

Finally, the threshold for parametric instabilities such as SRS [166], SBS [167], and TPD [168] scale as $I\lambda^2$. As a result, shorter wavelength drivers allow higher laser intensities before parametric instabilities begin to degrade the implosion performance.

4.5.2 Optimisation of Fast Ignition Implosions with ArF Drivers

To investigate the potential advantages of shorter-wavelength drivers, the fast ignition optimisation described in Section 4.4 was repeated using an ArF laser with wavelength $\lambda = 193$ nm. All other aspects of the optimisation framework were held fixed, including the capsule configuration, laser pulse parameterisation, and total laser energy constraint (485 kJ), enabling direct comparison with the baseline $\lambda = 351$ nm case.

The resulting optimised implosion is shown in Figure 4.9. In contrast to the 351 nm case, the central low-density hotspot is eliminated, and the compressed fuel density distribution lies almost entirely within the desired 300-450 g/cm^3 range. This represents a substantial improvement in the physical suitability of the compressed

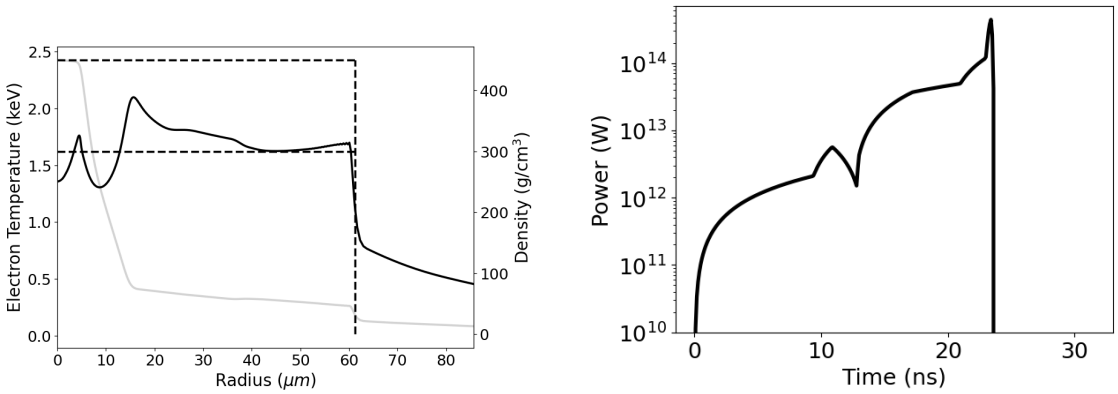


Figure 4.9: (a) Compressed fuel density and temperature distribution using an ArF driver. (b) Optimised laser pulse profile. The outer shock radius and target density window for fast ignition are indicated by the dotted lines. The central hotspot observed in the $\lambda = 351$ nm case is eliminated and the total compressed areal density is increased to $\rho R = 1.9$ g/cm^2.

fuel assembly for fast ignition. The compressed areal density increases from $\rho R = 1.53$ g/cm^2 to $\rho R = 1.90$ g/cm^2, approaching the high-gain threshold of 2 g/cm^2.

This increase in ρR is accompanied by a reduction in the radius of the compressed fuel, which decreases the total compressed fuel mass. While this reduction in mass may be expected to reduce the fusion energy yield, the increase in ρR leads to a higher burn-up fraction. The net effect is an increase in the expected fusion energy release from $\sim$ 19.6 MJ to $\sim$ 21.2 MJ.

It should also be noted that the optimised laser pulse profile for the ArF case is qualitatively similar to the $\lambda = 351$ nm case. The total energy used in both cases is 485 kJ. The pulse duration is slightly shorter, but with a larger peak intensity permitted by the shorter wavelength.

4.6 Summary

Re-optimisation of low-convergence-ratio wetted-foam implosions demonstrated that both Bayesian and PSO approaches outperform traditional sequential scan methods, achieving improved gain while reducing the total number of simulations required. Both advanced optimisation approaches achieved comparable final optima, with Bayesian optimisation requiring fewer evaluations, and particle swarm

optimisation converging more rapidly in terms of iteration count and wall-clock time. These results highlight a clear trade-off between sample efficiency and convergence speed, indicating that the optimal choice of optimiser depends on the available computational resources and the specific optimisation context.

The particle swarm optimisation framework was further applied to a substantially higher-dimensional fast ignition design landscape. In this regime, PSO proved capable of efficiently identifying a physically viable design that achieves a high compressed areal density of 1.5 g/cm^2. This high areal density was achieved using a laser pulse profile under strict intensity and energy constraints.

The fast ignition optimisation was extended to an ArF driver, yielding improved coupling. The central hotspot, which proved unavoidable in the use of a 351 nm driver was eliminated, with an associated increase in the ρR and expected fusion energy release. This result highlights the potential benefits of shorter wavelength drivers for ICF implosions.

While Gaussian process-based Bayesian optimisation becomes increasingly computationally expensive in highly-dimensional parameter spaces, particle swarm optimisation remains well suited to such high-dimensional problems. The complementary strengths of these two approaches suggest that no single optimisation strategy is universally optimal for ICF design, but rather that different regimes benefit from different algorithmic choices.

Looking forward, these optimisation frameworks provide a foundation for more sophisticated design studies, including capsule design, hybrid optimisation strategies, and experimental optimisation on high-power laser facilities. As ICF designs continue to grow in complexity, physics-motivated optimisation methods such as those presented here will play an increasingly important role in navigating complex design landscapes and accelerating progress toward high-performance fusion implosions.

Many of the low-convergence-ratio (Low-CR) implosions presented here, particularly the smaller capsule designs, achieve sufficiently large areal densities (ρR) for significant alpha-particle confinement, but remain below ignition conditions due to insufficient hotspot temperatures. This motivates the exploration of additional

heating mechanisms capable of augmenting the hotspot conditions beyond those achievable through implosion optimisation alone. The following chapters therefore investigate auxiliary heating using relativistic electron beams as a potential route to push these near-igniting implosions across the ignition threshold and improve overall implosion performance.

5

Collisionless Heating of Plasma Using Relativistic Electron Beams

Contents

5.1 Introduction 132
5.2 System Setup, Parameterisation and Linearisation . . 134
5.2.1 Multi-Fluid Theory and Governing Equations 135
5.2.2 Linearisation of the Governing Equations 136
5.2.3 Numerical Identification of Modes 139
5.3 Stable Plasma Modes (No Beam) 140
5.3.1 Characterising Electrostatic and electromagnetic Modes 140
5.3.2 Dispersion Relations and Comparison with Theory . . . 142
5.4 Beam-Plasma System and the Onset of Instabilities . 143
5.4.1 Beam Modes in Isolation 144
5.4.2 Splitting and Instability of the Beam Mode 144
5.4.3 Two-Stream Instability 147
5.4.4 Weibel Instability 147
5.4.5 Oblique Instabilities and Mixed Modes 148
5.4.6 Comparing Instability Growth Rates 149
5.5 Kinetic Effects and Vlasov-Maxwell Simulations 151
5.5.1 Initial Conditions and Simulation Setup 151
5.5.2 Linear Growth via Bump-on-Tail Instability 153
5.5.3 Saturation from Electron Trapping and Dephasing . . . 156
5.5.4 Non-Linear Mode Coupling 159
5.6 Coupling Efficiency 163
5.6.1 Electron Trapping Theory 164
5.6.2 Electron Momentum Evolution 165
5.6.3 Simulation Parameter Scan 166
5.7 Summary 169

5.1 Introduction

High-energy electron beams interacting with plasma are ubiquitous in both astrophysical phenomena and laboratory experiments. In astrophysics, such beams are produced in events like solar flares and coronal mass ejections [169], where magnetic reconnection accelerates electrons that subsequently propagate through the surrounding coronal plasma. Similarly, relativistic jets launched by accreting compact objects can interact with the interstellar medium [170], giving rise to beam-plasma systems. In laboratory settings, high-energy beams can be generated via laser-plasma interactions [171] and synchrotron accelerators [172], providing a platform for investigating beam-plasma dynamics under controlled conditions.

Beam-plasma systems are inherently unstable, with the electron beam acting as a source of free energy that can drive various plasma instabilities [33]. Among these are Weibel-type instabilities [173], which arise from anisotropies in the momentum distribution and generate magnetic fields transverse to the beam propagation direction. Two-stream instabilities, by contrast, are electrostatic in nature, driven by resonant interactions between the beam and background plasma along the direction of propagation [174]. Oblique instabilities occur when the wavevector is angled relative to the beam direction, and can be interpreted as a hybrid between the transverse Weibel and longitudinal two-stream modes [175].

In this chapter, the beam-plasma system is studied both analytically and numerically, using Vlasov-Maxwell simulations. Section 5.2 introduces the system setup and governing equations, followed by a detailed exploration of the initial evolution of the system via linearisation. Then, Section 5.3 explores the simple case of a stable plasma with no electron beams, demonstrating how modes are identified and describing the dispersion relations of the electromagnetic and plasma oscillation modes. Section 5.4 explores the more complex beam-plasma system,

delving into the Weibel, two-stream, and oblique instabilities. The conditions under which instability dominates the evolution of the system are determined.

Previous analytical and numerical studies have established the existence and behaviour of beam-plasma instabilities in a range of plasma environments. The objective of this chapter is therefore not to re-establish these mechanisms, but to extend their application toward collisionless plasma heating for ICF. In particular, the analytical treatment is extended to include electromagnetic modes, allowing the role of the Weibel instability and the transition between instability regimes to be characterised. This analysis is complemented by systematic Vlasov-Maxwell simulations that quantify the nonlinear evolution of the interaction and identify the beam parameters that maximise energy coupling to the background plasma. These results provide the physical foundation for the auxiliary heating scheme investigated in Chapter 6.

Initially, the system evolves through the exponential growth of wave modes. However, this growth eventually saturates. In Section 5.5, Vlasov-Maxwell simulations are introduced to evolve the system beyond the linear regime. The saturation mechanism is found to be due to the trapping of beam electrons in the potential wells of high-amplitude Langmuir waves [176]. In addition to electron trapping, high-amplitude Langmuir waves are capable of coupling with one another to generate new waves [177]. It is shown that this mode coupling results in beam energy being redistributed primarily to the background plasma electrons [178], resulting in an additional stage of background plasma electron heating.

For practical applications of electron beams in plasma heating, a key metric is the maximum fraction of energy transferred from the beam to the plasma, known as the coupling efficiency. Methods of predicting this coupling efficiency are presented in Section 5.6. Two approaches are discussed and compared with numerical simulations to determine the accuracy of their predictive capability. Finally, conclusions about the work presented in the chapter are given in Section 5.7.

5.2 System Setup, Parameterisation and Linearisation

To explore the dynamics of a beam-plasma system in a generalisable framework, consider a system consisting of two relativistic electron beams propagating through a background plasma. The background plasma, consisting of protons and electrons, is assumed to be spatially uniform. The electron beams are likewise uniform, with drift velocities $\mathbf{v}_1$ and $\mathbf{v}_2$. Since the beam velocities are relativistic, the electron momentum is given by $\mathbf{p} = \gamma m \mathbf{v}$, where the Lorentz factor is defined as $\gamma = (1 - v^2/c^2)^{-1/2}$.

The electron beams have number densities n_1 and n_2. To ensure quasi-neutrality, the background plasma has ion and electron densities given by $n_i = n_0$ and $n_e = n_0 - n_1 - n_2$ respectively. The ions are assumed to be immobile, and the background electrons are initialised with a velocity $\mathbf{v}_0$ such that the total system is neutral, with no net current.

The system is fully described by 12 parameters, consisting of the number densities and three-velocities of each of the three electron species: the background electrons and the two beams. These can be collected into a single parameter vector:

$$\mathbf{U}_0 = (n_e, \mathbf{v}_0,\, n_1, \mathbf{v}_1,\, n_2, \mathbf{v}_2) = (n_e, v_{x0}, v_{y0}, v_{z0},\, n_1, v_{x1}, v_{y1}, v_{z1},\, n_2, v_{x2}, v_{y2}, v_{z2})\,. \tag{5.1}$$

This 12-dimensional parameter vector completely defines the initial state of the system and forms the basis for the linear stability analysis presented in the following sections, where each species is treated as a separate fluid. The immobile ions provide a static neutralising background and do not contribute to the variables defining the system.

This configuration is highly versatile: setting the components of both beams to zero yields a stable plasma, while removing just one yields a single-beam setup. Varying the beam propagation angles allows exploration of different interaction geometries.

5.2.1 Multi-Fluid Theory and Governing Equations

A relativistic multi-fluid description of beam-plasma systems has previously been developed in [179], though that work focused exclusively on electrostatic interactions. In contrast, the analysis presented here extends the framework to a fully three-dimensional system that includes electromagnetic effects. This generalisation enables the exploration of a broader class of plasma modes, including electromagnetic instabilities such as those discussed in Sections 5.4.4 and 5.4.5.

To describe the dynamics of the system, the self-consistent coupling between the evolution of the electromagnetic fields and the motion of individual plasma species must be captured. This is achieved by combining Maxwell's equations with a set of fluid equations for each charged component. These governing equations are outlined below.

Maxwell's Equations

The system involves the interaction of charged particles and the associated electric and magnetic fields. As such, Maxwell's equations must be satisfied throughout the system's evolution. They are given by:

$$\nabla \cdot \mathbf{E} = \frac{\rho}{\epsilon_0}, \tag{5.2}$$

$$\nabla \cdot \mathbf{B} = 0, \tag{5.3}$$

$$\nabla \times \mathbf{E} = -\frac{\partial \mathbf{B}}{\partial t}, \tag{5.4}$$

$$\nabla \times \mathbf{B} = \mu_0 \mathbf{J} + \frac{1}{c^2}\frac{\partial \mathbf{E}}{\partial t}, \tag{5.5}$$

where $\mathbf{E}$ and $\mathbf{B}$ are the electric and magnetic fields; ρ is the charge density; and $\mathbf{J}$ is the current density.

Fluid Equations

To fully describe the plasma dynamics, Maxwell's equations must be coupled to a set of fluid equations governing the motion of each species. Each species can be modelled

as a distinct but interacting fluid within a multi-fluid framework. The motion of each species is governed by the continuity equation and the relativistic equation of motion:

$$\frac{\partial n_\alpha}{\partial t} + \nabla \cdot (n_\alpha \mathbf{v}_\alpha) = 0, \tag{5.6}$$

$$\left(\frac{\partial}{\partial t} + \mathbf{v}_\alpha \cdot \nabla\right)(\gamma_\alpha \mathbf{v}_\alpha) = \frac{q_\alpha}{m_\alpha}\left(\mathbf{E} + \mathbf{v}_\alpha \times \mathbf{B}\right), \tag{5.7}$$

where n_α, $\mathbf{v}_\alpha$, q_α, and m_α are the number density, velocity, charge, and mass of species α, and γ_α is the corresponding Lorentz factor.

These equations assume the system is collisionless and that the fluids are cold, meaning they have no thermal velocity spread. Particles within each species are therefore considered monoenergetic, represented by delta functions in velocity space.

In this framework, the total charge and current densities are expressed as sums over the individual species:

$$\mathbf{J} = \sum_{\alpha=0}^{2} q_\alpha n_\alpha \mathbf{v}_\alpha, \tag{5.8}$$

$$\rho = \sum_{\alpha=0}^{2} q_\alpha n_\alpha, \tag{5.9}$$

and are zero in the unperturbed system, i.e., $\mathbf{J}(t = 0) = 0$ and $\rho(t = 0) = 0$. Charge and current neutrality in the initial state ensures the absence of any pre-existing electric or magnetic fields.

5.2.2 Linearisation of the Governing Equations

Linearised Equations in Normalised Units

The process of linearisation, as described in Section 2.2.3, involves introducing small perturbations to the equilibrium state of the system and describing how these perturbations evolve under a plane wave ansatz. This simplifies the governing equations to a linearised form, which can be further simplified by adopting a set of natural units based on the plasma properties. Time is measured in units of the inverse plasma frequency, ω_p^{-1}, and speed is normalised to the speed of light, c, resulting in distance in units of c/ω_p. Mass and charge are normalised to the electron mass m_e and elementary charge e, respectively, while the number density is expressed

relative to the background plasma density n_0. With these choices, key physical constants simplify to: $\epsilon_0 = \mu_0 = c = 1$, $q_i = -q_e = 1$, $m_e = 1$ and $v \to \beta = v/c$.

Using this unit system and substituting the perturbed quantities into Maxwell's equations, the following linearised equations are obtained:

$$i\mathbf{k} \cdot \mathbf{E}_1 = \rho_1, \tag{5.10}$$

$$i\mathbf{k} \cdot \mathbf{B}_1 = 0, \tag{5.11}$$

$$i\mathbf{k} \times \mathbf{E}_1 = i\omega \mathbf{B}_1, \tag{5.12}$$

$$i\mathbf{k} \times \mathbf{B}_1 = \mathbf{J}_1 + (-i\omega)\mathbf{E}_1. \tag{5.13}$$

To linearise the fluid equation of motion, the Lorentz gamma factor must first be expanded to first order. This expansion is shown in Appendix D.1, and the result is given here along with the linearised fluid equations:

$$\gamma(\beta) = \gamma_0 + \gamma_1 = \gamma_0 + (\beta_0 \cdot \beta_1)\gamma_0^3, \tag{5.14}$$

$$(-i\omega)n_{\alpha,1} + i\mathbf{k} \cdot (n_{\alpha,0}\beta_{\alpha,1} + n_{\alpha,1}\beta_{\alpha,0}) = 0, \tag{5.15}$$

$$(-i\omega + \beta_{\alpha,0} \cdot i\mathbf{k})(\gamma_{\alpha,0}\beta_{\alpha,1} + \gamma_{\alpha,1}\beta_{\alpha,0}) = -\mathbf{E}_1 - \beta_{\alpha,0} \times \mathbf{B}_1. \tag{5.16}$$

Each species contributes a set of linearised equations. Since we consider three electron populations (background + two beams), and each equation has vector components, we obtain a coupled system with:

- 18 unknowns: three perturbed densities, three perturbed velocities per species ($3 \times 3 = 9$), plus 3 electric field components and 3 magnetic field components.
- 20 equations: three momentum equations per species (9 total), three continuity equations, and eight Maxwell equations.

However, two of Maxwell's equations (Gauss's laws, Equations 5.10 and 5.11) are constraint relations rather than evolution equations. These are automatically satisfied if the initial conditions and curl equations are consistent. Thus, the dynamical system is fully described by the remaining 18 equations.

This closed linear system governs the evolution of small perturbations in the beam-plasma system. It is possible to now recast this system in matrix form, allowing for a numerical approach to determine the supported wave modes and their stability.

Matrix Representation and Dispersion Condition

After linearisation, each equation contains only first-order terms in the perturbed quantities, and can be written as a linear combination of those perturbations:

$$\mathbf{m}_i \cdot \mathbf{u} = 0, \tag{5.17}$$

where $\mathbf{m}_i$ is a row vector containing the coefficients of the perturbed quantities, and $\mathbf{u} = \left[n_{0,1}, \boldsymbol{\beta}_{0,1}, n_{1,1}, \boldsymbol{\beta}_{1,1}, n_{2,1}, \boldsymbol{\beta}_{2,1}, \mathbf{E}_1, \mathbf{B}_1\right]^T$ is the column vector of perturbed quantities.

As an illustrative example, consider the x-component of the linearised Faraday law (Equation 5.12):

$$i\left(\mathbf{k} \times \mathbf{E}_1\right)_x - i\omega B_{x,1} = 0, \tag{5.18}$$

$$i\left(k_y E_{z,1} - k_z E_{y,1}\right) - i\omega B_{x,1} = 0, \tag{5.19}$$

$$(-ik_z)E_{y,1} + (ik_y)E_{z,1} + (-i\omega)B_{x,1} = 0, \tag{5.20}$$

$$[\,0,\,0,\,0,\,0,\,0,\,0,\,0,\,0,\,0,\,0,\,0,\,0,\,0,\,-ik_z,\,ik_y,\,-i\omega,\,0,\,0\,] \cdot \mathbf{u} = 0, \tag{5.21}$$

where the nonzero entries correspond to $E_{y,1}$ (14th), $E_{z,1}$ (15th), and $B_{x,1}$ (16th) element in the vector $\mathbf{u}$.

By applying this procedure to all of the linearised equations, each can be expressed in the form of a row vector multiplied by the column vector $\mathbf{u}$. By stacking these equations together, the full linearised system can be written compactly as a matrix equation:

$$\mathbf{M} \cdot \mathbf{u} = \mathbf{0}, \tag{5.22}$$

where $\mathbf{M}$ is the matrix formed by collecting all the coefficient row vectors $\mathbf{m}_i$ into rows. The full matrix, $\mathbf{M}$, is therefore 18×18 and is given in Appendix D.2, with the example in Equation 5.21 (Faraday's x-component) corresponding to row $i = 16$ of $\mathbf{M}$.

This equation is always valid for $\mathbf{u} = 0$, referred to as the trivial solution, which corresponds to the unperturbed state of the system. However, the interest lies in

finding non-trivial solutions which correspond to perturbed states supported by the system. These non-trivial (null vector) solutions exist only when the matrix $\mathbf{M}$ is singular, which corresponds to when $\det \mathbf{M} = 0$. Physically, this condition determines the dispersion relation of the system, which relates the frequency ω and wavevector $\mathbf{k}$ of any supported wave mode.

This matrix approach offers a powerful way to analyse wave propagation and stability: given any set of background parameters and a chosen wavevector $\mathbf{k}$, one can search for the frequencies ω that satisfy $\det \mathbf{M} = 0$. The corresponding null vectors $\mathbf{u}$ reveal the structure of each mode, showing which physical quantities are involved in the instability or wave dynamics.

5.2.3 Numerical Identification of Modes

A mode is a plane wave perturbation characterised by a pair $(\mathbf{k}, \omega)$ and a corresponding null vector $\mathbf{u}$, which satisfies the linear system $\mathbf{M}(\mathbf{k}, \omega) \cdot \mathbf{u} = 0$. Non-trivial solutions exist only when the determinant of the coefficient matrix vanishes, i.e., $\det \mathbf{M} = 0$, which defines the dispersion relation of the system.

Modes can therefore be identified numerically by scanning through $(\mathbf{k}, \omega)$, calculating the determinant of the matrix $\mathbf{M}$, and locating the points where the determinant falls to zero. Each of these points (solutions) corresponds to a wave supported by the system, with the pair $(\mathbf{k}, \omega)$ describing the spatial and temporal structure of the perturbation and the null vector $\mathbf{u}$ specifying the relative amplitudes and phases of the perturbed quantities involved in the mode.

Particularly interesting modes are those that are unstable, that is, modes that grow exponentially in amplitude over time. To understand instability, recall that a perturbation takes the form:

$$\mathbf{A}_1 \propto \mathbf{u} \exp\left[i(\mathbf{k} \cdot \mathbf{x} - \omega t)\right], \tag{5.23}$$

where $\mathbf{u}$ encodes the relative amplitudes and phases of the perturbed quantities. If the frequency ω has a non-zero imaginary component, then the perturba-

tion evolves as

$$\exp(-i\omega t) = \exp(-i\omega_r t)\exp(\omega_i t). \tag{5.24}$$

Hence, when $\omega_i > 0$ the mode grows exponentially in time. Conversely, $\omega_i < 0$ leads to decay, while purely real ω ($\omega_i = 0$) describes a stable, undamped wave.

For a given wavevector $\mathbf{k}$, modes with the largest $\mathrm{Im}(\omega)$ grow the fastest and dominate at late times, and are therefore of primary interest.

5.3 Stable Plasma Modes (No Beam)

Before studying the influence of electron beams, it is useful to first analyse the background plasma in isolation. This provides a baseline case consisting of a uniform, cold, collisionless, and quasi-neutral plasma, with equal densities of electrons and immobile ions. The system supports only stable modes and will serve as a reference for identifying beam-driven instabilities later.

The system has no net charge or current. Using the notation introduced in Equation 5.1, the parameter vector simplifies to:

$$\mathbf{U}_0 = (n_0,\, 0,\, 0,\, 0,\, 0,\, 0,\, 0,\, 0,\, 0,\, 0,\, 0,\, 0)\,. \tag{5.25}$$

Since this system is isotropic, the direction of the wavevector $\mathbf{k}$ is arbitrary. As a result, the wave modes can be characterised entirely by the magnitude of the wavevector, $k = |\mathbf{k}|$.

5.3.1 Characterising Electrostatic and electromagnetic Modes

To identify modes numerically, the determinant of the coefficient matrix $\mathbf{M}(\mathbf{k}, \omega)$ is evaluated over a grid in the complex frequency plane, $\omega = \omega_r + i\omega_i$, for a fixed wavevector magnitude. Figure 5.1 shows a contour plot of the determinant of $\mathbf{M}$ as a function of ω_r and ω_i for $k = 2$, with local minima ($\det \mathbf{M} = 0$) indicated by red points. These minima indicate the complex frequencies of the modes supported by the plasma at the chosen wavevector.

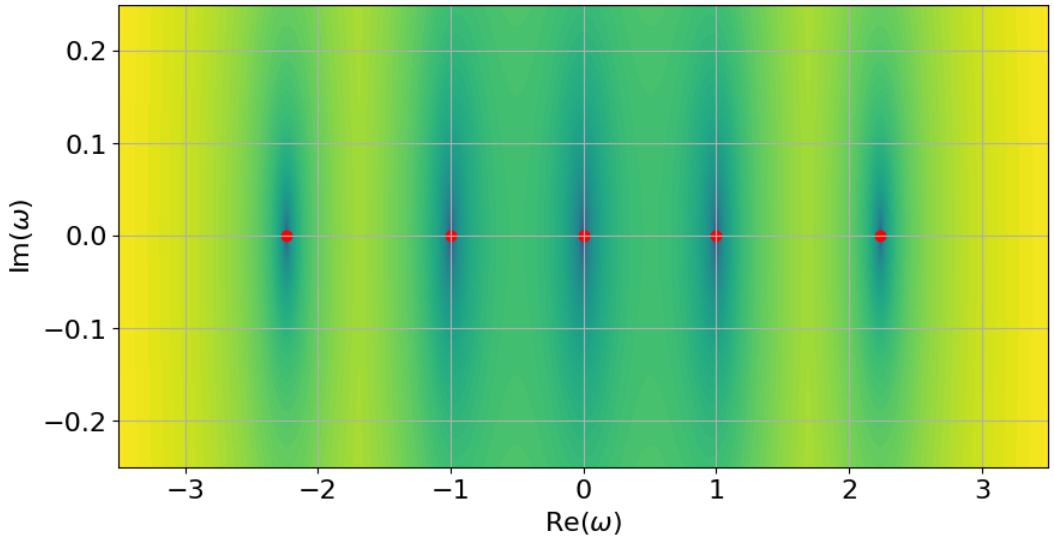


Figure 5.1: The determinant of the coefficient matrix $\mathbf{M}$ as a function of the real and imaginary parts of the frequency ω for a stable plasma with $k = 2$. The red dots indicate the local minima, corresponding to the complex frequencies of the modes supported by the plasma.

It is no surprise that the modes supported by a stable plasma are all stable, with $\omega_i = 0$. While there appear to be distinct modes with positive and negative ω_r, they are not physically independent. The negative modes appear due to the system's symmetry under time reversal. There are therefore three distinct modes that can be identified in this case at $(k, \omega_r) = (2, 0)$, $(2, 1)$, and $(2, 2.24)$.

Using these mode $(\mathbf{k}, \omega)$ pairs, the equation $\mathbf{M} \cdot \mathbf{u} = 0$ can be solved numerically to find the null vector $\mathbf{u}$ corresponding to each supported mode. Each element of $\mathbf{u}$ is complex, encoding both the amplitude and relative phase of its associated quantity. This vector therefore reveals the physical structure of the mode, describing which quantities are oscillating, and how they are coupled.

For a wavevector $\mathbf{k} = (2, 0, 0)$, the null vector associated with the $(2, 1)$ mode has non-zero amplitude for the density, and the x-components of velocity and electric field. This corresponds to plasma oscillations: particles displaced along $\mathbf{k}$ create a restoring electric field due to the resulting charge imbalance. The particle motion is then a simple harmonic oscillation about the equilibrium: an electrostatic mode, as described in Section 2.2.3.

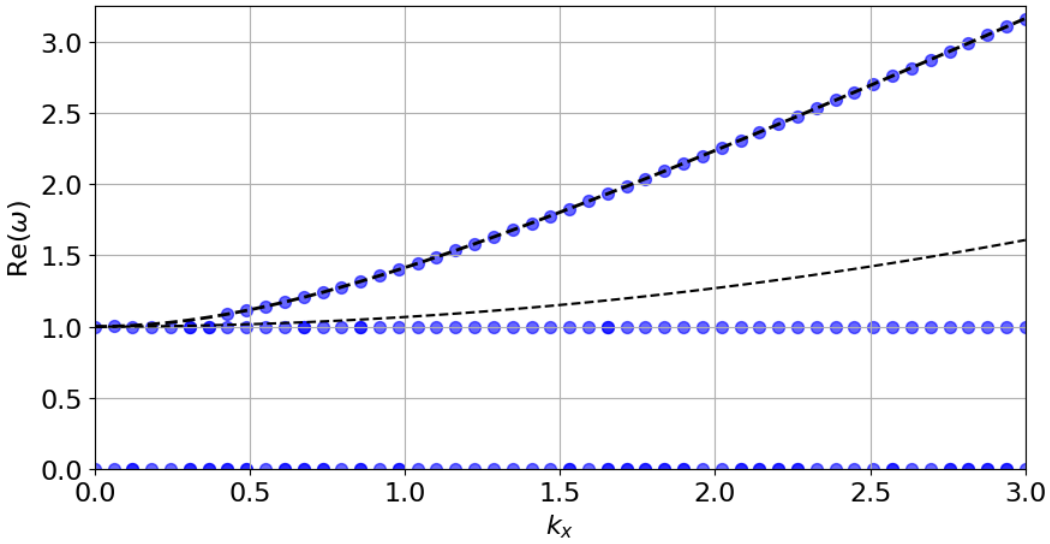


Figure 5.2: Dispersion curves indicating how the frequency of the modes supported by the system vary with k. Two modes are tracked, the electromagnetic mode (upper curve) corresponding to the propagation of light, and the plasma oscillation mode (lower curve). Plots are included comparing these modes to analytic results. The plasma oscillation modes differ from the analytic dispersion relation due to the assumption of a cold plasma.

The second mode, at $(2, 2.24)$, has a different structure. For $\mathbf{k} \parallel \hat{x}$, the null vector has non-zero components in the transverse directions (v_y, v_z, E_y, E_z, B_y, B_z). This mode involves both electric and magnetic field oscillations transverse to the direction of propagation, and corresponds to the propagation of electromagnetic waves through the plasma: a light wave modified by plasma effects, as introduced in Section 2.3.4.

5.3.2 Dispersion Relations and Comparison with Theory

Having identified the supported modes at a fixed wavenumber and determined the type of wave they represent, it is also interesting to see how their frequencies vary as the magnitude of the wavevector changes. This constructs the dispersion relations for each mode, which can be compare to theoretical expectations. Figure 5.2 shows the real part of the frequency, ω_r, for the two supported modes (the electrostatic plasma oscillation and the electromagnetic wave) as a function of wavevector magnitude $k = |\mathbf{k}|$.

The plasma oscillation mode remains at a constant frequency $\omega = \omega_p$, independent of k. This is a characteristic feature of cold, collisionless plasmas where thermal

effects are neglected. Since the pressure gradient (i.e., thermal) term is absent from the fluid equation of motion, the oscillation frequency does not depend on wavelength. In more realistic settings where finite electron temperature is included, the pressure term modifies the restoring force and introduces a k-dependence and leads to the Bohm-Gross dispersion relation (Equation 2.28). This dispersion relation has also been plotted in Figure 5.2 using a thermal velocity of $0.15c$ for comparison.

In contrast, the second mode (the electromagnetic wave) exhibits a clear dispersion. As shown in Figure 5.2, its frequency increases with k, asymptotically approaching the vacuum light line $\omega = kc$ at high wavenumber. This behaviour reflects the fact that, at small wavelengths (large k), the plasma becomes increasingly transparent to electromagnetic waves. This dispersion behaviour is consistent with the well-known cold plasma dispersion relation for electromagnetic waves (Equation 2.48), which has also been plotted for comparison.

5.4 Beam-Plasma System and the Onset of Instabilities

Having understood the identification and tracking of wave modes in the stable background plasma, consider the more dynamically rich system involving a relativistic electron beam propagating through a background plasma. This configuration introduces strong sources of free energy that can destabilise the system, giving rise to a range of growing modes.

The system is initialised with the following parameter vector:

$$\mathbf{U}_0 = \left(0.99,\, -\frac{0.01 \times 0.6}{0.99},\, 0,\, 0,\, 0.01,\, 0.6,\, 0,\, 0,\, 0,\, 0,\, 0,\, 0\right). \tag{5.26}$$

which describes a beam with a density of 1% the background plasma ion density and drift velocity of $0.6c$, corresponding to a kinetic energy just above 0.1 MeV. The background plasma electron density and drift velocity have been modified accordingly to ensure there is no net current or charge.

The addition of the beam breaks the isotropy of the background system and introduces a directional bias. This enables resonant interactions between the beam

and plasma, which can destabilise certain wave modes. As a result, instabilities may develop that grow exponentially in time, driven by the kinetic energy of the beam.

5.4.1 Beam Modes in Isolation

In the presence of a single electron beam drifting with velocity $\mathbf{v}_b$, the system can support a class of modes known as beam modes. These are simple solutions in which a perturbation propagates at the same speed as the beam itself, without distortion. Mathematically, they correspond to a dispersion relation of the form:

$$\omega = \mathbf{k} \cdot \mathbf{v}_b. \tag{5.27}$$

Physically, these modes represent fluctuations that are locked to the beam, where a local density or current perturbation that simply travels with the beam velocity. In the absence of any coupling to the background plasma, beam modes are neutral: they do not grow or decay in time, and manifest as purely real-frequency solutions

5.4.2 Splitting and Instability of the Beam Mode

When the beam interacts with a plasma, the beam mode does not appear as a single isolated root. Instead, a cluster of nearby solutions emerges, centred around the resonant condition $\omega = \mathbf{k} \cdot \mathbf{v}_b$. These solutions reflect the coupling between the beam and background plasma, which splits the beam mode into multiple hybrid modes with varying stability.

Interestingly, the structure of the mode cluster also evolves. At low k, the dominant splitting occurs along the imaginary axis producing complex-conjugate root pairs with non-zero imaginary components, where modes share the same real frequency but differ in growth rate. These represent unstable and damped modes, and their presence is a direct signature of energy exchange between the beam and background. As the beam modes approach the plasma oscillation mode, the modes strongly couple, leading to large imaginary components and hence large growth rates. As k increases beyond the plasma oscillation mode, the splitting transitions to the real axis, where the modes are stable but differ in oscillation frequency.

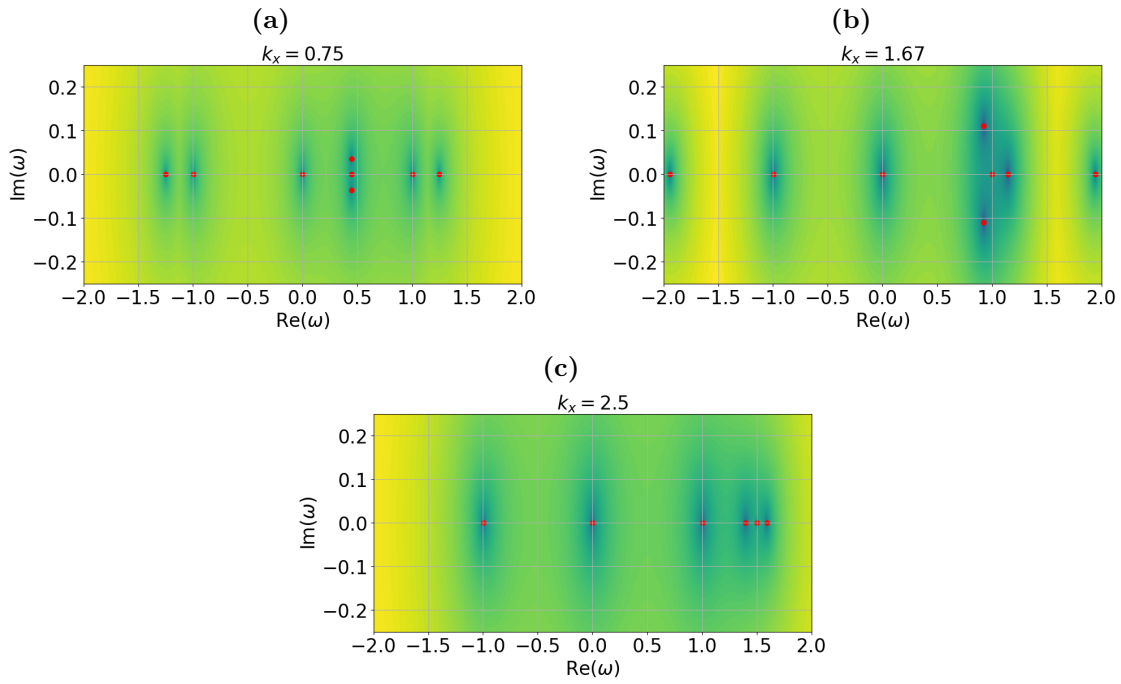


Figure 5.3: Evolution of the beam mode structure as a function of wavenumber k. Left: At low k, the beam mode splits along the imaginary axis, indicating instability. Centre: Near the plasma oscillation frequency, strong coupling produces large growth rates. Right: At higher k, the roots collapse onto the real axis and become stable, with distinct oscillation frequencies.

These three stages: imaginary-axis splitting at low k, strong coupling near the plasma oscillation frequency, and real-axis splitting at high k, are illustrated in Figure 5.3. The first panel shows the mode structure at low k, where the beam mode splits into complex-conjugate pairs, indicating unstable and damped modes. The second panel corresponds to the region where the beam mode interacts with the plasma oscillation mode, resulting in the largest imaginary components and maximal growth. The bottom panel shows the high-k regime, where the roots have collapsed onto the real axis and represent stable, propagating modes with differing frequencies. The progression of these plots reveals the continuous evolution of the beam mode as a function of wavenumber, and clearly demonstrates the transition from instability to stability.

Again, the beam modes can be tracked as k is varied. The mapping of the modes through (k, ω_r) and (k, ω_i) are shown in the left image of Figure 5.4. This behaviour of beam mode splitting can be determined analytically. Under the electrostatic

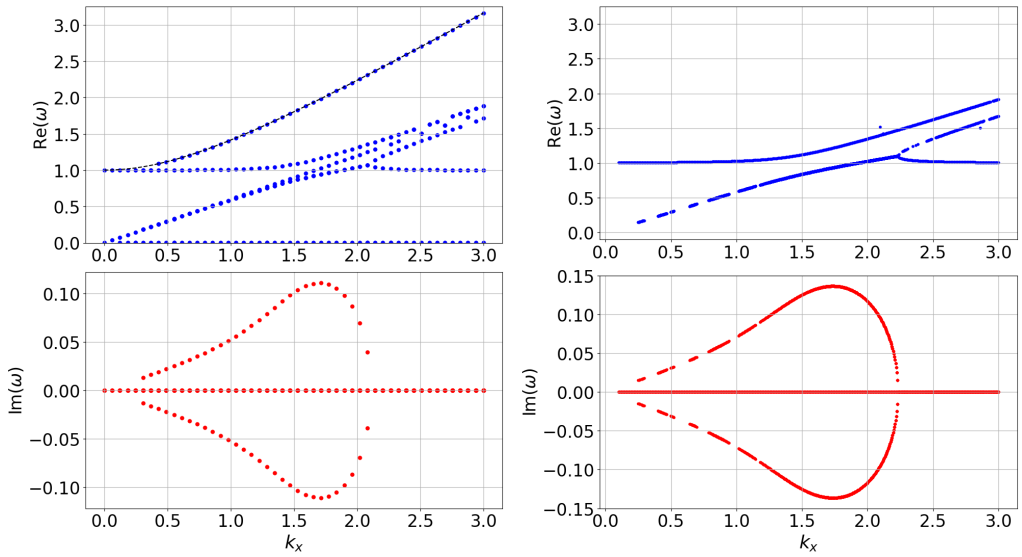


Figure 5.4: (Left) Numeric dispersion curves for the beam-plasma system. Beam modes are clearly visible, which interact with the plasma oscillation mode. These modes have non-zero imaginary components indicating damped (Im < 0) or growing (Im > 0) modes. (Right) Analytic dispersion curves for electrostatic modes under the cold plasma approximation. The analytic solution omits electromagnetic modes, such as the light wave and central beam mode.

and cold plasma assumptions, the dispersion relation [180]

$$1 - \frac{\omega_{p0}^2}{\omega^2} - \frac{\omega_{p1}^2}{(\omega - \mathbf{k} \cdot \beta_\mathbf{1})} = 0. \tag{5.28}$$

can be derived. Solutions to this equation are also plotted in the image on the right hand side of Figure 5.4. Note, that there are a few differences between these images. First, the light mode is missing from the analytic (right) dispersion relation. In addition, only two beam-related roots appear in the analytic dispersion relation, corresponding to the electrostatic side modes. The additional light mode and central beam mode, located at the phase velocity $\omega/k = v_b$, visible in the full matrix solution is absent in the analytic treatment due to the electrostatic approximation. The null vectors associated with these missing modes are electromagnetic, and are therefore not captured by the electrostatic analytic dispersion relation.

The strength of the splitting (and the degree of instability) depends on the beam parameters. Increasing the beam density enhances the coupling between species, leading to stronger mode splitting and larger growth rates. Reducing the beam velocity brings the beam mode into closer resonance with the background

plasma, also amplifying the instability. These effects are consistent with predictions from cold plasma theory.

5.4.3 Two-Stream Instability

The two-stream instability arises when an electron beam resonates with the natural plasma oscillations of the background plasma. This occurs when the phase velocity of an electrostatic wave matches the beam velocity, i.e., when $\omega/k = v_b$. In this regime, energy is efficiently transferred from the beam to the wave, destabilising the otherwise stable plasma mode and leading to exponential growth in its amplitude.

This instability is purely electrostatic and longitudinal, driven by the coupling of charge density fluctuations along the beam direction. In the context of the beam mode analysis presented earlier, the two-stream instability corresponds to the point in k-space where the imaginary part of the mode frequency, ω_i, reaches its maximum. This point is captured by the second panel in Figure 5.3 and is clearly visible in the numerical dispersion curves shown in Figure 5.4, where the (k, ω_i) profile reveals the strongest growth at $k_x = 5/3$. The growth rate at this point is $\omega_i \approx 0.11\,\omega_p$, and can be compared to the analytic growth rate

$$\omega_i = \frac{\sqrt{3}}{2^{4/3}} \left(\frac{n_b}{n_0}\right)^{1/3} \frac{\omega_p}{\gamma_b} \approx 0.12\,\omega_p, \tag{5.29}$$

provided by Bret et al. [181].

5.4.4 Weibel Instability

So far, only wavevectors parallel to the electron beam have been considered, i.e., $\mathbf{k} = (k_x, 0, 0)$. In this configuration, the dynamics are dominated by longitudinal (electrostatic) interactions, as seen in the two-stream instability. However, when the wavevector has a component transverse to the beam, a different class of instabilities emerges.

The Weibel instability arises due to anisotropy in the velocity distribution of the plasma, specifically when there is a net current or temperature difference along one direction. In the context of a beam-plasma system, the beam introduces a

strong velocity anisotropy: electrons stream along a single direction, while the background plasma remains isotropic. This anisotropy enables the growth of transverse perturbations, which are amplified through a magnetic feedback mechanism.

Physically, consider a small transverse current perturbation that generates a transverse magnetic field. This field deflects beam electrons, enhancing current filaments and reinforcing the magnetic field, resulting in a positive feedback loop that drives exponential growth. Unlike the two-stream instability, the Weibel instability is purely electromagnetic and transverse: the electric field is perpendicular to $\mathbf{k}$, and the instability does not require resonance with a plasma oscillation.

In terms of the dispersion relation, this manifests as a growing mode for wavevectors $\mathbf{k} \perp \mathbf{v}_b$, typically with $\omega_r \approx 0$ and $\omega_i > 0$, indicating purely growing modes rather than propagating waves. The null vector associated with these modes shows non-zero transverse velocity and magnetic field components (v_y, v_z, B_y, B_z), confirming their electromagnetic character.

The identification of these modes is shown in the first panel of Figure 5.5, for a wavevector $\mathbf{k} = (0, 2, 0)$. The second panel shows how ω_i varies with k_y. Whilst the imaginary component is increasing with k_y, it is clear that it asymptotes to $\sim 0.055\,\omega_p$ ($\omega_i(k_y = 10) = 0.053\omega_p$). This maximum growth rate can be compared to the analytic growth rate provided by Bret et al. [181],

$$\omega_i = \beta_b \sqrt{\frac{n_b}{\gamma_b n_0}}\,\omega_p \approx 0.056\,\omega_p, \tag{5.30}$$

demonstrating close agreement with the numerical solution of the fluid model.

5.4.5 Oblique Instabilities and Mixed Modes

In addition to purely longitudinal and transverse instabilities, the system can also support unstable modes at oblique angles to the beam axis. These oblique instabilities arise when the wavevector $\mathbf{k}$ has both longitudinal and transverse components (i.e., neither fully parallel nor fully perpendicular to the beam). They are typically characterised by hybrid modes that exhibit both electrostatic and electromagnetic features.

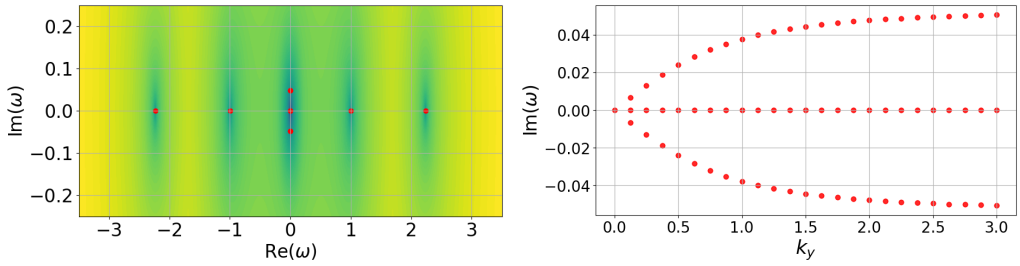


Figure 5.5: Identification of the Weibel instability. Left: Root structure in the complex frequency plane for $\mathbf{k} = (0, 2, 0)$, showing a purely growing mode. Right: Imaginary part of frequency ω_i as a function of transverse wavenumber k_y, showing unstable wave modes for $k_y > 0$.

Physically, oblique instabilities can be understood as a superposition of the two-stream and Weibel mechanisms. The longitudinal component of $\mathbf{k}$ allows for resonance with plasma oscillations, while the transverse component enables current filamentation and magnetic field amplification. This dual coupling allows oblique modes to tap into both types of free energy. As a result, the fastest-growing modes in many beam-plasma systems are neither purely longitudinal nor purely transverse, but oblique.

5.4.6 Comparing Instability Growth Rates

As mentioned previously, it is typically the wave that grows the fastest which dominates the early-time evolution of the system. It is therefore interesting, if not crucial, to compare the growth rates of the two-stream and Weibel instabilities in a range of beam-plasma configurations.

The analysis in this section has considered a system composed of a single cold electron beam propagating through an ion-electron background plasma. By varying the beam density and velocity, one can identify the regimes in which the dynamics of the system will be dominated by the Weibel instability - causing severe beam fillamentation and magnetic field growth - or the two-stream instability - driving strong electrostatic oscillations. Understanding the relative strength of these instabilities provides insight into which physical processes will shape the nonlinear evolution of the system under different beam conditions.

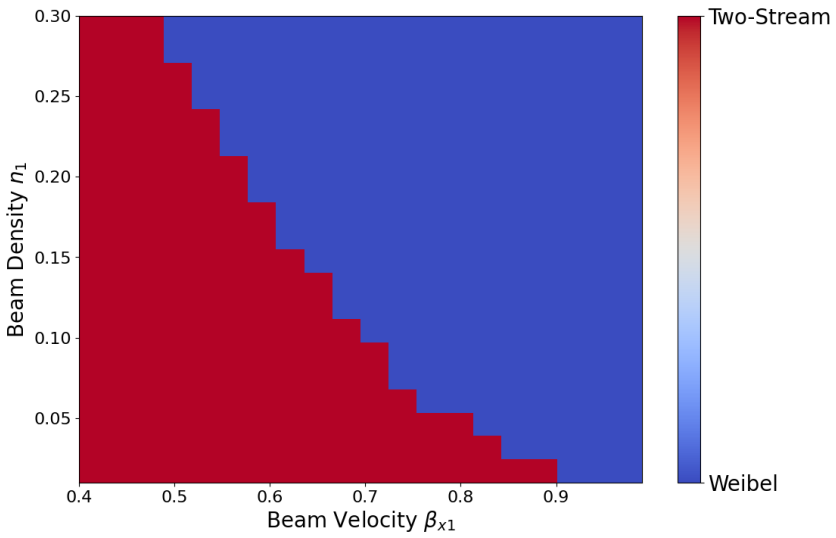


Figure 5.6: Binary map of instability dominance in the beam-plasma system. The coloured regions indicate whether the Weibel (blue) or two-stream (red) instability exhibits the highest growth rate for a given pair of beam parameters: velocity v_b and density n_b.

Figure 5.6 shows a binary map indicating which instability, the two-stream or the Weibel, dominates across the beam parameter space. Each pixel corresponds to a pair of beam density and velocity values. The instability with the highest maximum growth rate ω_i at that parameter point is identified and colour-coded.

Several clear trends emerge. At low beam velocities, the two-stream instability dominates, consistent with its electrostatic nature and reliance on resonance with background plasma oscillations. As the beam velocity increases, particularly into the relativistic regime, the Weibel instability becomes dominant over an increasing range of densities. This reflects the increased importance of velocity anisotropy and current-driven magnetic field generation in fast beams. The transition between the two regimes occurs smoothly and shifts depending on both beam parameters.

This mapping provides a valuable reference for predicting the nature of beam-plasma interactions in different physical regimes. In applications such as laser-plasma accelerators, astrophysical shocks, and high-intensity particle beams, understanding which instability dominates informs expectations about field generation, particle trapping, and energy transport.

5.5 Kinetic Effects and Vlasov-Maxwell Simulations

Up to this point, the plasma system has been described using a cold multi-fluid model, in which each species are treated as a continuous fluid characterised by bulk quantities such as density and velocity. While this approach captures many key wave phenomena and linear instabilities, it neglects critical kinetic effects related to velocity-space structure.

To capture these effects, a kinetic description of the plasma is required. To do this, the Vlasov-Maxwell simulation code `VALIS` [105], as outlined in Section 2.7.2, is employed to study the beam-plasma system from a fully kinetic perspective. This framework captures rich physics such as Landau damping [182], particle trapping, nonlinear phase mixing, and fine-scale velocity-space structure, all of which are essential in understanding the long-term evolution of beam-plasma instabilities.

5.5.1 Initial Conditions and Simulation Setup

To describe the beam-plasma system kinetically, the full initial distribution functions, $f_s(x, p, t = 0)$, of each species must be specified. The Vlasov-Maxwell system then governs the time evolution from this initial state.

Let the background plasma, composed of ions and electrons, be initially in thermal equilibrium and uniform in space. The plasma is in thermal equilibrium and is therefore described by a Maxwellian velocity distribution (Equation 2.1). The electron beams are also assumed to be spatially uniform and are initialised as drifting Maxwellians propagating with velocities $\pm\mathbf{v}_b$ through the background plasma:

$$F_{\pm b}(\mathbf{v}) \propto \exp\left(-\frac{m(\mathbf{v} \mp \mathbf{v}_b)^2}{2k_B T}\right). \tag{5.31}$$

The case initially considered consists of two relativistic electron beams propagating in opposite directions through the background plasma. The beams are initialised with velocities $\mathbf{v}_{\pm b} = \pm 0.94\, c$, corresponding to a kinetic energy $E_k = 1$ MeV and momentum $p = \pm 2.7\, m_e c$. Each beam has a density of $n_{\pm b} = 0.01\, n_0$, and their velocity spread is characterised by a temperature of $T_{\pm b} = 40$ keV.

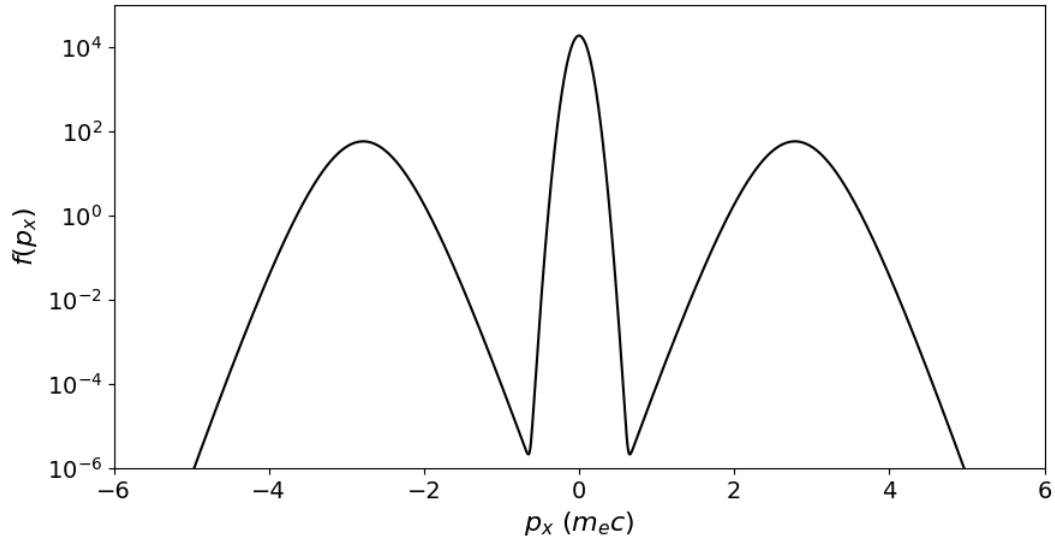


Figure 5.7: The inital electron momentum distribution, which includes the background plasma electrons and the two electron beams.

The background plasma is constructed to ensure there is no net charge or current density initially. This requires an ion density of $n_i = n_0$ and an electron density of $n_e = 0.98\, n_0$. No drift velocity is assigned to the background species, as the beam currents are equal and opposite, cancelling each other out. The background ion and electron temperatures are both set to $T_i = T_e = 4$ keV, which is comparable to conditions in the hotspot region of inertial confinement fusion (ICF) implosions. The total electron distribution, a superposition of the three electron populations ($f_e + f_{+b} + f_{-b}$), is shown in Figure 5.7.

The simulation used a spatial grid with 4096 cells spanning a domain of $400\, c/\omega_p$. The momentum space was discretised into 2048 cells covering the range $\pm 10\, m_e c$, which is sufficient to capture the full evolution of the distribution function. An ion mass of $1836\, m_e$ was used, corresponding to protons (i.e., a hydrogen plasma). The simulation was carried out for a duration of $1500\, \omega_p^{-1}$.

An overview of the energy evolution of the system is shown in Figure 5.8. This figure displays the change in kinetic energy of the beam, the energy of the background plasma electrons, and the electric field energy, all normalised to the initial beam kinetic energy ($\Delta E / E_{beam}^{\text{initial}}$). To improve clarity, these plots have been smoothed to remove high-frequency oscillations. Three distinct temporal regions can be identified.

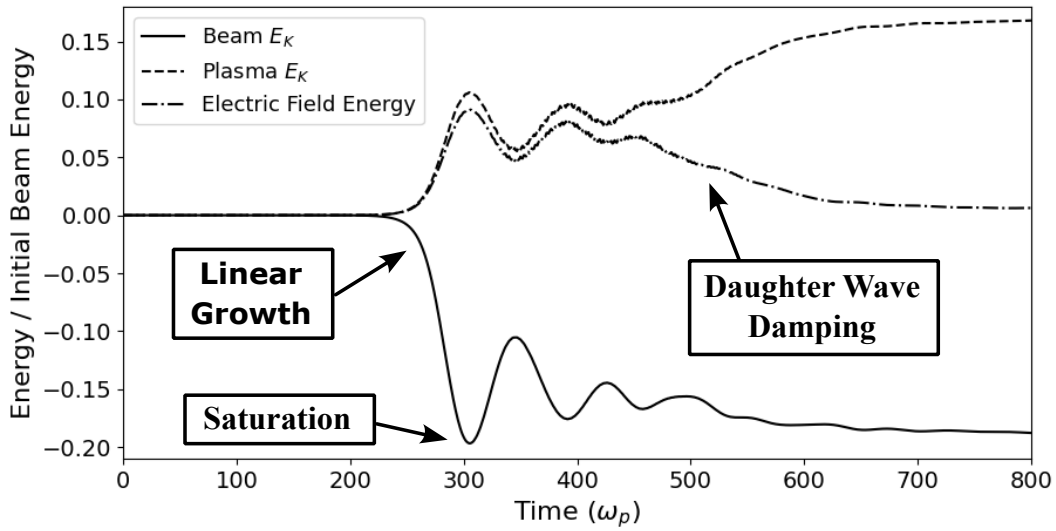


Figure 5.8: Evolution of the energy in the beam-plasma system. The kinetic energy in the electron beams (black), the kinetic energy of the background electrons (dark grey), and the energy stored in electric fields (light grey). These plots have been smoothed to remove high frequency oscillations.

In the first stage, there is an exchange of energy from the beam electrons to the background electrons along with energy being stored in electric fields. This energy exchange grows exponentially and corresponds to the linear growth phase described in the earlier analysis, and is driven by the resonant interaction between the beams and background plasma. Eventually, the growth saturates, marking the second stage. A brief oscillatory exchange of energy occurs, followed by the third stage: the transfer of electric field energy into the background plasma, initiating a second stage of plasma heating. Finally, the system stabilises, with no further energy exchange between the fields and particles.

These three phases, linear growth, saturation, and nonlinear relaxation, define the structure of the following subsections.

5.5.2 Linear Growth via Bump-on-Tail Instability

In the initial stage of the simulation, energy is transferred from the electron beams to the background plasma electrons, with a roughly equal amount of energy stored in the growing electric fields. This energy exchange grows exponentially in time and is indicative of a linear instability driven by resonant interactions.

As discussed previously, the early dynamics of the system are well described by linear theory, where perturbations are small compared to the unperturbed background. In the multi-fluid analysis of Section 5.4, it was shown that an electron beam propagating through a background plasma can give rise to unstable (exponentially growing) electrostatic plasma oscillations, or Langmuir waves when thermal effects are included.

A similar process is observed in the kinetic system here. The beams drive unstable Langmuir waves by resonantly interacting with the background electrons, extracting energy from the beam which then oscillates between electric field energy and electron kinetic energy. As the Langmuir wave grows exponentially in amplitude, the electric field energy and background electron energy also grow accordingly.

Up to this point, the instability has been referred to as "similar" to the two-stream instability described earlier. This is because the assumptions underlying the fluid model (particularly the assumption of single-velocity species) no longer apply. In the kinetic description, the instability responsible for wave growth is referred to as the bump-on-tail instability, which is a direct result of one of the most significant phenomena captured by kinetic theory: Landau damping [182].

Landau Damping/Growth

Landau damping refers to the resonant exchange of energy between plasma waves and particles moving near the wave's phase velocity. While the multi-fluid description assumes each species has a single velocity at a given point in space, the kinetic approach provides the full distribution of velocities, enabling a richer range of dynamics.

In the linearised Vlasov-Poisson system for electrostatic waves, the dispersion relation includes an integral over the particle velocity distribution. The imaginary part of the frequency, corresponding to the temporal growth or damping of the wave, arises from a pole in this integral when the phase velocity of the wave matches the velocity of particles in the distribution. The resulting growth or damping rate contributed by a given species s is [182]:

$$\gamma_s = -\frac{\pi}{2}\frac{\omega_{ps}^2}{k}\left.\frac{\partial f_s}{\partial v}\right|_{v=\omega/k}, \tag{5.32}$$

where ω_{ps} is the plasma frequency of species s, k is the wavenumber, and $\partial f_s/\partial v$ is the velocity derivative of the distribution function evaluated at the wave's phase velocity $v = \omega/k$. The total damping or growth rate is obtained by summing over all species:

$$\gamma = \sum_s \gamma_s. \tag{5.33}$$

This equation indicates that Landau damping and growth is separated by the sign of $\partial f_s/\partial v$. Landau damping occurs when $\partial f_s/\partial v < 0$, indicating that more particles are slower than the wave and can absorb energy from it. Conversely, if $\partial f_s/\partial v > 0$, the distribution can transfer energy to the wave, leading to Landau growth. This mechanism is particularly important in beam-plasma systems, where drifting populations introduce positive velocity gradients, driving instability via this mechanism.

Unlike fluid instabilities, which rely on differences in bulk flow, Landau damping and growth are intrinsically kinetic and depend on the phase-space structure of the distribution function, not just its moments. This makes them essential for understanding the stability and growth of Langmuir waves in collisionless plasmas.

Bump-on-Tail Instability

The bump-on-tail instability is a direct consequence of the Landau growth mechanism. As shown in Figure 5.7, the initial velocity distribution contains regions of positive gradient in velocity space introduced by the two drifting electron beams. These positive gradients drive the growth of Langmuir waves by feeding energy into waves whose phase velocities lie within the region of positive gradient.

Langmuir waves are longitudinal plasma oscillations in which electron displacements produce restorative electric fields, leading to coherent oscillations. Their presence can therefore be identified by taking a spatial Fourier transform of the electron density. This decomposes the density field into its spectral components. A strong Langmuir wave manifests as a peak in the amplitude spectrum at the wave's wavenumber.

Figure 5.9 shows this spectral decomposition. The left panel presents the spatial structure of the density, while the right panel shows the corresponding

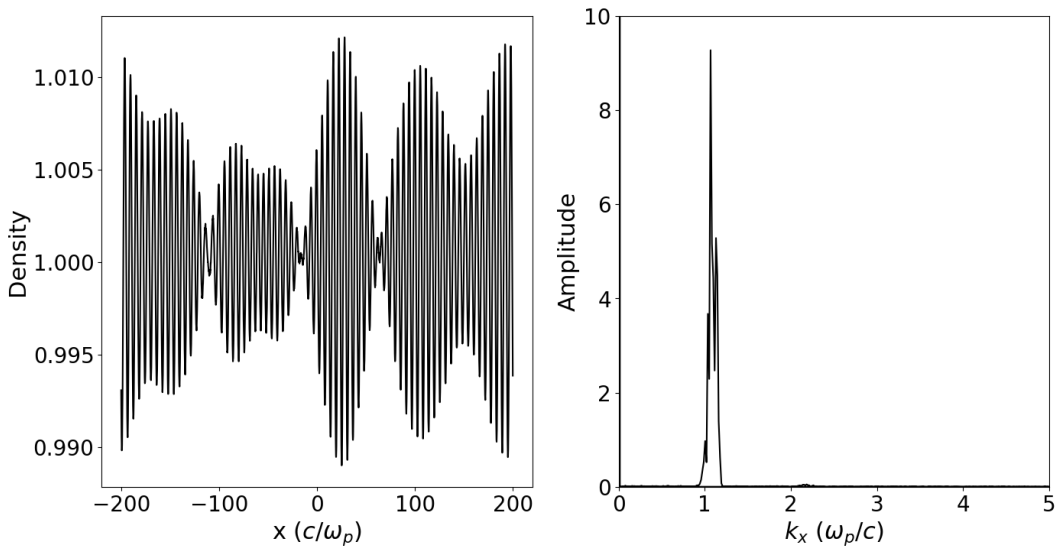


Figure 5.9: (Left) Background plasma electron density as a function of position at $t = 230\,\omega_p^{-1}$. (Right) Fourier transform of the density showing a clear spike at $k \sim 1.07$ corresponding to the dominant Langmuir wave driven by the bump-on-tail instability.

Fourier spectrum at $t = 230\,\omega_p^{-1}$. A clear spectral peak indicates the presence and growth of a Langmuir wave mode.

5.5.3 Saturation from Electron Trapping and Dephasing

The exponential growth of the Langmuir wave eventually saturates as nonlinear effects become significant. The critical transition occurs when the electron bounce frequency, ω_B, becomes comparable to the growth rate, ω_i, of the wave:

$$\omega_B = \sqrt{\frac{ekE}{m_e}} \sim \omega_i. \tag{5.34}$$

In the early (linear) stage, resonant electrons moving near the wave's phase velocity $v_\phi = \omega/k$ stay phase-locked with the wave, efficiently transferring energy and driving exponential growth. But as the wave amplitude increases, and trapping sets in, these same electrons begin to oscillate in the wave potential with frequency ω_B.

Once $\omega_B \sim \omega_i$ the electrons can no longer remain synchronously phased with the wave. Instead, they bounce in and out of phase, forming closed orbits in phase space. This trapping flattens the electron distribution function near the phase velocity, reducing the gradient in momentum space that drives the instability.

As a result, these electrons no longer contribute constructively to wave growth, and the instability saturates.

Electron Trapping

Electron trapping is most easily understood by considering the frame moving with the phase velocity of the Langmuir wave, $v_\phi = \omega/k$. In this frame, the wave appears as a stationary oscillating potential. For a single dominant wave mode, the electric field and corresponding potential in this frame are approximately given by:

$$E(x,t) = E_0(t)\sin(kx), \tag{5.35}$$

$$\phi(x,t) = -\frac{E_0(t)}{k}\cos(kx) = \phi_0(t)\cos(kx). \tag{5.36}$$

The total energy of an electron in this potential is conserved and can be written as:

$$\epsilon = \frac{1}{2}m_e(\Delta v)^2 + e\phi(x,t), \tag{5.37}$$

where $\Delta v = v - v_\phi$ is the electron velocity relative to the wave. As the wave amplitude grows, the depth of the potential well increases, and electrons whose kinetic energy in the wave frame is smaller than the potential barrier become trapped.

The condition for trapping is:

$$\frac{1}{2}m_e(\Delta v)^2 \lesssim e\phi_0(t), \tag{5.38}$$

which gives a trapping width in velocity space:

$$\Delta v_{\text{trap}} = \sqrt{\frac{2e\phi_0(t)}{m_e}}. \tag{5.39}$$

Once trapped, electrons oscillate in phase space with a characteristic bounce frequency given by:

$$\omega_B = k\Delta v_{\text{trap}} = \sqrt{2\frac{ek^2\phi_0(t)}{m_e}}, \tag{5.40}$$

which describes the nonlinear oscillation of electrons within the wave potential.

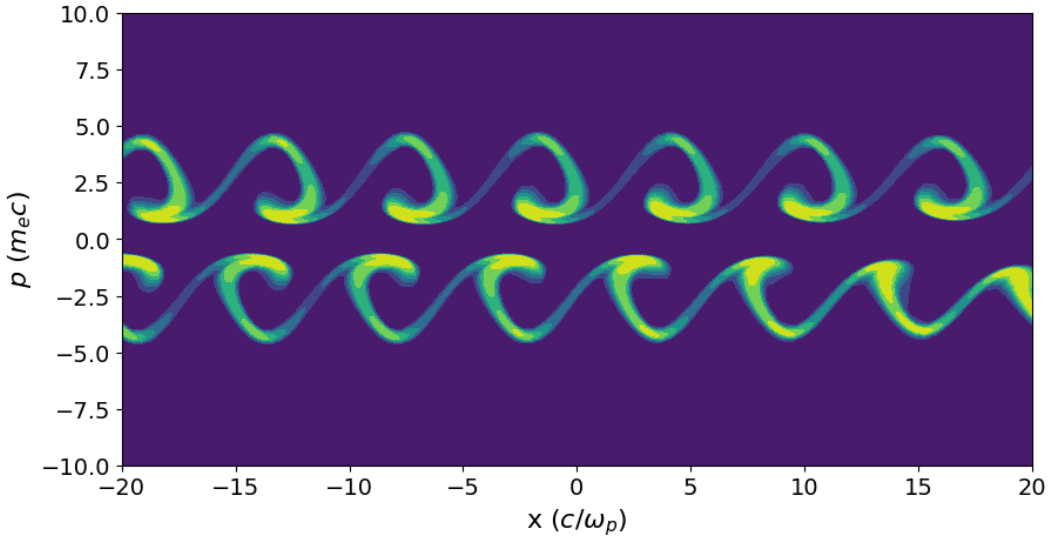


Figure 5.10: (x, p) phase space of the electron beams at $t = 300\,\omega_p^{-1}$. Voritces are formed due to the trapping of electrons in the high amplitude Langmuir wave.

Saturation and Phase-Space Signatures

Evidence of trapping and saturation is clearly visible in the simulation. Figure 5.10 shows the formation of phase-space vortices at the point of saturation ($t = 300\,\omega_p^{-1}$) where electrons oscillate within the wave potential. These are direct visual signatures of electron trapping and mark the transition to nonlinear behaviour.

Oscillations at the bounce frequency also appear in the temporal evolution of the system's energy. As the wave saturates, the electron beam energy begins to oscillate due to the coherent motion of trapped particles. These oscillations can be seen in Figure 5.8, shortly after the linear growth phase ends.

As trapping becomes significant, the wave ceases to grow exponentially and begins to saturate. The physical mechanism is the flattening of the velocity distribution function around the wave's phase velocity. This suppresses Landau growth, as the positive slope that previously drove the instability is eliminated by the redistribution of trapped particles. This effect is shown in Figure 5.11, where the momentum distribution of the background electrons develops a plateau in the trapping region, suppressing further energy exchange and stabilising the wave.

This saturated state represents a quasi-equilibrium between the trapped particle population and the wave potential. However, the system is not static: in the following

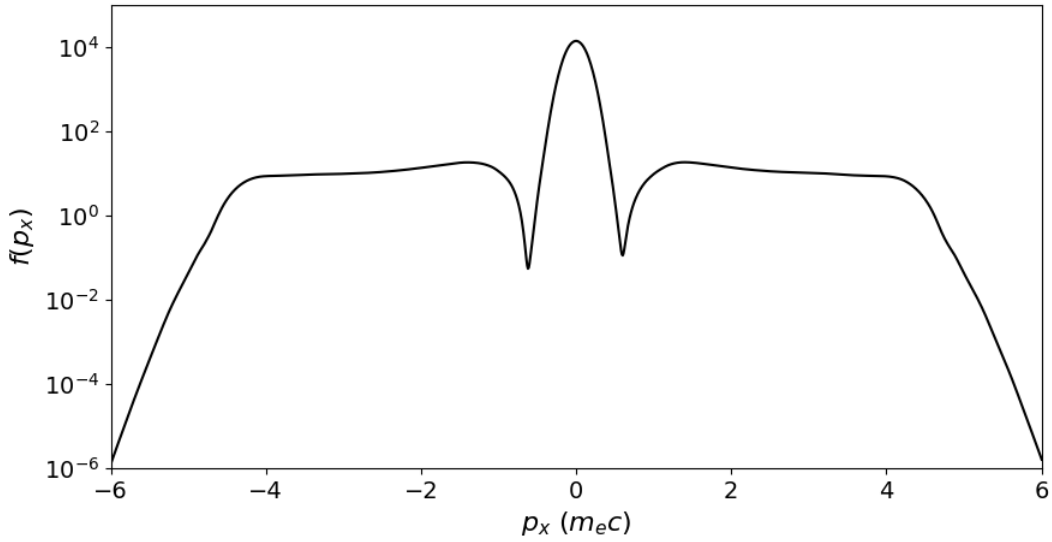


Figure 5.11: The total electron momentum distribution at $t = 300\,\omega_p^{-1}$. A plateau is starting to form between the peaks of the beams and the background plasma electrons. This is due to the trapping of electrons in the high amplitude Langmuir wave.

stage, nonlinear effects such as wave-wave coupling and spectral broadening begin to dominate, driving the system into a more turbulent regime.

5.5.4 Non-Linear Mode Coupling

Once the bump-on-tail instability saturates and the electron momentum distribution stabilises, no further energy is extracted from the electron beams. However, a significant amount of energy remains stored in the electric field of the residual high-amplitude Langmuir wave. As previously shown, this energy ultimately dissipates into the kinetic energy of the background electrons... but how?

Until now, the analysis has focused on linearised systems, where each physical quantity is assumed to consist of a background value plus a small perturbation. The critical assumption here is the smallness of the perturbation: that higher-order terms in the governing equations can be neglected. This assumption breaks down when the amplitude of the Langmuir wave becomes large, as is the case here. Once higher-order (nonlinear) terms become significant, the dynamics depart from linear theory, and we enter the nonlinear regime.

Non-linear Coupling

In the nonlinear regime, the plasma evolves not through simple linear superposition, but through interactions between wave modes, a process known as nonlinear coupling.

Fourier analysis remains a powerful tool for understanding this behaviour. However, unlike in the linear case, Fourier modes no longer evolve independently. Nonlinearities in the governing equations induce mode coupling, meaning the evolution of a given mode depends on the presence and properties of other modes. These interactions give rise to phenomena such as harmonic generation, beat-wave structures, and energy transfer between modes.

Despite this interdependence, expressing the system as a sum of plane waves remains both valid and insightful, particularly for identifying dominant wave activity, resonance behaviour, and pathways of energy redistribution. It is therefore appropriate to examine the evolution of individual plane wave components. For example, a perturbed quantity can be expressed as:

$$A \propto \mathrm{Re}\left[\exp\left(i(kx-\omega t)\right)\right] = \frac{1}{2}\exp\left(i(kx-\omega t)\right) + \frac{1}{2}\exp\left(-i(kx-\omega t)\right), \quad (5.41)$$

where the complex conjugate ensures real-valued physical quantities.

To illustrate nonlinear coupling, consider the second-order product of two waves:

$$\begin{aligned} A_1 A_2 \propto &\left[\exp\left(i(k_1 x-\omega_1 t)\right) + \exp\left(-i(k_1 x-\omega_1 t)\right)\right] \\ &\times \left[\exp\left(i(k_2 x-\omega_2 t)\right) + \exp\left(-i(k_2 x-\omega_2 t)\right)\right]. \end{aligned} \quad (5.42)$$

Multiplying these terms generates new components with wavevectors and frequencies at:

$$(k_1 + k_2,\ \omega_1 + \omega_2),\ (k_1 - k_2,\ \omega_1 - \omega_2), \quad (5.43)$$

and their complex conjugates. These correspond to wave modes whose wavevectors and frequencies are either the sum or difference of the parent waves. In general, any pair of interacting modes can produce new modes with:

$$(k_{\mathrm{new}}, \omega_{\mathrm{new}}) = (k_1 \pm k_2,\ \omega_1 \pm \omega_2), \quad (5.44)$$

along with their complex conjugates, provided the parent modes have sufficient amplitude to drive the interaction.

In this simulation, the dominant modes are the forward and backward Langmuir waves with wavevectors and frequencies approximately at $(\pm 1, 1)$. Through the coupling of these first-order terms, second-order modes can be driven at:

- $(\pm 2, 2)$ — higher harmonics, which are typically non-resonant and strongly damped,
- $(\pm 2, 0)$ — static modulations, contributing to wave-wave beat structures.

As more modes appear, a cascade begins: first-order modes couple with second-order modes to produce third-order products, such as:

$$(1, 1) + (2, 0) \to (3, 1), \tag{5.45}$$

$$(1, 1) + (2, -2) \to (3, -1), \tag{5.46}$$

$$(1, 1) + (2, 2) \to (3, 3). \tag{5.47}$$

Whether these modes persist depends on whether they lie near a natural mode supported by the plasma. The modes at $(3, 1)$ and $(3, -1)$ lie near the Langmuir dispersion curve and are weakly damped; these are resonant and thus accumulate energy. In contrast, modes like $(3, 3)$ are non-resonant and rapidly dissipate.

This hierarchical coupling continues to higher orders, though each successive step becomes weaker. This process is evident in Figure 5.12, which shows the 2D Fourier transform of the electron density in (k, ω)-space. The strongest features lie along the linear Langmuir dispersion curve, consistent with resonance. Weaker harmonics and mode-mixing features are visible at higher k and ω, especially near sum/difference positions such as $(2, 2)$ and $(2, 0)$.

This nonlinear spectral evolution marks the transition from coherent wave growth to a more turbulent regime where wave-wave interactions, energy redistribution, and damping compete, leading to a broad spectrum of modes.

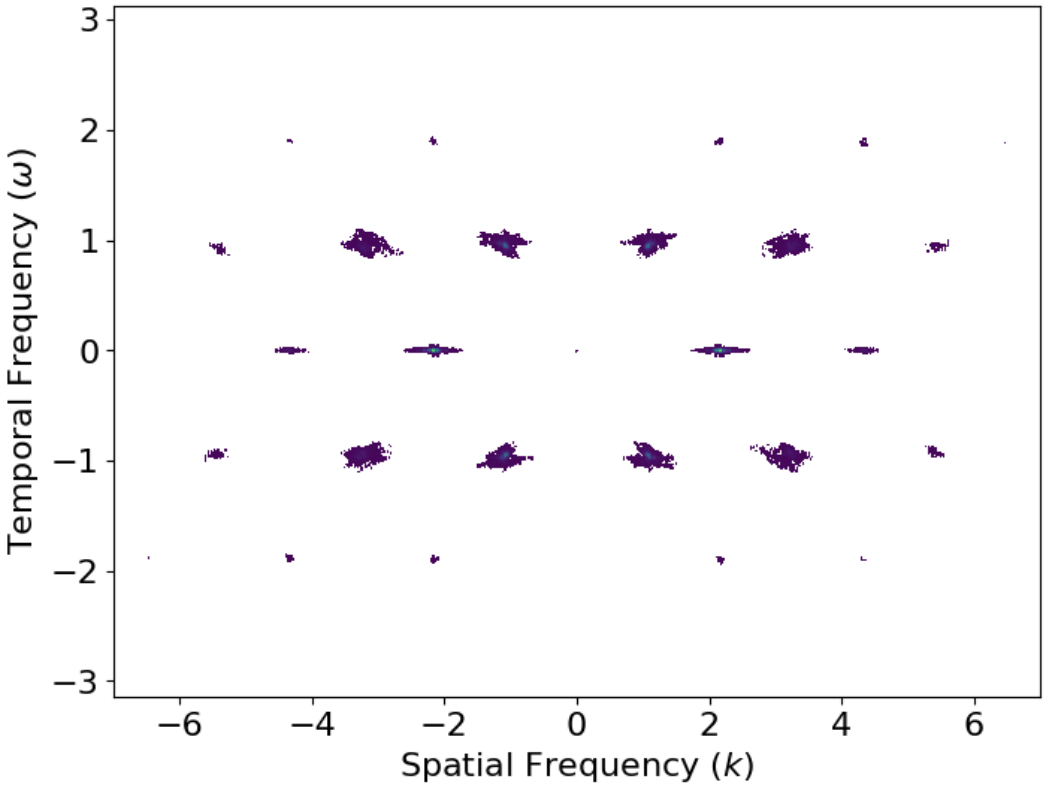


Figure 5.12: Two-dimensional Fourier transform of the background electron density, showing energy distribution in (k, ω)-space. Strong features lie along the Langmuir dispersion curve, with weaker nonlinear features at higher-order coupling locations.

Daughter Wave Damping

Nonlinear coupling transfers energy from the dominant Langmuir wave, $(1, 1)$, into higher-order Langmuir modes with odd-integer multiples of the wavevector, $(1, 2n + 1)$. Since the phase velocity of these waves is given by ω/k, daughter modes have lower phase velocities than the original wave and thus interact with lower-velocity regions of the electron distribution.

Like the primary wave, they can trap electrons, leading to the formation of additional plateaus in phase space, centred at:

$$v_\phi^{(n)} = \frac{\omega}{(2n+1)k}, \tag{5.48}$$

where n is the harmonic order of the daughter wave. Because these waves have lower amplitudes, the resulting plateaus are typically narrower than the original.

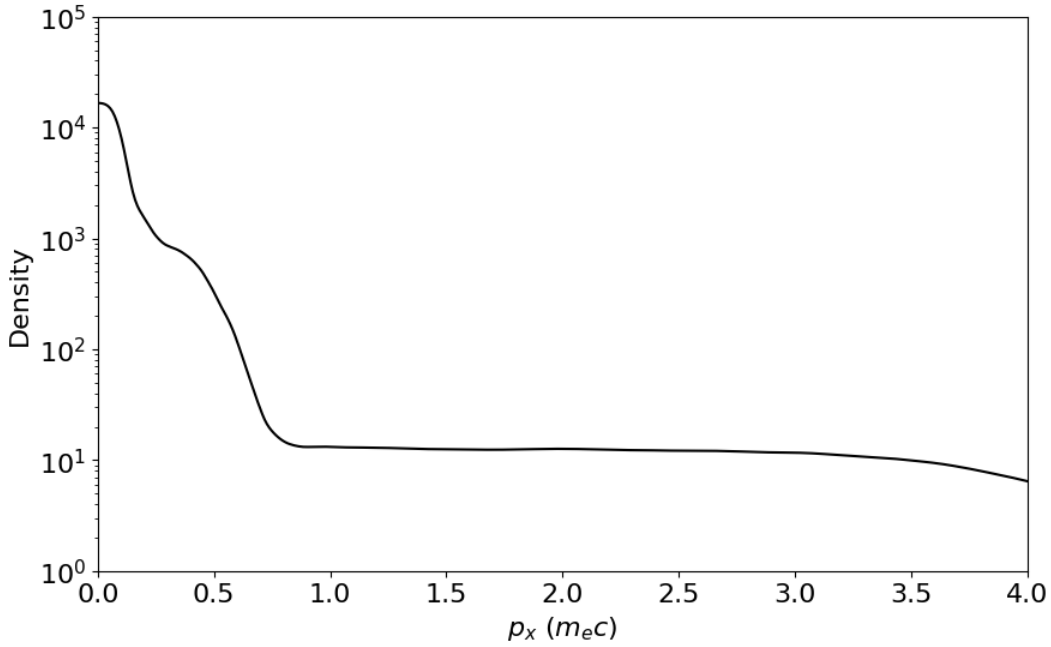


Figure 5.13: Final electron momentum distribution at the end of the simulation for electrons beams with energy and density 1 MeV and $0.01 n_0$ respectively. The distribution is stable (monotonically decreasing), and an additional plateau is visible at low momentum due to the trapping of electrons in higher order Langmuir waves.

An additional plateau can be observed in the final distribution function of the electrons (the combined distribution of beam and background electrons) around $p_x = 0.45\, m_e c$. This final momentum distribution is shown in Figure 5.13 at timestep $t = 1000\, \omega^{-1}$.

This mechanism effectively extends trapping and heating to successively lower-velocity regions of the background plasma. While the initial wave primarily interacts with fast electrons near the beam velocity, the daughter waves drive heating of bulk plasma electrons. In this way, energy initially stored in the electric field is gradually transferred to bulk electrons, resulting in preferential heating of the background plasma.

5.6 Coupling Efficiency

While the preceding analysis focused on describing the evolution of the system, many practical applications are ultimately concerned with a different question: how efficiently can an electron beam transfer energy to the background plasma?

This quantity, referred to as the coupling efficiency, is defined in this work as the fraction of beam kinetic energy lost.

Two approaches to estimating the coupling efficiency are explored in this section: one based on electron trapping theory, and the other derived from analysis of the momentum distribution.

5.6.1 Electron Trapping Theory

The amount of energy which can be extracted from the beam is limited by the saturation mechanism: electron trapping. It is therefore natural to use electron trapping theory to produce an estimate of the coupling efficiency.

Significant work has been conducted predicting the energy contained in the electric field at the point of saturation. Quasi-linear analysis, in which perturbations are assumed small but finite and alterations to the momentum distribution feed back on the wave growth, provides a framework for understanding this saturation process more quantitatively. As the wave amplitude grows, resonant electrons become trapped in the wave potential, causing a flattening of the distribution function near the resonant velocity. This flattening reduces the free energy available to drive the instability, thus limiting further wave growth.

Fainberg, Shapiro, and Shevchenko [183] derived an expression for the saturated electric field energy (normalised to the beam's initial total energy):

$$W = \sum_k \frac{|E_k|^2}{8\pi} \frac{1}{n_b \gamma_0 m c^2}, \tag{5.49}$$

which depends on the dimensionless parameter described by the beam and plasma parameters,

$$S = \beta_0^2 \gamma_0 \left(\frac{n_b}{2n_0} \right)^{\frac{1}{3}}, \tag{5.50}$$

yielding a saturation estimate:

$$W = 0.158S. \tag{5.51}$$

This result captures the nonlinear feedback and trapping dynamics, but is limited to the weak beam ($S \ll 1$) regime, and does not account for relativistic effects.

A more general semi-empirical model by Thode and Sudan [184] extends the theory into the relativistic regime:

$$W = \frac{1}{2} S (1+S)^{-\frac{5}{2}}, \tag{5.52}$$

indicating that in the weak beam ($S \ll 1$) regime $W \propto S$, whilst in the strong beam regime ($S \gtrsim 1$) relativistic effects lead to sharper saturation.

Assuming the extracted beam energy is equally split between the wave's electric field energy and the kinetic energy of background electrons, the total coupling efficiency is approximately $2W$. However, W is normalised to the beam's total (rest + kinetic) energy, while the quantity of interest is the fraction of available (kinetic) beam energy. This is therefore given by

$$2\tilde{W} = 2\frac{\gamma_0}{\gamma_0 - 1} W. \tag{5.53}$$

5.6.2 Electron Momentum Evolution

An alternative estimate for coupling efficiency is based on the initial and final electron momentum distributions. Regardless of the saturation mechanism, stability requires the velocity distribution to become flat, satisfying:

$$\sum_s -\frac{\pi}{2}\frac{\omega_{ps}^2}{k}\left.\frac{\partial f_s}{\partial v}\right|_{v_\phi = \omega/k} \leq 0 \quad \forall v_\phi. \tag{5.54}$$

It is therefore possible to estimate the final momentum distribution by modifying the initial distribution to be flat whilst conserving number density.

In qualitative terms, the "bump" in the beam distribution must fill the "trough" between it and the background plasma to remove the positive slope that drives instability. By shifting beam electrons into this trough (conserving electron number), a flat, stable distribution is formed. The coupling efficiency can then be estimated by calculating the difference in kinetic energy between the initial and final distributions.

This is illustrated in Figure 5.14, where the initial distribution is shown along with the region of beam electrons that can be shifted into the trough. By moving the beam electrons from the blue region into the red region between the background and beam, a perfectly flat distribution will be formed.

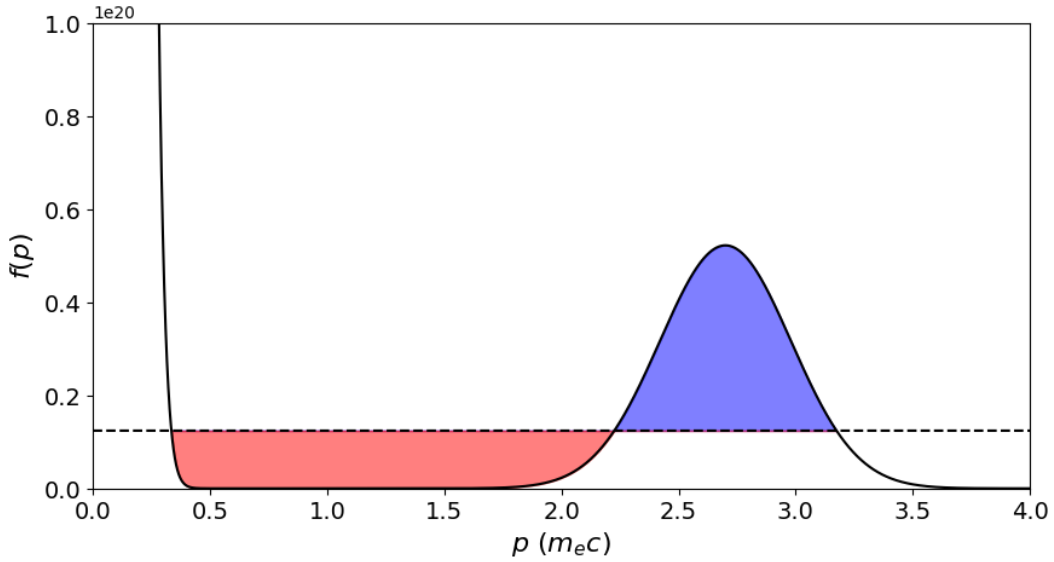


Figure 5.14: Schematic of how the initial electron distribution can be modified to stabilise the system. The blue region represents excess beam electrons that can be redistributed into the red trough to form a flat, stable distribution. This process removes the positive velocity gradient that drives the instability, enabling an upper-bound estimate of energy transfer from the beam to the background plasma.

This method does not conserve energy, as beam electrons are shifted to lower energies without accounting for energy redistribution. In reality, the background distribution broadens, contributing to the filling of the trough. Moreover, this approach assumes the high-energy tail of the beam remains unchanged. In practice, the plateau width is governed by the amplitude of the dominant Langmuir wave and can extend beyond the beam's initial tail, potentially moving some electrons to higher energies. By neglecting both of these effects, this approach acts as an upper bound on the possible coupling efficiency.

5.6.3 Simulation Parameter Scan

To evaluate the accuracy of these models, a series of `VALIS` simulations were performed across a range of beam energies and densities. Since the parameter S depends only on the beam velocity and density, these simulations allow direct testing of the models across both beam parameters.

A total of 25 simulations were performed, grouped into three sets:

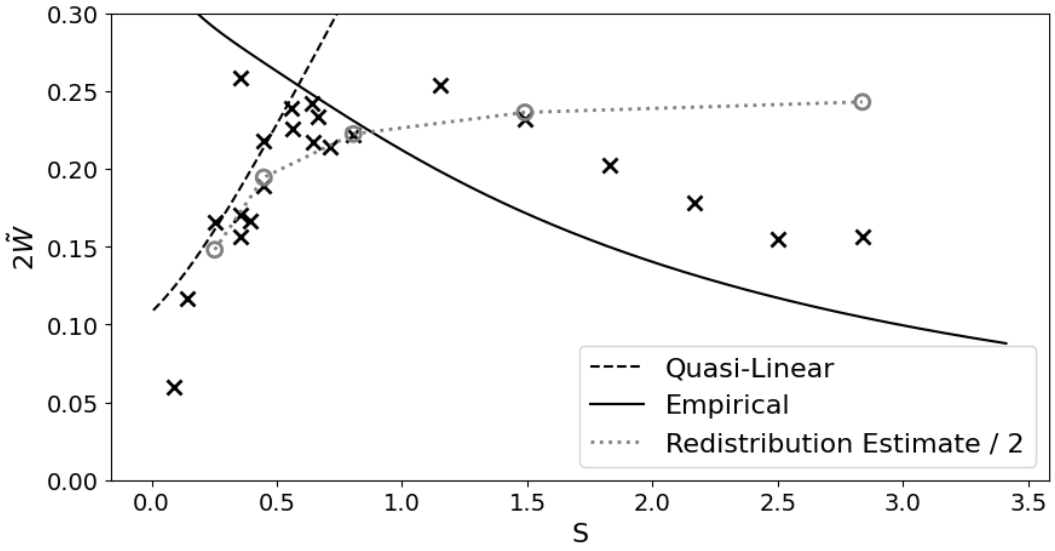


Figure 5.15: Coupling efficiency of the beam-plasma system as a function of S. Results are compared against three theoretical predictions: the quasi-linear trapping model (Equation 5.51), the Thode-Sudan model (Equation 5.52), and the redistribution estimate halved.

- 15 simulations at fixed beam density $n_b = 0.01n_0$, with beam energy varied from 0.15 to 8 MeV.
- 5 simulations at fixed beam energy of 1 MeV, with n_b varied from 0.005 to $0.04n_0$.
- 5 simulations at 2 MeV, again scanning n_b over the same range.

Together, these span values of S from 0.09 to 2.83.

As shown in Figure 5.15, simulation results agree well with the quasi-linear theory (Equation 5.51) for $S < 0.45$. For higher values, where relativistic and nonlinear effects become significant, the predictions of Thode and Sudan (Equation 5.52) better match the simulation data. At very high S (> 1.5), simulations exceed the predictions for the relativistic electron trapping model. This discrepancy arises due to the trapping-based model predicting the maximum electric field energy. Whilst the source term for the electric field energy is the coupling from the electron beam, there are also sink terms extracting energy from the electric field. Most notably, for particularly high amplitude electric fields, non-linear coupling into heavily damped

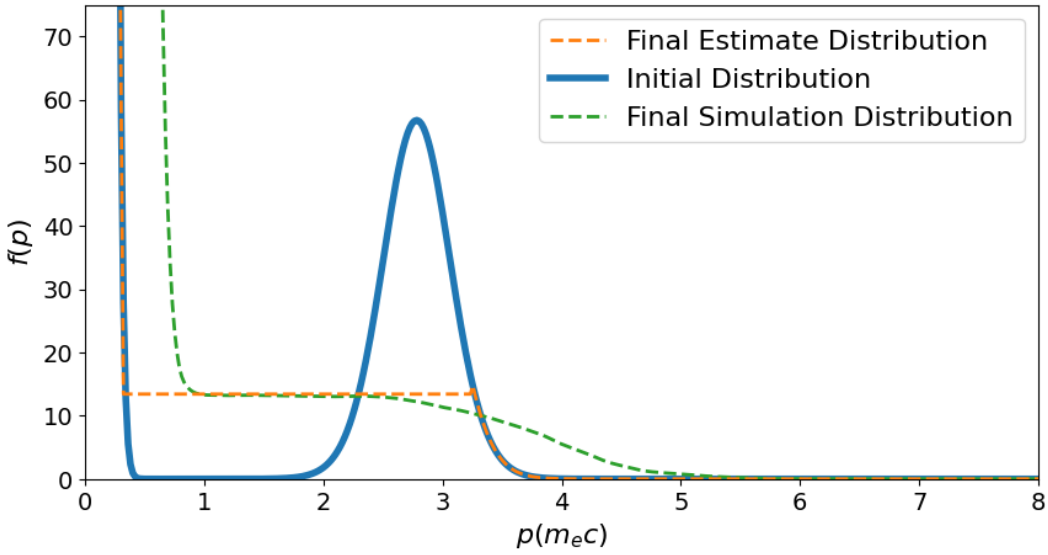


Figure 5.16: Comparison of simulated and predicted final electron distribution functions for a 1 MeV, $n_b = 0.01\, n_0$ beam. The simulation shows a significantly broader plateau than the predicted distribution from the momentum evolution model. This is due to eletron trapping accelerating electrons to energies beyond the tail of the beam. This broader plateau results in a notable difference in predicted vs. simulated coupling efficiency. Notably, the momentum-based model overestimates the coupling efficiency by a factor of two.

modes (e.g. the $(2, 2)$ mode) is significant. This energy moves swiftly from the electric field to background plasma electrons. Once the electric field energy sink term balances the source term, the electric field energy saturates, but this do not require the source term to be 0. Because of this, energy can still be extracted from the electron beam beyond the saturation of the electric field energy. This is particularly relevant for high strength interactions which drive high amplitude Langmuir waves, and is why this effect becomes noticeable above an S threshold.

Comparing the simulation results to the momentum evolution estimates yields several notable differences. Figure 5.16 compares the final simulated electron distribution with the distribution predicted by the momentum redistribution method outlined in Section 5.6.2 for the 1 MeV beam with density $n_b = 0.01\, n_0$. In this simulation, the dominant Langmuir wave driven by the bump-on-tail instability has a phase velocity of approximately $u_\phi \approx 2.5$.

For the distribution to reach a stable configuration, the electron distribution function must flatten over the range $0.5 \lesssim u \lesssim 2.5$. Achieving this requires

the Langmuir wave (centred at $u_\phi \approx 2.5$) to trap electrons with a trapping half-width of approximately $\Delta u/2 \approx 2$. Because this trapping width also applies to electrons with velocities greater than the wave phase velocity, the final simulated plateau extends to higher velocities, spanning approximately $0.5 \lesssim u \lesssim 4.5$. This plateau is therefore significantly wider than that predicted by the momentum redistribution model and includes electrons that have been accelerated by the wave. As a result, the final coupling efficiencies differ substantially between the predicted and simulated distributions.

Interestingly, the predicted coupling efficiency from the momentum redistribution model is consistently approximately twice the value observed in the simulations. When the predicted values are reduced by a factor of two, the agreement with the simulation results becomes remarkably good. Figure 5.16 shows the predicted coupling efficiency using this redistribution model in grey, demonstrating that halving the predicted value yields a much closer match to the simulated coupling efficiency. The physical origin of this empirical factor remains unclear, and whether this behaviour persists for other beam distributions remains an open question requiring further investigation.

5.7 Summary

The unstable beam-plasma system has been presented in this chapter and the instabilities driven by the system described. Linearisation of the multi-fluid equations was used to investigate modes supported by the system, the early time evolution of the system, and identifying the most unstable mode for a range of beam parameters. The Vlasov-Maxwell simulation code `VALIS` was then utilised to determine the non-linear evolution of the system. Electron trapping was identified as the mechanism saturating the two-stream instability, and mode coupling was shown to preferentially deposit electric field energy in the background plasma electrons. The timescale for the evolution of the system was determined to be on the order of $\sim 1000\,\omega_p$.

Methods of estimating the fraction of energy that can be extracted from the electron beam were presented and compared to numerical simulation results. These approaches indicate that quite significant coupling efficiencies of ~ 20 % can be achieved - a key parameter when considering the application of such heating schemes. This study, however, considers only idealised beam and plasma distributions. The next chapter explores the generation and propagation of electron beams in more realistic scenarios, and the implications of these results for beam-driven plasma heating in ICF contexts.

6

Auxiliary Heating Inertial Confinement Fusion Implosions

Contents

6.1 Introduction 172
6.2 Electron Beam Generation, Transport, and Hotspot Coupling 173
6.2.1 Electron Beam Generation 174
6.2.2 Channel Formation 176
6.2.3 Electron Collisional Scattering and Distribution Evolution 178
6.2.4 Beam Divergence and Density Evolution 181
6.2.5 Electron Beam Properties at the Hotspot 183
6.2.6 Improving Electron Beam Transport and Divergence . . 187
6.3 Implementation of Auxiliary Heating in Low-CR Simulations 188
6.3.1 Previous Exploration of AH in ICF Implosions 188
6.3.2 Interface vs Hotspot Energy Deposition 190
6.3.3 Optimum Deposition Time 190
6.3.4 Energy Scan and Implosions Performance 191
6.3.5 Tritium Lean Implosions with Auxiliary Heating 192
6.4 Summary 194

6.1 Introduction

The beam-plasma system explored in the previous chapter describes the ability to transfer energy from an electron beam into a background plasma collisionlessly via the excitation of Langmuir waves. It is therefore natural to consider whether this mechanism can be harnessed to deposite additional thermal energy into an ICF hotspot via an externally driven electron beam to augment the fusion yield.

This line of reasoning is not unique, and has been explored substantially in the context of Fast Ignition (FI) ICF [70, 185, 186]. However, the energy deposition mechanism in FI is primarily collisional, where high energy charged particles (ions or electrons) deposite their energy via Coulomb collisions with a dense, isochorically compressed fuel assembly. In contrast, the auxiliary heating (AH) approach explored here relies on collective beam-plasma interactions to deposite energy into a lower-density, higher-temperature (and, therefore, less collisional) hotspot.

This chapter first (Section 6.2) explores the feasibility of the AH scheme by considering the generation of the electron beam, its transport through the plasma, and its coupling to the hotspot. The key requirement for the AH scheme is that the beam must arrive at the hotspot with a momentum distribution capable of driving the bump-on-tail instability, and with sufficient density to ensure a significant growth rate. The analysis presented here therefore focuses on assessing whether these conditions can be met in a realistic ICF implosion scenario, and identifying the key processes that govern the effectiveness of the scheme.

The potential performance benefits of auxiliary heating are then explored in Section 6.3. In particular, a modified energy deposition model is introduced, in which heating is localised at the hotspot-shell interface rather than distributed throughout the hotspot volume. The impact of this model on fusion yield is assessed, and the potential of auxiliary heating to enhance implosions with reduced tritium content is also investigated.

6.2 Electron Beam Generation, Transport, and Hotspot Coupling

The AH scheme relies on the ability of a relativistic electron beam to deposit energy into the hotspot via collective beam-plasma interactions. As demonstrated in Chapter 5, this requires not only that energetic electrons reach the hotspot, but that they do so with sufficient density and with a momentum distribution capable of driving unstable modes.

The central question addressed in this section is therefore: can a realistically generated electron beam obtain/retain the properties required to drive instability after transport through the surrounding plasma?

To answer this, the evolution of the electron beam is considered as a sequence of coupled processes. Electrons are first generated at the relativistic critical density surface through intense laser-plasma interactions. They must then propagate through the coronal plasma and dense fuel shell, during which collisional effects, scattering, and collective fields modify their energy and momentum distribution. A key feature enabling the propagation of these high-current electron beams through dense plasma is the formation of a background return current. As the relativistic electron beam propagates, the background plasma electrons respond on plasma timescales by establishing a return current that approximately cancels the beam current, thereby maintaining quasi-neutrality and preventing the build-up of large electrostatic and magnetic self-fields that would otherwise rapidly disrupt beam propagation. Although the net current is therefore small, the finite resistivity of the background plasma gives rise to resistive electric and magnetic fields that influence beam transport, collimation, and energy loss. These collective transport effects are discussed where relevant throughout the following sections. Simultaneously, the beam diverges, reducing its density and causing some of the beam to miss the hotspot.

These processes have competing effects on the viability of the AH mechanism. Collisional scattering can reshape the electron distribution in a manner crucial to the success of the scheme, while beam divergence and transport losses act to

dilute the beam. The effectiveness of the scheme therefore depends on the balance between distribution shaping and density reduction during transport.

In this section, these effects are combined to construct a self-consistent estimate of the electron beam properties at the hotspot, including its density, energy spectrum, and spatial profile. From this, the ability of the beam to drive instability and deposit energy into the hotspot is assessed.

6.2.1 Electron Beam Generation

The AH scheme requires the generation of a population of energetic electrons capable of penetrating to the hotspot. In the configuration considered here, these electrons are produced through high-intensity laser-plasma interactions at the relativistically corrected critical density surface.

At sufficiently high laser intensities, the electron motion in the oscillating electric field becomes relativistic. By considering the electrons oscillating in a large amplitude electric field and averaging over an oscillation cycle, the relativistic Lorentz factor is given by $\gamma = \sqrt{1 + a_0^2/2}$, where $a_0 = eE_0/m_e c$ is the normalised vector potential. This corresponds to a shift in electron mass, and therefore critical density, given by $n_{c,rel} = \gamma n_c$. As a result, laser energy is absorbed at a higher density surface closer to the hotspot than would otherwise be possible.

The interaction of a short-pulse, high-intensity laser with the plasma at the relativistic critical density surface leads to the generation of a forward-directed population of fast electrons. The characteristic energy of this population can be described by a hot electron temperature T_{hot}, which depends on the laser intensity and wavelength. There are two commonly used scalings for T_{hot}. The first, derived by Wilks et. al. [187], follows the ponderomotive scaling [188], which gives a characteristic temperature of

$$T_{\mathrm{hot}}^{\mathrm{pond}} = m_e c^2 \left(\sqrt{1 + \frac{I\lambda_{\mu\mathrm{m}}^2}{1.37 \times 10_{18}}} - 1\right). \tag{6.1}$$

This was later shown experimentally to be an overestimate of the characteristic temperature. A more accurate experimental fit was provided by Beg et. al. [189],

and further validated experiementally at intensities up to $I \sim 10^{20}$ W/cm^2 [190], which follows the scaling

$$T_{\text{hot}}(\text{keV}) = 215(I_{18}\lambda_{\mu\text{m}}^2)^{\frac{1}{3}}. \tag{6.2}$$

The initial electron distribution is then described by a thermal bulk and a high energy tail characterised by the hot electron temperature. The distribution of this high energy tail can be modelled as

$$f_E(E) \propto \exp\left(-\frac{E}{T_{\text{hot}}}\right). \tag{6.3}$$

While this high energy tail exhibits a region of positive gradient in momentum space, when including the low energy bulk component, the overall distribution is monotonically decreasing with momentum, and therefore stable to the bump-on-tail instability. Modification of this distribution during transport is therefore required for instability-driven heating to occur.

The laser absorption at the relativistic critical density surface can be quite efficient, with $\sim$ 10%-40% [191, 192] of the laser energy being deposited. While the actual efficiency of the laser deposition depends on a range of parameters, such as the plasma lengthscale and pulse duration, the empirical fit [193]

$$\xi = 0.108(I_{18}\lambda_{\mu\text{m}}^2)^{\frac{1}{4}}, \tag{6.4}$$

provides a useful estimate for the coupling efficiency in the range of intensities relevant to short-pulse ignition (10^{18} - 10^{20} W/cm^2). The total number of electrons in the beam is determined by the fraction of laser energy converted into fast electrons and the average electron energy

$$N = \frac{\xi E_{\text{laser}}}{E_{\text{avg}}}. \tag{6.5}$$

In addition to setting the energy spectrum and particle number, the generation process also determines the initial angular spread of the beam, which influences its subsequent transport and coupling to the hotspot. The primary cause of beam divergence is not fully understood, with a range of possible mechanisms being

possible. A leading candidate is the interaction surface between the laser and plasma. As the laser bores into the plasma, the interaction surface is not flat, and the electrons are accelerated normal to the surface, leading to a range of initial angles. This picture qualitatively explains both the observed increase in divergence with laser intensity and the correlation between electron energy and emission angle, whereby electrons emitted at larger angles generally possess lower energies.

Experimental measurements provide an empirical scaling for the divergence angle as a function of laser intensity [193, 194]:

$$\theta_{\frac{1}{2}}(I) = 15 + 13\log_{10}(I_{18}\lambda_{\mu\mathrm{m}}^2). \tag{6.6}$$

This equation is valid for intensities relevant to short-pulse ignition (10^{18} - 10^{21} W/cm^2) using wavelengths of $\lambda = 1.053$ μm, and corresponds to divergence angles in the range 15° to 54°, implying substantial beam spreading over typical transport distances. More recent experimental studies have continued to investigate fast-electron beam divergence under ignition-relevant conditions and suggest that the empirical scaling of Equation 6.6 may overestimate the divergence in some regimes [195].

These equations define the initial conditions for the electron beam: its energy distribution, characteristic temperature, total particle number, and angular spread. These properties form the input to the transport and evolution processes considered in the following sections, and determine whether the beam can retain sufficient density and an appropriate distribution to drive instability at the hotspot.

6.2.2 Channel Formation

To maximise the density of the beam at the hotspot and ensure a significant portion of the laser energy arrives at the critical density surface, the laser must be guided through the surrounding, underdense coronal plasma and be deposited as close to the hotspot as possible. This can be acheived through the use of high-Z cones [196–199], which can be inserted into the capsule to reduce the hydrodynamic flow into the

cone, producing a low density channel for the laser to propagate through. This is the approach often considered in the FI scheme, and could be similarly applied here.

However, hotspot based implosion designs, to which AH would be applied to, typically have higher symmetry requirements than FI designs, which may inhibit the use of cones. An alternative approach is to use laser-driven hole boring [200, 201], which uses the ponderomotive force of the laser itself:

$$F_p = -\frac{e^2}{4m_e\omega^2}\nabla E^2. \tag{6.7}$$

The ponderomotive force acts to move charged particles away from regions of high laser intensity to regions of lower intensity. Therefore, a laser pre-pulse [202], or multi-pulse, approach [203, 204] with a radial spatial intensity gradient can act to evacuate a channel, producing a low density path that the main ignitor pulse can propagate through.

Low density channels formed via laser hole boring have been shown to provide additional benifits, including electron beam collimation effects, self focusing, and instability mitigation. Due to the refractive index (Equation 2.50) of a plasma being less than unity, as the material is expelled from the high intensity regions of the laser pulse, the refractive index in these regions increases. This produces a waveguide effect, where rays bend towards the centre of the channel, focusing the laser pulse and increasing its intensity. This self-focusing effect is not diffraction limited, and can produce extremely high intensities within the channel. Due to the increased intensity, the critical density surface is pushed further into the plasma, allowing the laser to bore deeper. This feedback mechanism, referred to as the HKT mechanism, was first explored by Habara [205], Kodama [192], and Tanaka [191], who demonstrated 'super-penetration' of high intensity laser beams into overdense plasma at the Institute of Laser Engineering (ILE), Osaka University.

Through the HKT mechanism, the pulse can also be made robust to filimentation and hosing instabilities. Filliments form on lengthscales comparable to the plasma wavelength, and cause the pulse to break up into multiple separate filaments. If the spatial width of the pulse is reduced below the characteristic width of a

single filament, the pulse can propagate without breaking up [206]. As the beam travels down the channel, the plasma density increases, and the plasma wavelength decreases. For the beam to remain stable against filamentation, the beam width must self focus at least as fast as the plasma wavelength decreases. This has been demonstrated by Spiers et al. [206] on the Orion laser facility at AWE, demonstrating the formation of long-lived channels formed by the HKT mechanism.

The hosing instability arises from transverse perturbations across the beam profile, causing different regions of the beam to experience different refractive indices, leading to the beam bending. It is known that hosing becomes more severe as the density of the plasma increases [207, 208]. However, work by Ceurvorst [209] has shown that the increased effective electron mass in relativistic laser pulses also acts to mitigate hosing.

The primary role of channel formation in the context of auxiliary heating is to reduce the distance between the electron source and the hotspot. As shown in Section 6.2.5, beam divergence leads to a rapid reduction in beam density with propagation distance. Channel formation and the HKT mechanism can therefore act to mitigate this density loss by enabling the laser pulse, and hence the electron source, to penetrate deeper into the plasma.

6.2.3 Electron Collisional Scattering and Distribution Evolution

While an energetic electron distribution can be generated at the relativistic critical density surface, it must then propagate through the remaining coronal plasma and dense fuel shell to reach the hotspot. During this transport, the electrons are subject to a range of effects including collisional scattering. While the majority of these processes will be detrimental to the AH process, collisional scattering is expected to be crucial for the success of the scheme. As explained in Section 5.5.2, the heating mechanism relies on the bump-on-tail instability, in which a positive gradient in the electron momentum distribution function drives the growth of Langmuir waves. The initial electron beam distribution is monotonically decreasing with momentum, and

therefore stable to the bump-on-tail instability. However, as the beam propagates through the plasma, lower energy electrons are preferentially scattered out of the beam due to their higher collisional cross-section (see e.g. Equation 2.40). It is therefore possible that a positive gradient in the momentum distribution function can form as the electrons travel towards the hotspot.

Quantifying the evolution of the electron distribution function is a complex task, as it requires a fully kinetic treatment of electron transport through the plasma between the absorption surface and the hotspot, including the high-density shell surrounding the hotspot. The primary challenge in modelling this region is the extremely short Debye length (Equation 2.8), which is on the order of 10 pm. Resolving this scale is necessary to accurately capture the response of electric fields to charge imbalances and to avoid numerical instabilities and spurious heating. However, since the characteristic hotspot and shell scales are ~ 100 μm, this would require on the order of 10^7 spatial cells per dimension, along with ~ 100 particles per cell, which is currently computationally infeasible.

Therefore, to gauge the feasibility of generating a suitable distribution function via collisional scattering, an analytic collisional model was used and solved numerically. The collisional model uses the relativistic treatment for Coulomb scattering first introduced by Mott [210], and refined by Rosen [211]. The leading-order differential cross section is

$$\frac{\mathrm{d}\sigma}{\mathrm{d}\Omega}(\theta) = \frac{Z^2 e^4}{4\left(4\pi\epsilon_0\right)^2 m^2 v^4}\left(1-\beta^2\right)\csc^4\left(\frac{\theta}{2}\right). \tag{6.8}$$

In the case of scattering, the longitudinal momentum loss from scattering through an angle θ is $\Delta p_{\parallel} = (1-\cos(\theta))p$, and in the limit of small angle scattering events, $\theta \ll 1$, $\sin(\theta) \sim \theta$, $\csc^4(\theta/2) \sim 16/\theta^4$, and $1-\cos(\theta) \sim \theta^2/2$. The momentum-transfer cross section is therefore given by

$$\sigma_{tr} = 2\pi \int (1-\cos(\theta))\frac{\mathrm{d}\sigma}{\mathrm{d}\Omega}(\theta)\sin(\theta)\,\mathrm{d}\theta \tag{6.9}$$

$$= \frac{\pi Z^2 e^4}{2\left(4\pi\epsilon_0\right)^2 m^2 v^4}\left(1-\beta^2\right)\int \frac{\theta^2}{2}\frac{16}{\theta^4}\theta \mathrm{d}\theta \tag{6.10}$$

$$= \frac{\pi Z^2 e^4}{2\left(4\pi\epsilon_0\right)^2 m^2 v^4}\left(1-\beta^2\right)\int_{\theta_{\min}}^{\theta_{\max}} \frac{8}{\theta}\mathrm{d}\theta. \tag{6.11}$$

This integration from $\theta_{\min}$ to $\theta_{\max}$ is similar to the $b_{\min}$ and $b_{\max}$ impact parameters used in Section 2.2.4 defining the Coulomb logarithm (Equation 2.38). As a final result, the momentum transfer cross section is therefore given by

$$\sigma_{tr} = \frac{4\pi Z^2 e^4}{\left(4\pi\epsilon_0\right)^2 m^2 v^4}\left(1-\beta^2\right)\ln\left(\Lambda\right) = \frac{Z^2 e^4}{4\pi\epsilon_0^2 m^2 c^4}\frac{\ln\left(\Lambda\right)}{\beta^4\gamma^2}, \tag{6.12}$$

where $\gamma = (1-\beta^2)^{-1/2}$.

The equation for the scattering cross section can then be used to derive a probability that an electron propagates without collisions. Assuming that collision events are independent and occure over a small length dx, they can be treated as a Poisson process. The probability of an electron with momentum p surviving to position x is then given by:

$$P(p,x) = \exp\left[-\int_{x_0}^{x}\sigma_{tr}(p,x')\,n(x')\,\mathrm{d}x'\right], \tag{6.13}$$

where $n(x')$ is the background plasma density at position x'. The electron beam flux, $\Gamma(p,x)$, as a function of momentum, p, and position, x, can then be modelled as

$$\Gamma(p,x) = \Gamma_0(p)\,P(p,x), \tag{6.14}$$

where $\Gamma_0(p)$ is the initial beam flux at the point of generation.

For an initial electron flux distribution, $\Gamma_0(p)$, and a plasma temperature, $T(x)$, and density, $\rho(x)$, profile that the electron beam propagates through, the electron flux evolution can be determined numerically. The electron distribution is discretised into momentum bins, and the temperature and density profiles can be divided into cells of width Δx and an average taken. For each momentum bin, the corresponding β and γ factors are calculated.

Within each spatial cell, the Coulomb logarithm is computed locally. The minimum impact parameter is taken as the maximum of the classical closest-approach distance (Equation 2.35, evaluated using the relativistic kinetic energy) and the quantum de Broglie wavelength (Equation 2.36), as defined in Equation 2.37. The maximum impact parameter is given by the Debye length (Equation 2.8). The

Coulomb logarithm is then obtained from the ratio of these scales, allowing the scattering cross section to be evaluated within each cell.

The probability that an electron with momentum p traverses a given cell without undergoing a collision is computed from Equation 6.13, using the local density and cross section. The cumulative survival probability to a position x' is obtained by replacing the integral in Equation 6.13 with a discrete sum over all upstream cells. Repeating this procedure for each momentum bin allows the flux distribution to be evolved according to Equation 6.14. This procedure yields the evolution of the electron momentum distribution as the beam propagates through the plasma, allowing assessment of whether a positive gradient sufficient to drive instability can be generated at the hotspot.

6.2.4 Beam Divergence and Density Evolution

A major factor limiting the effectiveness of electron beam driven heating mechanisms is the divergence of the beam during transport. As the beam propagates, it will naturally diverge due to an initial angular spread. This divergence reduces the density of the beam as it reaches the hotspot, which in turn reduces the growth rate of the instability.

This reduction in beam density has a direct impact on the auxiliary heating mechanism. As shown in Chapter 5, the growth rate of beam-plasma instabilities depends on the ratio of beam density to background plasma density (Equation 5.50). In the limit of sufficiently low density, the instability is suppressed altogether, and no collisionless heating will occur. Beam divergence therefore acts as a key degradation mechanism for auxiliary heating.

To quantify this effect, consider a beam with an initial radius of r_0 and a half-angle divergence of $\theta_{1/2}$. After propagating a distance L, the beam radius will have increased to

$$r = r_0 + L \tan\left(\theta_{1/2}\right), \tag{6.15}$$

where the divergence half-angle is defined as in Equation 6.6. The severity of this effect depends strongly on the transport distance, L, with longer propagation distances leading to significantly greater beam dilution.

Assuming the beam electrons are travelling close to the speed of light and travelling distances around 500 μm, they have a time of flight of ~ 1 ps. As the beam is driven for ~ 10 ps, the beam can be assumed not to diverge along the axis of propagation. Using this, and that the total number of electrons in the beam remains constant, the density of the beam will decrease as $n_b \propto (r_0/r)^2$. Even modest divergence angles can therefore result in significant dilution of the beam over the transport distance to the hotspot.

The total number of electrons in the beam is

$$N = \int_0^{c\tau} n(x)A(x)\,\mathrm{d}x, \tag{6.16}$$

where the maximum distance the electrons can propagate to is $c\tau$ where c is the speed of light, and τ is the laser pulse duration. Assuming the electron beam does not diverge parallel to the propagation axis (i.e. all electrons travel forwards at the same speed), and that the beam does not undergo collisions, then the beam density satisfies the continuity equation $n(x)A(x) = \text{const.}$, giving

$$n(x) = \frac{N}{A(x)c\tau}. \tag{6.17}$$

This describes the electron beam occupying a cone of increasing radius with length $c\tau$ and a cross sectional area of $A \propto r^2$.

The total number of beam electrons generated from the interaction is also given by Equation 6.5 by considering the conservation of energy. Putting these two equations together results in the following expression for the beam density as a function of position:

$$n(x) = \frac{\xi E_{\text{laser}}}{c\tau E_{\text{avg}}}\frac{1}{A(x)}. \tag{6.18}$$

In addition to density reduction, large divergence angles introduce a further loss mechanism: electrons may fail to intersect the hotspot entirely. This geometric

loss becomes increasingly important for large divergence, small hotspot radii, or large propagation distances, effectively reducing the fraction of the beam that contributes to heating.

Mitigation of beam divergence is therefore essential for the viability of the auxiliary heating scheme. As discussed in Section 6.2.2, the formation of low-density channels via laser hole boring provides a mechanism for reducing the effective transport distance, L. The use of a multi-pulse approach, also discussed in Section 6.2.2, can further enhance beam collimation by generating magnetic fields that act to focus the beam. However, as will be shown in the following section, even with these mitigation strategies, beam divergence remains a dominant limitation on the achievable beam density at the hotspot.

6.2.5 Electron Beam Properties at the Hotspot

This section combines the processes described in the previous sections into a practical example to assess the feasibility of a realisitically generated electron beam driving the AH process in the hotspot. In particular, the key requirement identified in Chapter 5 is that the beam must arrive with a momentum distribution capable of driving the bump-on-tail instability, and with sufficient density to ensure a significant growth rate.

The analysis is based on the optimised configuration corresponding to the smallest wetted foam capsule considered in Chapter 4. This is the capsule size in which the least fusion energy is produced, and therefore the case most likely to benefit from additional heating. Futhermore, the reduced capsule size minimises the transport distance and therefore mitigates beam divergence. As such, it provides a best-case estimate for the viability of the scheme.

Density and temperature profiles from this implosion are extracted at peak compression and are used to define the background plasma through which the electron beam propagates. These profiles span the region from the relativistically corrected critical density surface to the hotspot and form the basis for modelling both collisional transport and distribution evolution. They are shown in Figure 6.1. The

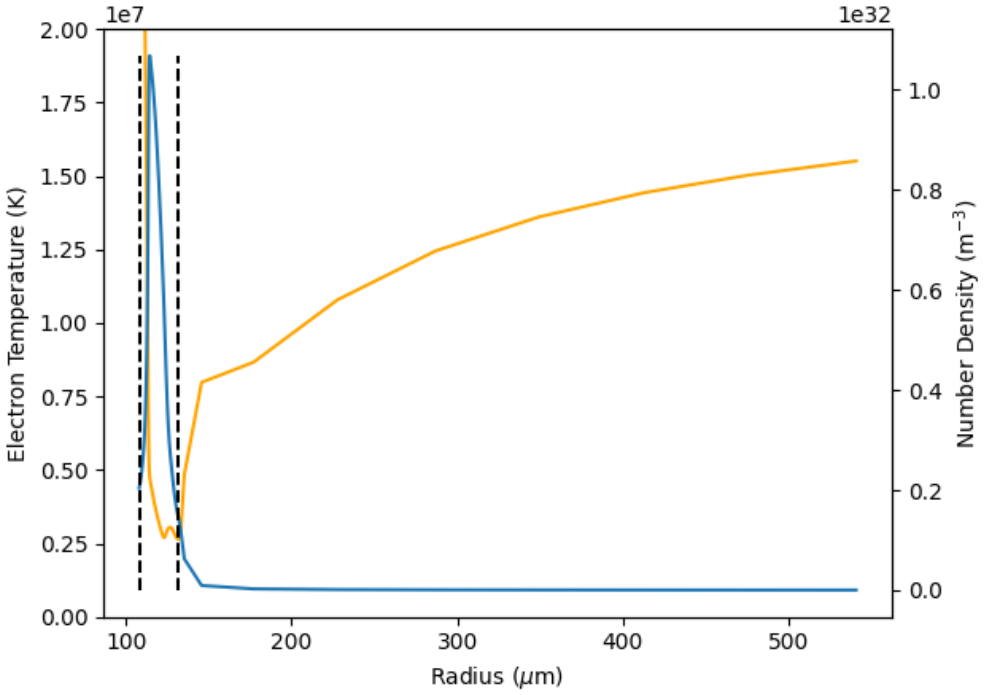


Figure 6.1: Density and temperature profiles extracted from the optimised $\sim$ 2 MJ implosion at peak compression. The relativistically corrected critical density surface is located at $\sim$ 550 μm, and the hotspot boundary is located at $\sim$ 150 μm.

background plasma density in the hotspot is just below 10^{31} m^{-3}, therefore requiring electron beam densities around 10^{29} m^{-3} to drive significant instability heating.

Using FI relevant laser parameters, the properties of the initial electron beam can be estimated. For the configuration considered here, a laser energy of $\sim$ 20 kJ, duration $\sim$ 15 ps, wavelength $\sim$ 1.053 μm, and spot size radius $\sim$ 20 μm. This corresponds to an intensity of $I \sim 10^{20}$ W/cm^2 and power $P \sim 1.3$ PW.

From this, the characteristic hot electron temperature is calculated using Equation 6.2, yielding $T_{\text{hot}} \sim 1$ MeV. The initial electron energy distribution is therefore described by a thermal bulk with an exponential high-energy tail of the form $f_E(E) \propto \exp(-E/T_{\text{hot}})$, as introduced in Section 6.2.1. The average energy of this distribution (calculated on the range of 0 - $10T_{\text{hot}} \approx 10$ MeV) is $E_{\text{avg}} = 1.37$ MeV.

The predicted coupling efficiency from laser energy to energetic electrons is $\xi = 35.6$ %, calculated using Equation 6.4. The initial density of the beam at the point of generation can then be calculated using Equation 6.18, yielding $n_b = 5.76 \times 10^{27}$ m^{-3}. This corresponds to $n_b/n_0 \sim 10^{-3}$ at the point of generation,

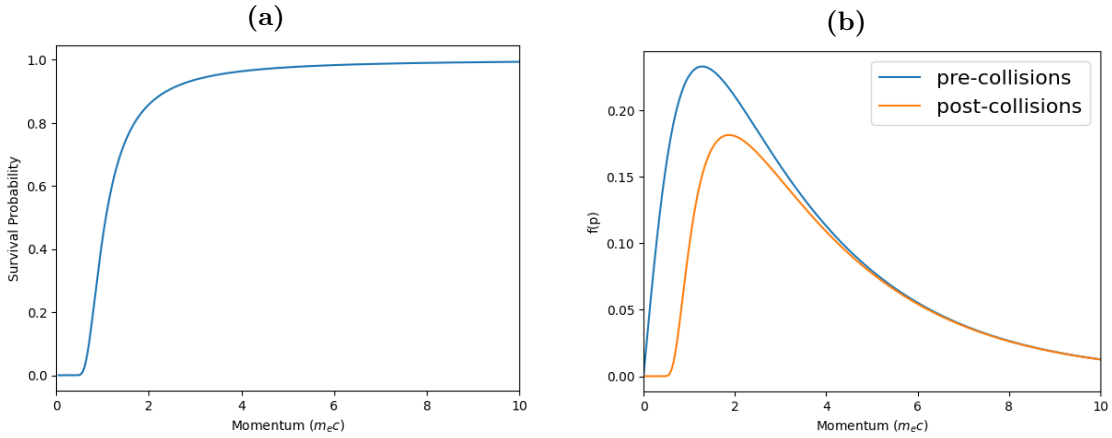


Figure 6.2: (a) Survival probability of electrons propagating to the hotspot as a function of momentum. (b) Initial and resulting momentum distributions pre- and post-collisions. The initial distribution is shown in blue, while the post-collisional distribution is shown in orange.

placing the system firmly in the weak beam regime discussed in Chapter 5. As such, even before transport effects are considered, the beam density is already close to the lower limit required for meaningful instability growth.

As the beam propagates towards the hotspot, its properties are modified by two primary mechanisms: collisional scattering and geometric divergence. The collisional scattering is modelled using the approach described in Section 6.2.3, which calculates the survival probability of electrons as they propagate through the plasma. The survival probability as a function of electron momentum is shown in Figure 6.2(a), while the initial and resulting momentum distributions pre- and post-collisions are shown in Figure 6.2(b). It can be seen that electrons below $\sim 1\, m_e c$ are heavily depleted by collisional scattering, while the high energy tail of the distribution is largely unaffected. This results in a positive gradient forming in the momentum distribution between $\sim 0.5\, m_e c$ and $\sim 2\, m_e c$.

The major limiting factor for the AH scheme is the divergence of the beam during transport, which both acts to reduce the density of the beam at the hotspot, and causes a portion of electrons to miss the hotspot entirely. Equation 6.6 gives a divergence half angle of 41.3°. The laser must then propagate a distance of ~ 500 μm from the critical density surface to the hotspot, which results in the

beam radius increasing from 20 μm to 443 μm. This corresponds to a reduction in beam density by a factor of $443^2/20^2 = 471$. After accounting for divergence, the beam density at the hotspot is reduced to $n_b \sim 10^{25}$ m^{-3}, corresponding to $n_b/n_0 \sim 10^{-6}$. At this level, the growth rate of the instability is expected to be negligible, effectively suppressing collisionless heating.

Furthermore, the radius of the hotspot is only 80 μm, and the radius of the beam at the hotspot is 443 μm, meaning that only ~ 3 % of the beam will actually intersect the hotspot. Divergence is therefore a critical issue for the AH scheme, and mitigation stratagies are crucial for the success of the scheme. As it stands, the impact of divergence is so severe that the AH scheme is unlikely to be a viable approach to augmenting ICF implosions.

Overall, this analysis demonstrates that a fraction of the electron beam can reach the hotspot with a modified momentum distribution capable of driving instability. However, beam divergence drastically reduces the beam density, leading to values far below those required for significant instability growth. As a result, divergence emerges as the dominant limiting factor for the auxiliary heating scheme.

Attempts to mitigate this through simple adjustments to laser parameters are ineffective. Increasing the laser energy primarily raises the characteristic electron temperature and divergence angle, rather than the beam density. Similarly, increasing the pulse duration increases the total number of electrons, but not the density, as the beam length increases proportionally. Reducing the spot size leads to higher intensities, again increasing both temperature and divergence. Therefore, improving the viability of the auxiliary heating scheme requires fundamental changes to the properties of the electron beam, particularly reductions in divergence and/or modifications to the initial distribution. However, even if the propagation distance, L, and angle of divergence, θ, are both reduced by a factor of 2, the beam density at the hotspot increases by an order of magnitude, but still remains well below the required level for significant instability growth. This highlights the severity of the divergence issue and the need for significant mitigation strategies to make the AH scheme viable.

6.2.6 Improving Electron Beam Transport and Divergence

As is evident from the above analysis, the transport and divergence of the electron beam is a critical issue for the viability of the auxiliary heating scheme. This is a general issue for most heating mechanisms that rely on energetic electron beams, including the electron-FI scheme. Significant work has therefore been conducted on improving the transport and collimation of laser-driven electron beams.

As discussed in Section 6.2.2, the formation of low-density channels via the HKT mechanism can reduce the transport distance, and significantly reduce the beam divergence angle. The use of a pre- or multi-pulse approach to generate the channel can further enhance beam collimation. The electrons accelerated by the pre-beam are lower energy and more collimated, and generate magnetic fields that act to focus the primary electron beam [202, 204]. Work by Scott et al. [212] has shown that a pre-pulse with energy on the order of 500 J can generate magnetic fields that can collimate a main drive beam over distances of hundreds of microns.

Resistivity gradients in the plasma have also been explored a potential method to collimate electron beams [213]. By using a high-resistivity core and a low-resistivity cladding, magnetic fields are generated at the interface, providing a waveguide effect to confine the electron beam along the core. By embedding resistivity gradients into the capsule design, this approach could be used to collimate the electron beam as it propagates towards the hotspot. This may, however, have severe implications in the symmetry of the implosion, and therefore may not be suitable for hotspot-based implosions.

Finally, externally applied magnetic fields have been considered [214]. As the capsule implodes, the magnetic field is compressed and amplified, providing a mechanism to collimate the electron beam as it propagates towards the hotspot.

While these transport-enhancement mechanisms are not considered in the present work, they demonstrate that electron beam divergence may not necessarily be a fundamental limitation of auxiliary heating approaches. Continued developments in relativistic electron beam transport and guiding physics may therefore substantially

improve the feasibility and coupling efficiency of auxiliary heating schemes for inertial confinement fusion implosions.

6.3 Implementation of Auxiliary Heating in Low-CR Simulations

The conclusions of the above section reflects limitations in beam transport rather than the fundamental effectiveness of the heating mechanism itself. This section considers the potential impact of AH on an implosion isolated from the challenges associated with beam generation and propagation.

6.3.1 Previous Exploration of AH in ICF Implosions

Previous studies exploring the impact of AH in ICF implosions have typically assumed a uniform energy deposition model, in which the electron beam deposits its energy throughout the hotspot volume [215, 216]. However, as shown in Chapter 5, the instability responsible for this energy transfer operates on extremely short timescales, implying a correspondingly small spatial scale of the energy deposition. As a result, the energy is expected to be deposited in a thin layer near the edge of the hotspot, rather than uniformly throughout its volume. This section therefore revisits previous AH studies under a more physically motivated energy deposition model, and assesses how this modification alters the predicted impact of AH on implosion performance.

The AH scheme was previously applied to low-CR wetted foam implosions similar to those discussed in Chapter 4, with long-pulse drive energies ranging from 0.1 to 1.7 MJ [215]. In these studies, thermal energy was deposited into the hotspot over a short timescale (7 ps) near peak compression. A linear scan over deposition time was performed to identify the optimal injection point, showing that the optimum occurs shortly before peak compression. While the precise timing varies with implosion scale, it was found to be largely independent of the thermal energy deposited.

Due to uncertainty in the practical implementation of the heating scheme, two coupling efficiency values were considered: a ‘high’ coupling efficiency of 9.4 %, and

a 'low' coupling efficiency of 2.7 %. The high coupling efficiency case corresponds to the best-case scenario for the AH scheme, where the laser energy is aborbed with an efficiency of 52 %, while the low efficiency case corresponds to a more pessimistic scenario, where the laser energy is absorbed with an efficiency of 15 %. Regardless of the energy absorption efficiency, the resulting electron beam is assumed to deposit its energy in the hotspot with an efficiency of 18 %. In both of these cases, there are no assumptions about the efficiency of the beam transport, which was found to be a major limiting factor in the previous section. These efficiency values are relevant when considering the total gain of the implosion. For the results presented in this thesis, only the yield amplification will be considered, as no suitable transport coupling efficiency can be assumed.

Further studies [216] have investigated which implosions are most likely to benefit from auxiliary heating. These works identified a strong correlation between the effectiveness of AH and the burn-up fraction of the fuel. The burn-up fraction, introduced in Equation 4.6, depends inversely on the parameter H_b, which is typically taken to be ~ 7 g/cm^2 for a burn temperature of $T_{\text{burn}} \sim 20$ keV. However, a more accurate temperature-dependent expression for H_b is given by [217]:

$$H_b = 0.866\sqrt{T_{\text{burn}}}\exp\left(0.572\ln\left(\frac{64.2}{T_{\text{burn}}}\right)^{2.13}\right), \tag{6.19}$$

where T_{burn} is the burn averaged (spatially and temporally) ion temperature in keV.

The AH scheme deposits energy directly into the hotspot and therefore does not modify the compressed areal density, ρR. Consequently, any increase in burn-up fraction must arise through a reduction in H_b, i.e. through an increase in the hotspot temperature. The temperature dependence of H_b implies that the burn-up fraction is most sensitive to heating in the range $T_{\text{burn}} \approx$ 5-20 keV, with the strongest response occurring near $T_{\text{burn}} \approx 10$ keV. This corresponds closely to the temperature regime of marginally igniting or near-igniting implosions. Auxiliary heating is therefore expected to be most effective in implosions that are close to ignition, where modest increases in hotspot temperature can lead to disproportionately large increases in burn-up fraction and yield.

6.3.2 Interface vs Hotspot Energy Deposition

As was shown in Chapter 5, the instability operated on timescales of 1000 ω_{p}^{-1}. Using Equation 2.26, the duration of the interaction is therefore ~ 11 fs. Assuming the beam moves at the speed of light, the interaction width is on the order of microns. The impact this reduced energy deposition volume has on the implosion performance is therefore assessed. In the new approach, the energy is deposited in a thin 5 μm shell at the edge of the hotspot. This alternate energy deposition model will be referred to as interface heating, with the previous approach being hotspot heating.

Snapshots of the hotspot and interface heating can be seen in Figures 6.3 and 6.4 respectively. In each image, the black dashed lines mark the inner and outer boundaries of the shell. The hotspot is defined as the region enclosed by the inner shell boundary, while the shaded region indicates the cells over which the auxiliary heating energy is deposited. In these simulations, a total of 20 kJ of thermal energy is deposited. It can be seen in the case of interface heating that the thermal energy is initially deposited in a thin region at the interface between the hotspot and shell, but rapidly expands to fill the hotspot. After $\mathcal{O}(10$ ps) the interface and hotspot temperature profiles are very similar. However, once the thermal energy has diffused through the hotspot, the interface heating case results in a slightly lower temperature. At the point of injection, however, a significantly higher peak temperature is achieved in the interface heating case. This higher peak temperature appears in the region of highest fuel mass (largest hotspot radius and density). The increased temperature in this region outweighs the reduced temperature in the centre of the hotspot, leading to comparable, if not slightly improved, performance in the interface heating case.

6.3.3 Optimum Deposition Time

Further comparisons can be made between the optimum time of energy deposition. A linear scan over deposition time is run for both models and are compared in Figure 6.5. This is done for depositing 2, 8, and 20 kJ of thermal energy around the bang time, which is defined as the time where the most fusion reactions occur. As

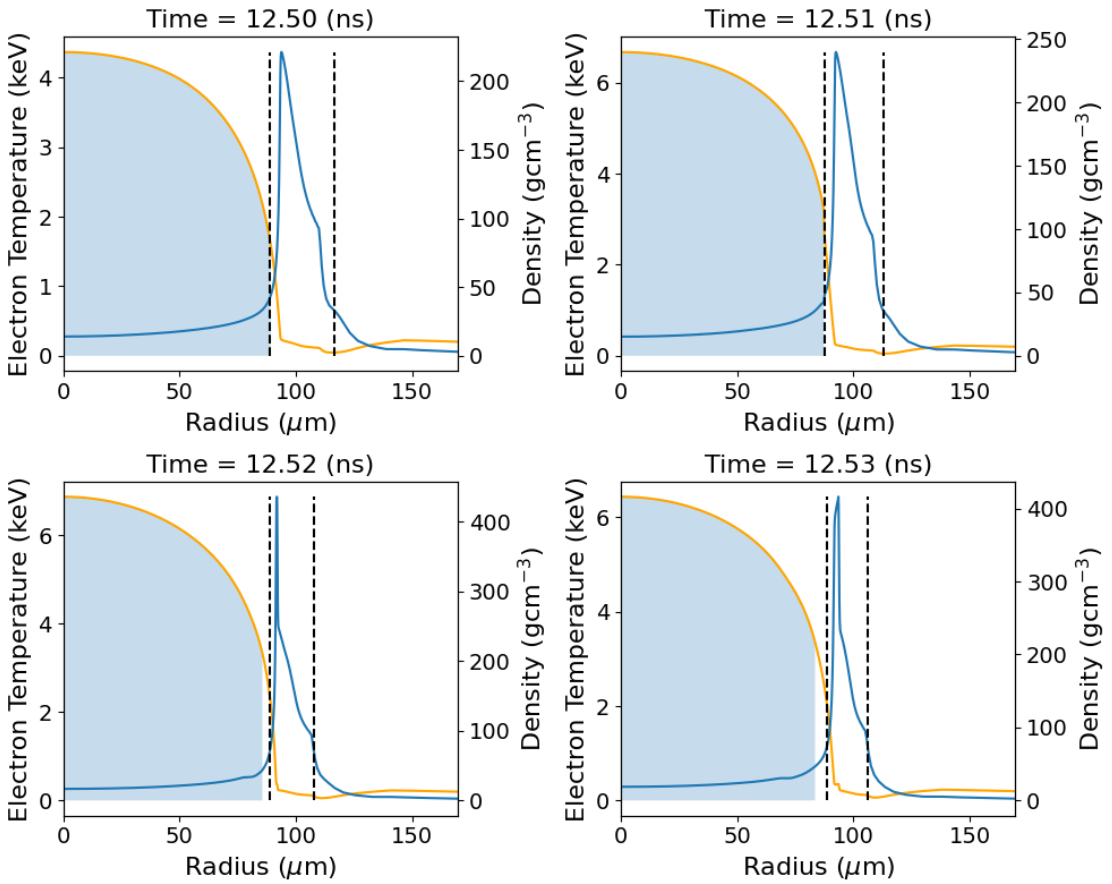


Figure 6.3: Snapshots of the hotspot heating case, thermal energy being deposited in the shaded region.

was stated in [215], the optimum deposition time is roughly constant with respect to the amount of thermal energy deposited. Furthermore, both interface and hotspot heating are in agreement with optimal deposition time.

6.3.4 Energy Scan and Implosions Performance

A linear scan over total deposited thermal energy, ranging from 0 to 100 kJ, is shown in Figure 6.6. From these simulations, for the 0.74 MJ implosion to reach the same yield as the 1.8 MJ implosion, the auxiliary heating approach must amplify the yield by $\sim$ 26 times. To acheived this, the AH scheme must deposit $\sim$ 65 kJ of thermal energy into the hotspot. Therefore, for this approach to be competitive with simply increasing the long-pulse drive energy (in this case by 1.06 MJ), the overall efficiency of the AH scheme (including laser absorption, beam generation, transport,

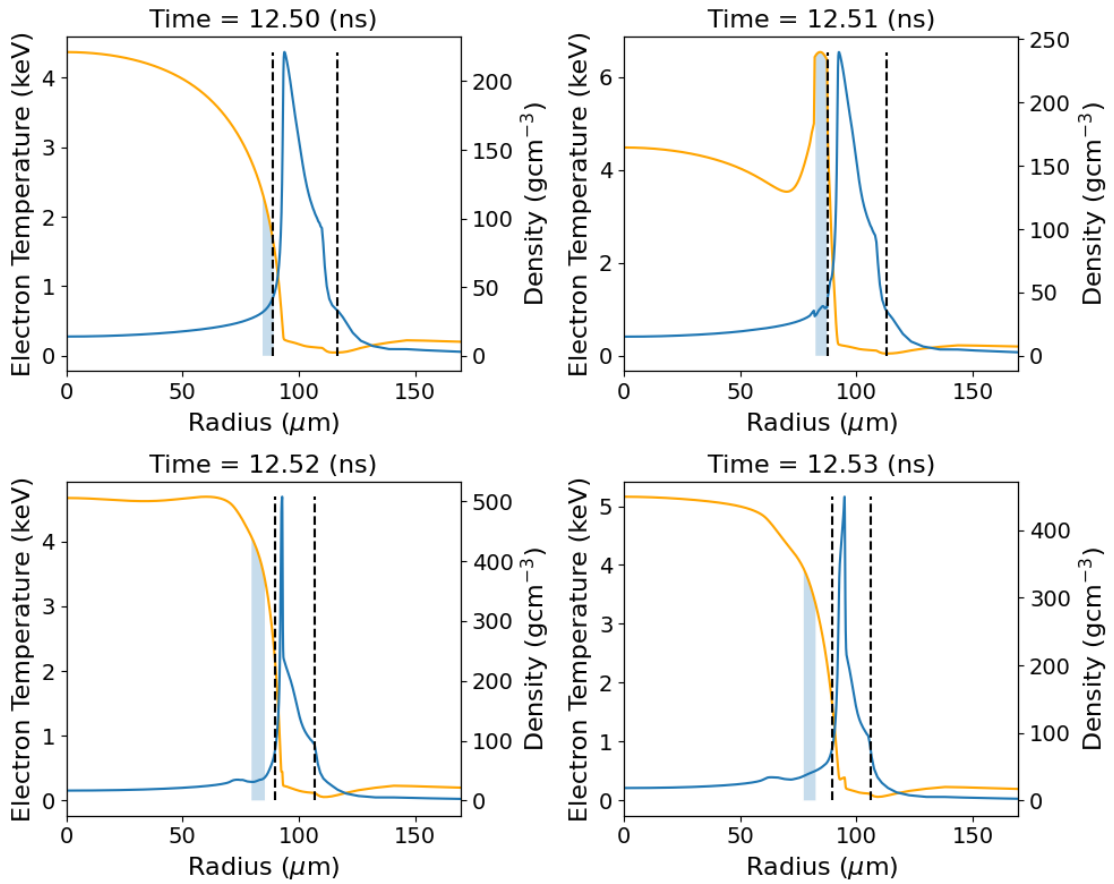


Figure 6.4: Snapshots of the interface heating case, thermal energy being deposited in the shaded region.

and energy deposition) must be larger than $\eta > 6.1$ %. Even under the optimistic assumptions of the previous study, this limit is approached, and surpassed in the less optimistic case. This further highlights the challenges associated with the practical implementation of the AH scheme, and the need for significant improvements in coupling efficiency to make it a viable approach for augmenting ICF implosions.

6.3.5 Tritium Lean Implosions with Auxiliary Heating

As was mentioned in Section 2.5.3, the use of tritium in ICF implosions presents a range of challenges, including its scarcity, handling difficulties, and neutron fusion products. As such, there is significant interest in developing ICF implosions that use a reduced amount of tritium, referred to as tritium lean implosions. In these implosions, there is a higher reliance on the DD fusion reaction, which has a higher

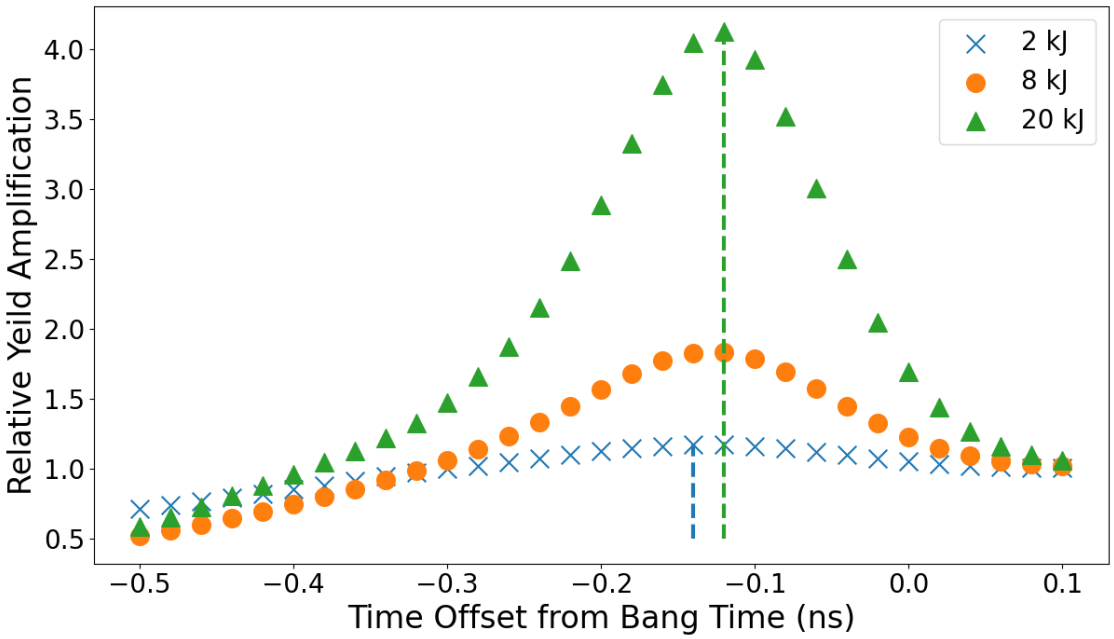


Figure 6.5: Linear scan over deposition time for interface heating. The optimum deposition time is roughly constant with respect to the amount of thermal energy deposited, and both hotspot and interface models are in agreement with optimal deposition time.

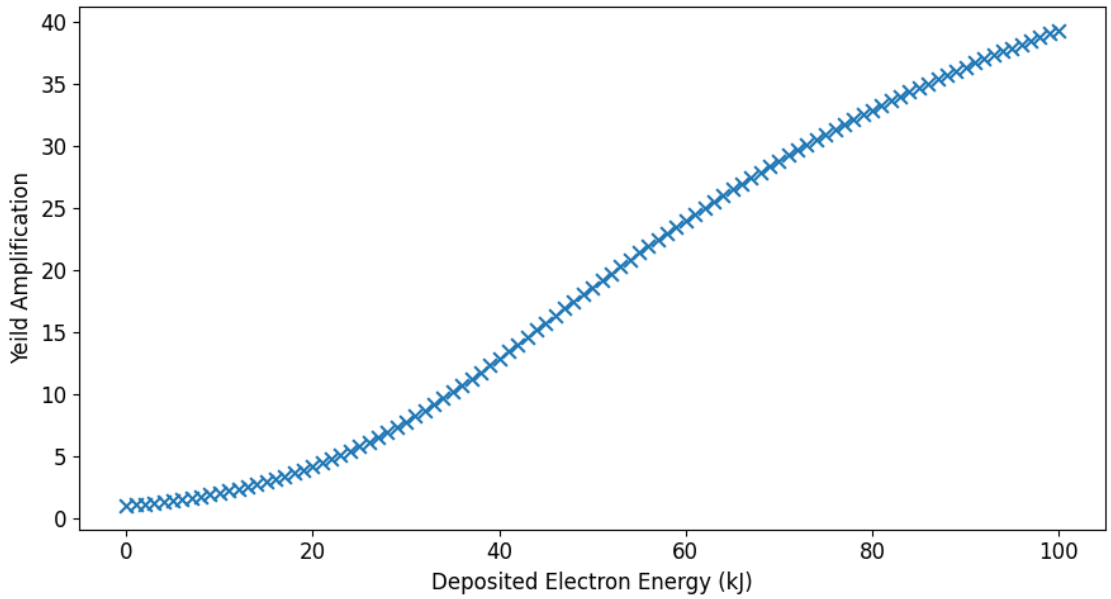


Figure 6.6: Linear scan over total deposited thermal energy for interface heating.

ignition temperature than the DT reaction. As a result, tritium lean implosions are more difficult to ignite than their 50-50 DT counterparts.

The reduced volume over which the AH scheme deposits energy may be particularly beneficial for tritium lean implosions. This reduced volume allows the AH scheme to raise the local temperature of the fuel to much high values, which may ignite DD reactions which propagate through the fuel. Here, the potential for the AH scheme to enable tritium lean implosions is explored.

First, time deposition scans similar to those presented in Section 6.3.3 were performed to confirm that the energy deposition time did not vary with the fuel mixture. This was found to be true, and was tested using a range of deposition energies (2 kJ, 8 kJ, 20 kJ, and 60 kJ).

Performance was then compared between the interface and hotspot heating models for a range of tritium fractions, with the results shown in Figure 6.7. The performance of the implosion is found to decrease as the tritium fraction is reduced, which is expected due to the higher ignition temperature of the DD reaction. However, the relative performance between the interface and hotspot heating models remains similar across all tritium fractions. This indicates that there is no substantial difference in the effectiveness of the AH scheme between these two energy deposition models, even in the case of tritium lean implosions. The reduced deposition volume of the interface heating model does not appear to provide a significant advantage for tritium lean implosions compared to hotspot heating, at least within the range of parameters explored here.

6.4 Summary

The practical feasibility of the auxiliary heating scheme through the generation and transport of an laser-driven electron beams into the hotspot and an ICF implosion has been assessed. An analytic treatment of collisional effects indicates that a suitable electron distribution, with a positive gradient capable of driving the bump-on-tail instability, is generated during the transport. Beam divergence, however, remains a major limiting factor, critically limiting the beam density at

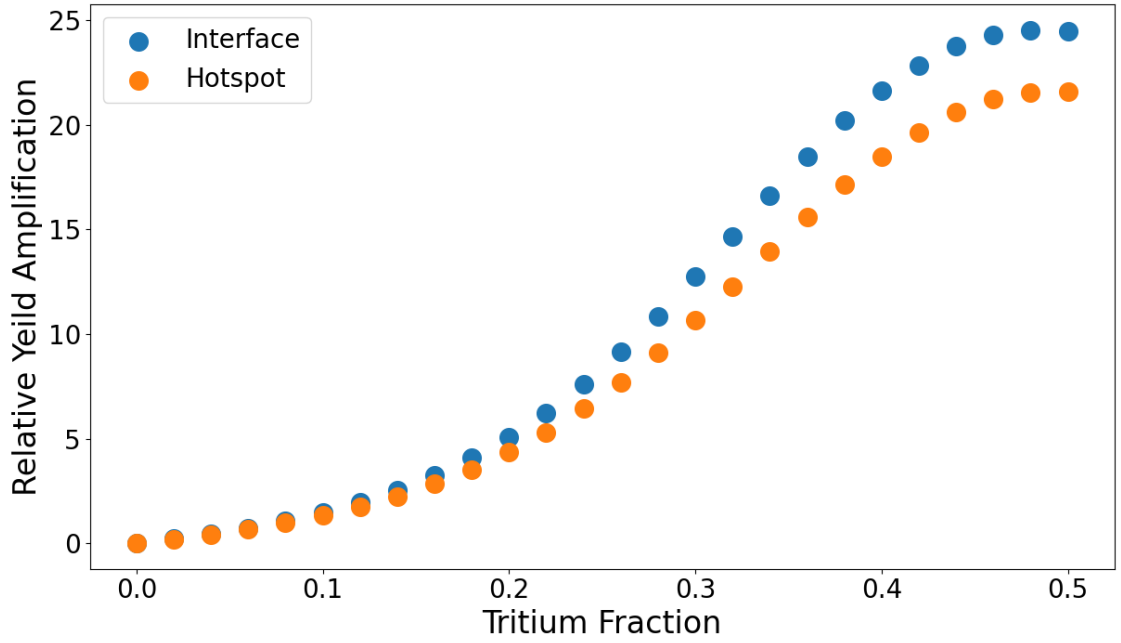


Figure 6.7: Comparison of performance between interface and hotspot heating models for a range of tritium fractions. The performance of the implosion decreases as the tritium fraction is reduced, but the relative performance between the two heating models remains similar across all tritium fractions.

the hotspot and causing a signficant fraction (> 95 %) of the beam to miss the hotspot. This issue is not easily addressed through altering laser parameters, and fundamental changes to generation and/or transport of the electron beam is necessary for the success of the scheme.

Assuming the issues around beam transport to the hotspot can be addressed, the potential impact of the heating scheme has been explored. This has been explored in the past by depositing additional thermal energy uniformly across the hotspot. Due to the short timescales of the instability, the instability would deposit energy in a thin shell on the outskirts of the hotspot. A comparison between the two energy deposition models has been made, showing very little impact on the performance between these two assumptions, with performance increasing mildly in the interface heating case. The amount of deposited energy required to boost the yield of a 0.7 MJ implosion to that of the 2 MJ was found to be 65 kJ. The use of the AH scheme therefore requires an efficiency of 6.1 % to be comparable to just increasing the long-pulse drive energy.

Finally, the application of the AH scheme to tritium lean implosions has been considered. The increased temperatures achieved in the interface heating model do not appear to provide a significant advantage for tritium lean implosions compared to hotspot heating.

7

Equation of State Measurements of Low-Density Silica Aerogel Foam

Contents

7.1 Introduction 198
7.2 Experimental Setup 199
7.2.1 Target Design 199
7.2.2 Laser System 200
7.2.3 Deviation From the Plan 201
7.3 VISAR Diagnostic 202
7.4 Determining the Foam Shock State 205
7.4.1 Shock Variables and the Rankine-Hugoniot Relations . . 205
7.4.2 Imepdance Matching 205
7.4.3 Error Propagation 208
7.5 VISAR Analysis and Results 209
7.5.1 VISAR Analysis 209
7.5.2 Shock Data 211
7.6 Vacuum Gap 213
7.6.1 Vacuum Gap Evidence 213
7.6.2 Characterising the Vacuum Gap 213
7.6.3 Vacuum Gap Simulations 214
7.6.4 Updated Hugoniot 217
7.7 Summary 220

7.1 Introduction

The interest in wetted foams for ICF applications has been discussed in Section 4.2, where wetted foam targets provide a platform for which ingiting implosions with reduced susceptibility to instability growth can be acheived. Foams are highly versitile, and have additional potential applications to mitigate key ICF implosion challanges. Both over- [218, 219] and under-dense [220, 221] foams have been proposed as ablation materials to suppress hydrodynamic instabilities by reducing laser imprint effects. Low-density foams have also been proposed as hohlraum filling material, to restrict the expansion of the hohlraum walls into the laser entrance holes [222].

Despite their promise, the behaviour of foams under shock compression remains poorly characterised, limiting the ability to accurately model or simulate their responses. A common simplification treats foams as homogeneous materials of lower density, but this approach breaks down at very low densities (e.g., 10 mg/cm^3 [223]), where pore collapse and the transition to a homogeneous plasma significantly influences the shock propagation [224, 225]. Although shock compression experiments have been performed over a range of foam densities, compositions, and shock strengths [226–229], experimental data in this low-density regime remain scarce. Additional measurements are therefore required to constrain the equation of state (EOS) of these materials, enabling more accurate hydrodynamic simulations and improved ICF target design.

This chapter presents an experiment performed at the GEKKO XII laser facility to extend this database by reporting experimental measurements of the EOS of 90 mg/cm^3 silica (SiO_2) aerogel foam [230, 231] under laser-driven shock compression up to 160 GPa. Silica aerogel was selected because silica-based foams are already relevant to several inertial confinement fusion concepts, particularly as low-density hohlraum fill materials. Although the relatively high atomic number of SiO_2 makes it less attractive as a fuel carrier within the capsule, where low-Z materials are preferred to minimise radiative losses, its EOS remains important for accurately modelling these alternative ICF applications. Shock states were determined via velocity

interferometry (VISAR) [232] measurements of shock velocities and impedance matching [233, 234] to an α-quartz standard [235]. Results are compared with previous aerogel results of a higher density, and a Quotidian Equation of State (QEOS) model [236].

In addition to providing new Hugoniot data in a previously underexplored density regime, this study introduces a diagnostic- and simulation-based methodology for identifying and correcting vacuum-gap effects at material interfaces. A key experimental observation from the VISAR diagnostic was the presence of a persistent vacuum gap between the quartz and the foam. The vacuum gaps were characterised for each shot and the impact on the measured shock velocity was evaluated through comparison with hydrodynamic simulations. Once accounted for, the measured states show strong agreement with QEOS predictions.

This chapter is organised as follows. Section 7.2 describes the target design and laser system, followed by a description of the VISAR diagnostic in Section 7.3. Section 7.4 outlines how the full shock state is experimentally determined, including the treatment and propagation of uncertainties. The experimental results are presented and discussed in Section 7.5, with comparisons to previous higher-density silica aerogel foam studies and QEOS predictions. The presence and impact of vacuum gaps are examined in Section 7.6, along with supporting hydrodynamic simulations and their influence on the inferred shock states. Finally, Section 7.7 summarises the chapter.

7.2 Experimental Setup

7.2.1 Target Design

The target design, illustrated in Figure 7.1, consists of four distinct layers, with the drive laser pulse incident from left to right. The leftmost CH layer serves as the "pusher," absorbing laser energy and converting it into a high-pressure shock. Immediately behind it, the aluminium (Al) layer absorbs energetic electrons and high-frequency radiation produced in the coronal plasma, preventing preheating of the subsequent layers. The third layer is alpha-quartz, chosen as a reference

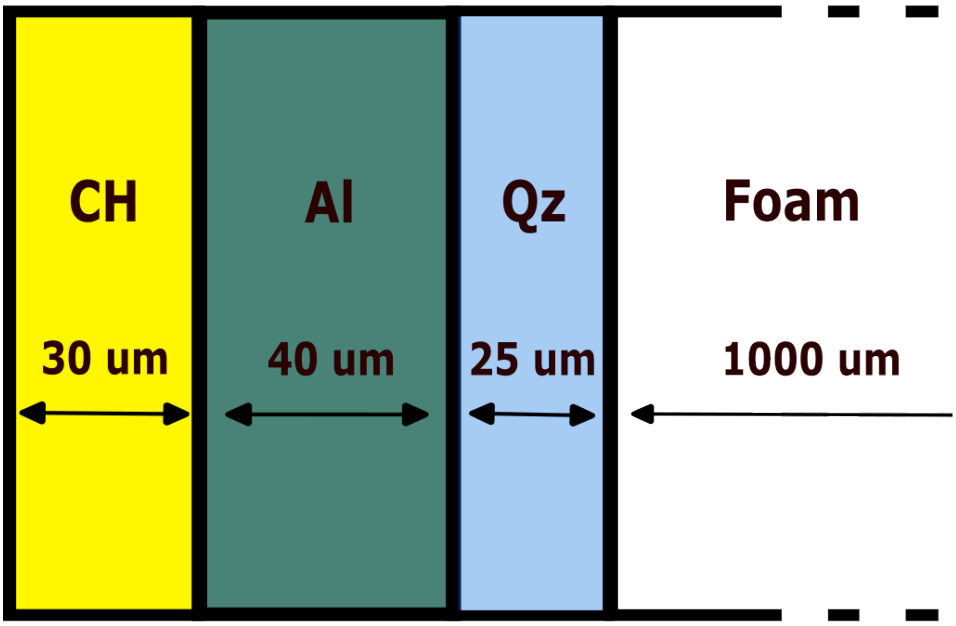


Figure 7.1: Schematic of the multi-layered target used in the experiment. The diameter of the target is 4 mm. The drive laser pulse is incident from left to right, and the probe laser pulse for both the VISAR diagnostic is incident from right to left.

material due to its well-characterised shock and release behaviour. This layer enabled the impedance matching analysis described in Section 7.4.2, which was used to determine the shock state of the foam. The final layer is the low-density aerogel foam, the material under investigation.

7.2.2 Laser System

The experiment was conducted at the Institute of Laser Engineering (ILE), University of Osaka, where the GEKKO XII laser was used to drive a strong shock through a multi-layered target [237–241]. The facility supported a total of 12 Nd:YAG laser beams: 9 frequency-tripled and 3 frequency-doubled. The frequency-tripled beams delivered a combined maximum energy of approximately 900 J, with a spot size of 600 μm. The frequency-doubled beams provided a higher per-beam energy on target, with a total maximum output of 750 J within a spot size of 1000 μm. For this experiment, all beams were operated with a fixed pulse duration of 1 ns. For the lower-intensity shots, only the 3ω beams were used, while the 2ω beams were added for the highest-intensity shots. In these cases, the 2ω and 3ω beams were

delivered simultaneously and spatially overlapped on the target, allowing drive intensities exceeding those achievable with the 3ω beams alone. This setup is shown in Figure 7.2 and enabled intensities on target up into the 10^{14} W/cm^2 order of magnitude, capable of driving strong shocks up to 160 GPa into the foam.

7.2.3 Deviation From the Plan

The original experimental plan was to field both dry foams and foams wetted with isopropanol (IPA). A side by side comparison of the shock propagation through the two materials would have provided insight into the influence of the liquid on the shock propagation, and if/how EOS models should be modified to account for the presence of the liquid. However, the targets required wetting at the Rutherford Appleton Laboratory (RAL) in Oxfordshire prior to transport to Osaka. The wetted targets were therefore flown to Osaka in hand luggage (not by me!) and were damaged during transit. The IPA had been absorbed by the glue layer sealing the target, causing the glue to fail and the IPA to leak out and evaporate. Wetted targets are therefore not included in the results presented in this chapter, although a follow-up experiment performed at the LULI2000 [242] laser facility in Paris including wetted foam targets has been conducted since.

The target design in Figure 7.1 was also not the original design of the targets. Initially an additional quartz layer was affixed to the rear of the foam. The primary purpose of this added quartz layer was to sandwich the foam layer, allowing for the foam to be wetted and sealed, without the liquid leaking. While the wetted targets did not survive transit, the dry targets, for which this added layer was not necessary, were still constructed with this additional layer. When first fielding these targets, the presence of this additional quartz layer resulted in significant ghost fringes appearing in the VISAR data. These ghost fringes compromised the ability to accurately determine the shock velocities. This layer was therefore removed (by me!) in the ILE laboratory in Osaka. As will be discussed in Section 7.6 evidence for the presence of vacuum gaps between the remaining quartz layer and foam were observed. It is inconclusive, but the relatively large vacuum gaps may have been

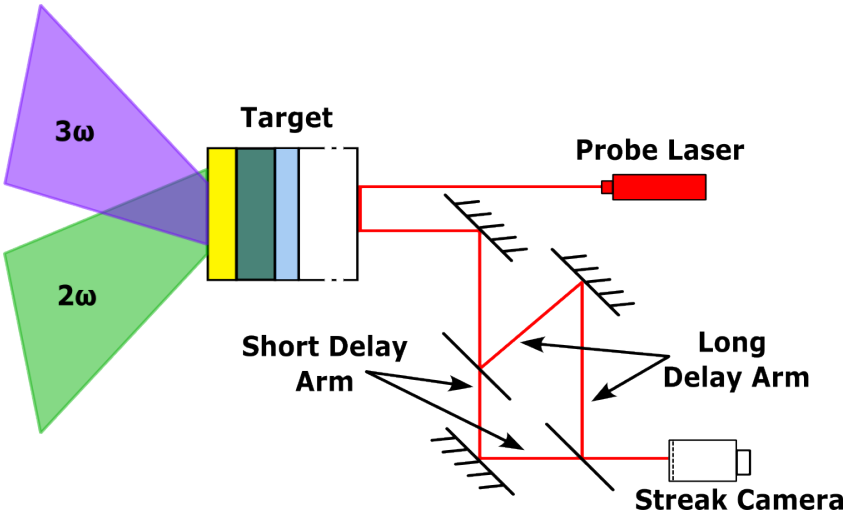


Figure 7.2: Simplified schematic of the experimental setup. The target is driven by overlapped frequency-doubled (green) and frequency-tripled (purple) laser beams. For the VISAR diagnostic, a probe laser pulse reflects off the shock front and is split down two separate paths of different optical lengths. A phase shift is introduced that is dependent on the velocity of the surface. The two beams are recombined, interfere, and the surface velocity is inferred from the measured intensity.

a result of this removal process, which may have caused the foam to be lifted off the quartz, creating a gap between the two layers.

While these deviations from the initial experiment plan were unfortunate, the experiment still provided valuable data on the shock response of low-density silica aerogel foam, and outlines a methodology for accounting for the presence of vacuum gaps in experiments similar to these.

7.3 VISAR Diagnostic

The primary diagnostic used was Velocity Interferometer System for Any Reflector (VISAR) [232]. VISAR measures the velocity of a moving reflective surface by detecting phase shifts in a probe laser beam due to changes in path length induced by the moving surface.

A probe laser pulse is directed at the reflective surface of interest. The reflected light is split into two beams, each sent through optical delay paths (etalon arms) of different lengths. The path difference between these two arms, Δx_ϵ, introduces a

corresponding time delay $\Delta t = \Delta x_\epsilon / c$, and hence a phase shift between the beams:

$$\Delta\phi_\epsilon = \frac{2\pi\Delta x_\epsilon}{\lambda}. \tag{7.1}$$

When recombined, the two beams interfere, either constructively or destructively depending on the phase difference. If the reflective surface is stationary, the phase difference remains constant.

However, if the surface moves during the time delay Δt, an additional path difference arises. A surface velocity u_{avg} over time Δt results in a displacement $\Delta x_s = u_{\text{avg}}\Delta t$ of the surface, modifying the total path difference to $\Delta x = \Delta x_\epsilon + 2n\Delta x_s$, where n is the refractive index of the material. This leads to a total phase shift of:

$$\Delta\phi_\epsilon = 2\pi\Delta x/\lambda = \frac{2\pi}{\lambda}(\Delta x_\epsilon + 2nu_{\text{avg}}\Delta t), \tag{7.2}$$

assuming Δt is much longer than the time light takes to travel Δx_s. This indicates that the phase evolves in proportion to the surface velocity.

The detector only records the interference intensity [243]:

$$I = I_0\left[1 + \cos(\Delta\phi_\epsilon)\right], \tag{7.3}$$

where I_0 is twice the intensity of an individual beam. Because the cosine function is 2π periodic, many different phase values produce the same measured intensity. Since velocity is inferred from the time evolution of $\Delta\phi$, this implies that multiple possible velocities are consistent with any single intensity measurement. This ambiguity is illustrated in Figure 7.3, showing how different phase values correspond to identical detector intensity signals. The periodicity of the intensity as a function of velocity is referred to as the velocity per fringe (VPF). This is the velocity change required to induce a 2π phase shift and hence a full fringe shift in the interference pattern, and is given by:

$$\text{VPF} = \frac{\lambda}{2n\Delta t}. \tag{7.4}$$

In a standard configuration in which the beam recombiner is placed normal to the pulse, the recombined light has a constant phase difrerence, giving only a

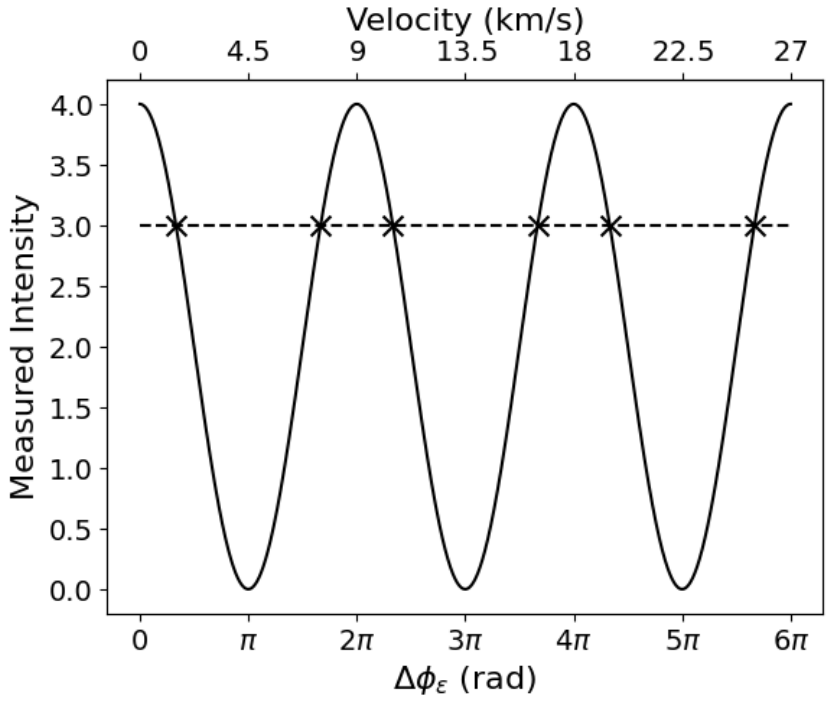


Figure 7.3: Illustration of the ambiguity in velocity inference from VISAR intensity measurements. For a scalar instensity measurement (dotted line), two possible velocity measurements exist for each 2π phase shift. This is resolved by tilting the detector (for gradient information) and using two VISARs with different VPFs (for absolute velocity information).

scalar intensity value with two possible phase solutions per 2π. To partially resolve this, the detector is tilted slightly so that the recombined beams form a spatial interference fringe across one axis of the detector. This provides knowledge of the gradient of the phase, reducing the ambiguity to one solution per 2π phase shift.

To fully resolve the ambiguity, a second VISAR with a different etalon delay (and hence a different VPF) is used. The VPFs are chosen such that they share no common multiples within the velocity range of interest. The overlap between the candidate velocity sets from each VISAR then uniquely determines the true surface velocity. In the experiment, VPFs of 5.5 and 9 km/s were used, for which the first common multiple within error bars occurs at ~ 27 km/s (27.5km/s for the 5.5 km/s system). This allows an unambiguous velocity within a ~ 27 km/s window to be identified.

7.4 Determining the Foam Shock State

7.4.1 Shock Variables and the Rankine-Hugoniot Relations

As discussed in Section 2.4, a shock is defined as a moving surface which introduces a discontinuous jump in material properties. The Rankine-Hugoniot jump conditions (Equations 2.56–2.58) can then be derived by considering the conservation of mass, momentum, and energy. These equations ensure that if any two of the shocked material variables (ρ_1, P_1, ϵ_1, u_1, U_s) are known, then all the shock variables can be calculated, defining the unique shock state of the material.

7.4.2 Imepdance Matching

Whilst two foam shock variables are required to define the full foam shock state, only one shock variable, the shock velocity, U_s, can be directly measured using the VISAR diagnostic. Impedance matching is a method to determine the other variables by leveraging known material behaviour in the reference material, in this case alpha-quartz [233, 234].

When a shock transits from the quartz standard into the foam, a release wave is launched back into the quartz, while the foam is compressed by the transmitted shock. At the interface, both the pressure and particle velocity must be continuous to prevent separation; this condition defines the impedance matching constraint. By characterising the release behaviour of the quartz, it is possible to determine the final matched state at the interface, which corresponds to the shock state in the foam.

First, the shock state of quartz at the interface is determined from the measured quartz shock velocity. Quartz has a well-characterised Hugoniot, given by a fitted expression [244]:

$$U_s = a_0 + a_1 u_p - a_2 u_p \exp\left(-a_3 u_p\right), \tag{7.5}$$

where the variables a_i are given in Table 7.1. This equation defines the quartz principal Hugoniot, the locus of final states achievable by a single shock from an unshocked state. The measured U_s^{Qz} allows the calculation of u_p^{Qz} via numerical

a_0 (km/s)	a_1	a_2	a_3 $(\text{km/s})^{-1}$
5.477	1.242	2.453	0.4336

Table 7.1: Coefficients for the quartz principal Hugoniot fit in the form $U_s = a_0 + a_1 u_p - a_2 u_p e^{-a_3 u_p}$, used to determine the shocked state of quartz from the measured shock velocity [244].

methods, and from this the full quartz shock state can be calculated using the Rankine-Hugoniot relations.

Once the quartz shock state is determined, the release path of the quartz is estimated by a Mei-Grüneisen model. The Mei-Grüneisen model is a first-order Taylor expansion of pressure around the Hugoniot:

$$P(V,E) \approx P_H(V) + \left(\frac{\partial P}{\partial E}\right)_V \cdot (E - E_H(V)). \tag{7.6}$$

where $P_H(V)$ and $E_H(V)$ are the pressure and energy along the Hugoniot. The Grüneisen parameter, $\Gamma(V)$, is then defined as

$$\frac{\Gamma(V)}{V} = \left(\frac{\partial P}{\partial E}\right)_V. \tag{7.7}$$

This approach is detailed in reference [245]. It has been shown that the Grüneisen parameter can be taken as constant along the release path, but varies based on the initial shock state, U_s, using an 'effective' Grüneisen parameter, $\Gamma_{\text{eff}}(U_s)$. An empirical fit for Γ_{eff} is given in Reference [244] to be:

$$\Gamma_{\text{eff}} = \begin{cases} -1.4545 + 0.1102\,U_s, & U_s \leq 14.69, \\ 0.579\left[1 - \exp\left(-0.129\,(U_s - 12.81)^{3/2}\right)\right], & U_s > 14.69, \end{cases} \tag{7.8}$$

with a constant uncertainty of ± 0.036.

Assuming isentropic release, the energy evolution with volume is given by:

$$\frac{\mathrm{d}E}{\mathrm{d}V} = -P(V,E) = -P_H(V) - \frac{\Gamma_{\text{eff}}}{V}(E - E_H(V)). \tag{7.9}$$

This differential equation can be solved numerically, for example via Euler's method:

$$E_{i+1} = E_i - \left[P_H(V_i) + \frac{\Gamma_{\text{eff}}}{V_i}(E_i - E_H(V_i))\right] \cdot \Delta V, \tag{7.10}$$

$$P_{i+1} = P_H(V_{i+1}) + \frac{\Gamma_{\text{eff}}}{V_{i+1}}(E_{i+1} - E_H(V_{i+1})). \tag{7.11}$$

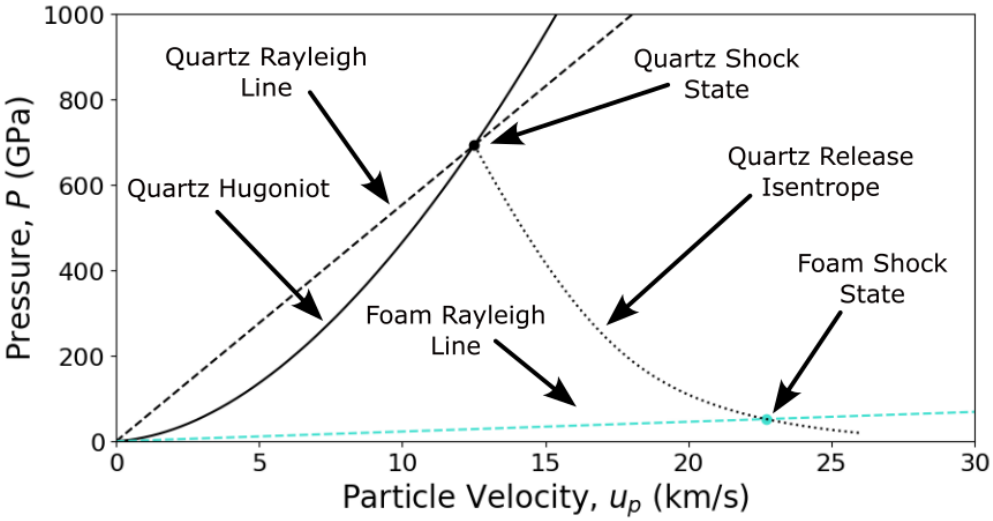


Figure 7.4: Pictorial demonstration of the impedance matching calculation.

to map a release isentrope curve describing the locus of possible release states of the quartz.

The particle velocity along this isentrope, $u_p(V)$ can then be found from the Reimann integral:

$$u_p = u_{p,1} + \int_{P_1}^{P} \frac{V}{C_s}\,\mathrm{d}P, \tag{7.12}$$

where $u_{p,1}$ is the initial shock state particle velocity and C_s is the sound speed along the isentrope given by:

$$C_s = -V^2 \frac{\partial P}{\partial V}. \tag{7.13}$$

This procedure yields a release curve in (u_p, P) space from the shocked quartz state. Because the quartz and foam remain in contact, the interface must lie along this release path and simultaneously satisfy the Rankine-Hugoniot conditions for the foam. As discussed in Section 2.4.2, the foam shock state must lie along the foam Rayleigh line (Equation 2.57), which is a linear relationship in the (u_p, P) plane with gradient $\rho_0^{\mathrm{foam}} U_s^{\mathrm{foam}}$. Thus, the shock state of the foam is found where the foam's Rayleigh line intersects the quartz release isentrope. The whole impedance matching calculation is illustrated in Figure 7.4.

7.4.3 Error Propagation

The uncertainty in the impedance matching calculation was quantified using a Monte Carlo approach [246]. For each pair of measured shock velocities in the quartz and foam, the impedance matching procedure was repeated 10^5 times. In each iteration, all input parameters with associated uncertainties were randomly sampled from normal distributions defined by their respective means and variances.

Each Monte Carlo realisation yielded a shock state for the foam, characterised by the shock velocity U_s^{foam}, particle velocity u_p^{foam}, and pressure P^{foam}. The ensemble of results from the full set of realisations formed a distribution for each of these quantities. The reported shock state corresponds to the mean of each distribution, while the associated uncertainty is given by the standard deviation.

The uncertainty enters the impedance matching calculation through four sources. The first two arose from direct experimental measurements: the shock velocities in the quartz and foam, U_s^{Qz} and U_s^{foam}, respectively. These were treated as normally distributed with means equal to the measured values and standard deviations of $\sigma_U^{\text{Qz}} = 1$ km/s and $\sigma_U^{\text{foam}} = 2$ km/s. The larger uncertainty used for the foam shock velocity stems from additional material uncertainties. For example, the uncertainty in foam density propagates into the refractive index of the foam and hence the VPF used for VISAR analysis. This effect introduces a corresponding uncertainty in the inferred foam shock velocity, which is encapsulated in the larger 2 km/s error estimate.

The remaining sources of uncertainty stem from the material models used in the calculation. The first is the quartz Hugoniot fit in the (U_s, u_p) plane (Equation 7.5). The coefficients used in this calculation were sampled from a multivariate Gaussian with mean values presented in Table 7.1 and a covariance matrix as described in Table 7.2. The second model-based uncertainty is the calculation of the effective Grüneisen parameter, (Equation 7.8). Again, this was described by a multivariate Gaussian with the mean determined by the equation, and a constant standard deviation of ± 0.036.

	σ_0	σ_1	σ_2	σ_3
σ_0	3.028×10^{-3}	-1.49×10^{-4}	-3.715×10^{-3}	-6.275×10^{-4}
σ_1	-1.49×10^{-4}	7.839×10^{-6}	1.448×10^{-4}	2.752×10^{-5}
σ_2	-3.715×10^{-3}	1.448×10^{-4}	1.729×10^{-2}	1.605×10^{-3}
σ_3	-6.275×10^{-4}	2.752×10^{-5}	1.605×10^{-3}	1.907×10^{-4}

Table 7.2: Covariance matrix for the coefficients used in the quartz Hugoniot fit (Table I). The coefficients are treated as a multivariate Gaussian distribution during Monte Carlo sampling to account for model uncertainty [244].

7.5 VISAR Analysis and Results

7.5.1 VISAR Analysis

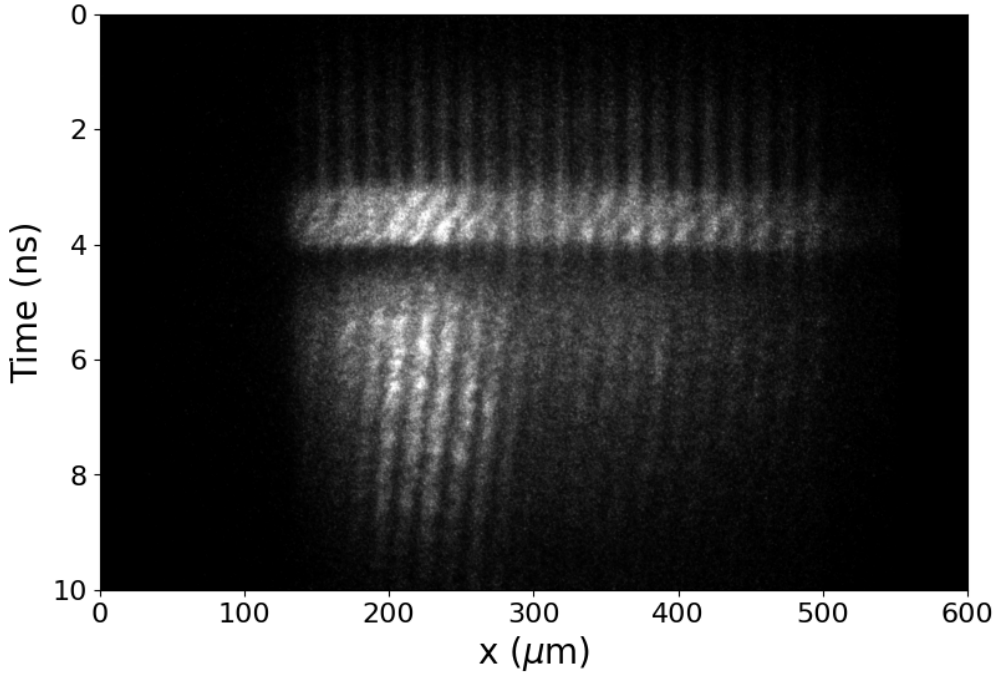


Figure 7.5: Representative raw VISAR data corresponding to shot 12 in Table 7.3, with time increasing down the vertical axis. The phase of the fringe pattern encodes the velocity of the moving surface. The image exhibits four distinct regions: (1; $0-3$ ns) vertical fringes indicating reflection from the stationary Al rear surface; (2; $3-4$ ns) diagonally shifting fringes as the shock transits from the Al through the quartz; (3; $4-4.5$ ns) a dark band associated with a vacuum gap between quartz and foam; and (4; >4.5 ns) further diagonal fringe motion as the shock propagates through the foam.

As described in Section 7.3, the VISAR diagnostic records phase shifts in a fringe pattern, which correspond to changes in the velocity of a moving surface [247].

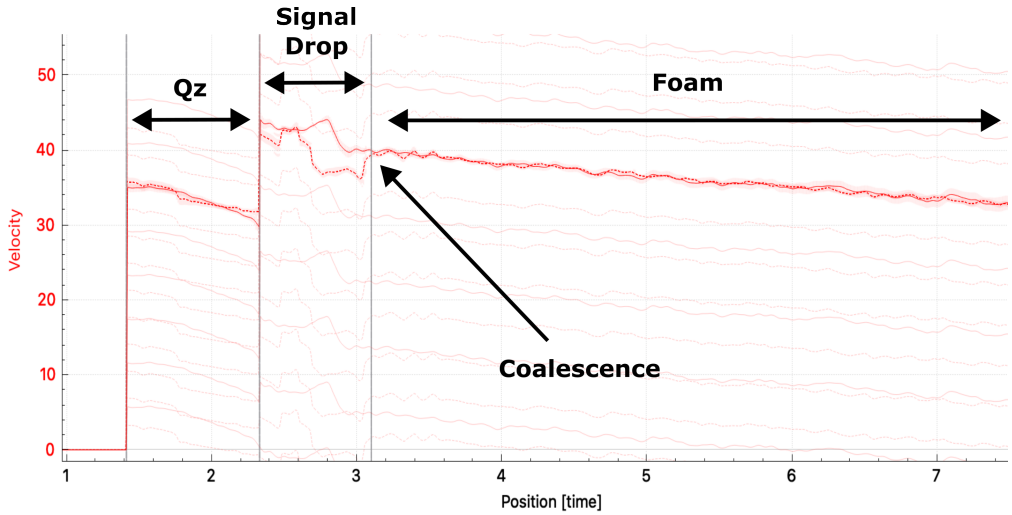


Figure 7.6: The velocity matching process for shot 12 using two VISAR systems. Candidate velocities derived from each VISAR overlap at the true surface velocity.

A representative raw VISAR image is shown in Figure 7.5, with time progressing downward and the horizontal axis representing the spatial dimension across the target. The streak camera operated with a time window of 10 ns, and the probe laser pulse duration was 7 ns.

Initially, the foam and quartz layer were transparent to the probe laser pulse, and the probe passed through the target and reflected off the rear surface of the Al layer. However, once shocked, both the quartz and foam become reflective [248]. This resulted in the probe laser pulse reflecting off the shock front. These interactions can be identified in Figure 7.5 as an example. At the top of the image ($0-3$ ns), the fringes are vertical, indicating a constant phase and therefore constant velocity; this corresponds to the probe reflecting from the rear surface of the stationary Al layer before the shock arrives. A sharp transition follows, where the signal strength increases and the fringes begin to move diagonally ($3-4$ ns)–this marks the breakout of the shock from the Al into the underlying quartz. Here, the probe reflects from the moving shock front, and the phase shift indicates a non-zero velocity. The negative slope of the fringe pattern indicates that the shock is decelerating in the quartz.

A narrow dark band appears further down ($4-4.5$ ns), corresponding to the shock reaching the interface between quartz and foam. This band indicates the presence of a vacuum gap, which is discussed further in Section 7.6. Immediately

following the dark band, the signal intensity rises again, along with a sudden fringe displacement, followed by continued diagonal motion at a shallower angle (> 4.5ns). The rise in signal intensity is consistent with the shock reforming in the foam and beoming reflective to the probe pulse again. The abrupt shift in the fringe pattern corresponds to the velocity of the shock in the foam, while the reduced fringe slope implies slower deceleration compared to the propagation through quartz.

Following this qualitative interpretation, the raw phase shift data from both VISAR systems, configured with VPF values of 9.024 km/s and 5.554 km/s, were first aligned in time by matching the abrupt phase jumps observed at shock breakout. The data was then converted into corresponding sets of possible surface velocities. The true surface velocity was determined by identifying where these sets overlap within a physically reasonable range. The refractive indices of the alpha-quartz and aerogel foam that were used for this analysis were 1.54 and 1.03 [249] respectively. This velocity matching process is illustrated in Figure 7.6.

For the impedance matching calculation, the shock velocity in the quartz was taken as the shock velocity immediately before the dark band. In this example the quartz shock velocity was measured around 2.3 ns to be 30.5 km/s. The foam shock velocity was harder to determine due to the the lack of signal immediately after the shock breaks out of the quartz layer. In this work, the foam shock velocity was measured immediately after the dark band, once the VISAR signals re-align. This will likely result in a slightly under-estimated foam shock velocity, although correcting for this is again addressed in Section 7.6.

7.5.2 Shock Data

From ten fielded targets, a total of nine valid shots were obtained. One low-energy shot was excluded because the resulting shock was too weak to produce a measurable VISAR reflection, preventing accurate determination of the shock velocity. The corresponding foam shock states are tabulated in Table 7.3 and displayed in the (u_p, P) plane in Figure 7.7. Note that the shot numbers in Table 7.3 exceed ten because additional calibration and diagnostic targets were fielded during the

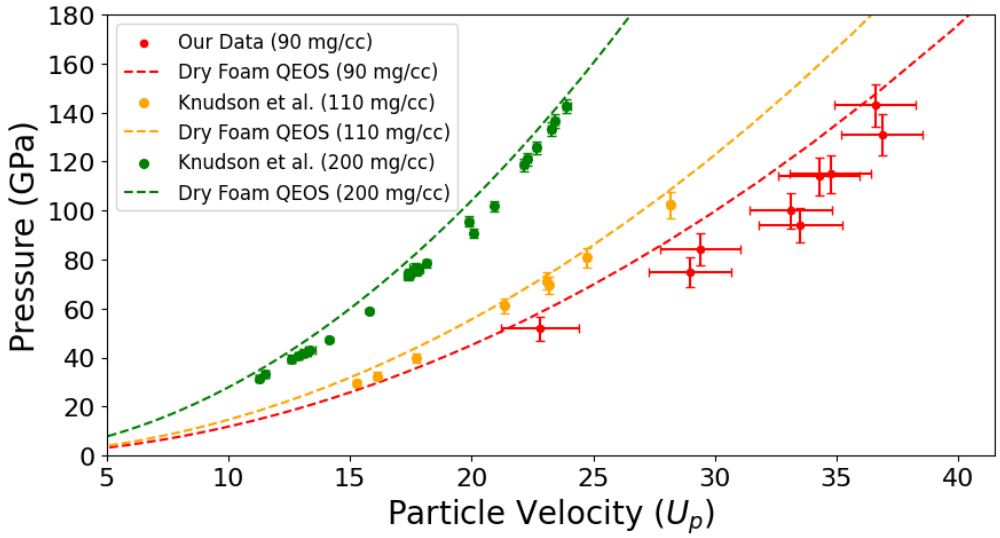


Figure 7.7: Measured shock states in 90 mg/cm^3 silica aerogel foam plotted in the (u_p, P) plane (red). Data points represent mean values from nine shots with uncertainties derived from Monte Carlo analysis. The dashed line shows the predicted Hugoniot from the QEOS model. Previous experimental data and QEOS curves for 110 mg/cm^3 (yellow) and 200 mg/cm^3 (green) aerogels from Knudson *et al.* [245, 250] are included for comparison. The experimental data for the 90 mg/cm^3 foam lies slightly below the QEOS prediction, consistent with systematic underestimation due to the presence of vacuum gaps.

Shot	U_s^{Qz} (km/s)	u_p^{Qz} (km/s)	P^{Qz} (GPa)	Γ_{eff}	U_s^{foam} (km/s)	u_p^{foam} (km/s)	P^{foam} (GPa)
4	28.75 ± 1.00	18.75 ± 0.80	1430 ± 111	0.579 ± 0.036	37.44 ± 2.00	34.28 ± 1.68	114 ± 7.70
5	25.4 ± 1.00	16.07 ± 0.80	1084 ± 96.2	0.577 ± 0.036	32.04 ± 2.00	29.40 ± 1.67	84.1 ± 6.57
7	27.81 ± 1.00	18.00 ± 0.80	1328 ± 107	0.579 ± 0.036	31.41 ± 1.99	33.51 ± 1.71	93.9 ± 7.03
8	27.78 ± 1.00	17.97 ± 0.80	1325 ± 107	0.579 ± 0.036	33.90 ± 2.00	33.12 ± 1.69	100 ± 7.22
9	20.9 ± 1.00	12.53 ± 0.78	696 ± 76.3	0.549 ± 0.036	25.49 ± 1.99	22.81 ± 1.60	51.8 ± 5.06
11	29 ± 1.00	18.95 ± 0.80	1459 ± 112	0.579 ± 0.036	37.10 ± 2.01	34.75 ± 1.69	115 ± 7.74
12	30.5 ± 1.00	20.15 ± 0.81	1631 ± 119	0.579 ± 0.036	39.95 ± 2.00	36.87 ± 1.69	131 ± 8.30
13	30.6 ± 1.00	20.23 ± 0.80	1643 ± 119	0.579 ± 0.036	43.81 ± 1.99	36.59 ± 1.66	143 ± 8.61
16	24.9 ± 1.00	15.67 ± 0.80	1036 ± 94.3	0.576 ± 0.036	28.90 ± 2.00	28.99 ± 1.69	74.8 ± 6.21

Table 7.3: Summary of experimental results from all valid shots. Mean values and uncertainties for shock velocity U_s, particle velocity u_p, and pressure P are given for both quartz and foam. The effective Grüneisen parameter Γ_{eff} is calculated for each quartz shock state.

experimental campaign. The measured shock states are compared against the Hugoniot derived from the QEOS model, where the foam is described as a very low-density, homogeneous α-quartz.

Also in Figure 7.7 is the experimental data and the associated QEOS models for both 200 mg/cm^3 (green) and 110 mg/cm^3 (yellow) density aerogel foam obtained by Knudson et al. using planar plate impact experiments at the Sandia Z machine [245]. A similar trend is observed, with the experimental measurements lying slightly below the QEOS prediction; however, the deviation is more pronounced for the 90 mg/cm^3 dataset.

7.6 Vacuum Gap

7.6.1 Vacuum Gap Evidence

Evidence of a vacuum gap between the quartz and foam layers was consistently observed in the VISAR diagnostic data for all but one shot (Shot 13). The gap manifests as a drop in signal intensity immediately following shock breakout from the rear surface of the quartz layer. This feature is attributed to the presence of a vacuum or low-density region between the quartz and foam layers, resulting in a loss of coherent reflection from a well-defined shock front and hence a temporary absence of reflectivity and signal. Similar vacuum gaps have been reported in previous experiments [251] and are known to reduce the transmitted shock velocity in the downstream foam, thereby disrupting the impedance-matching calculations.

7.6.2 Characterising the Vacuum Gap

The duration of the VISAR signal loss provides a diagnostic of the vacuum-gap dynamics. This interval corresponds to the time required for the shock to reform in the foam after shock breakout from the rear surface of the quartz.

In the presence of a vacuum gap, when the shock reaches the rear surface of the quartz, material is released into the gap and streams toward the foam. Upon arrival of the material at the foam surface, a period of ramped compression is driven in the foam, which subsequently steepens and coalesces into a well-defined shock.

During this ramped-compression phase, no continuous reflective surface exists, and the interaction is therefore non-reflective to the VISAR diagnostic. Consequently, the point at which the VISAR velocity traces re-align (after the interval of signal loss) corresponds to the formation of a fully developed shock within the foam. The time at which the foam shock velocity is measured (see Section 7.5.1) therefore marks the completion of shock reformation rather than the initial material arrival at the foam surface. The temporal interval between shock breakout from the rear quartz surface and this coalescence time was measured for each shot from the diagnostic data and is reported in Table 7.4.

Shot	U_s^{Qz} (km/s)		Δt_{coal} (ns)		
	Exp.	Sim.	Exp.	Sim.	Sim. Gap Width (μm)
4	28.75	28.87	0.55	0.57	18.5
5	25.40	25.79	0.78	0.80	26
7	27.80	27.79	1.78	1.76	70
8	27.78	28.01	1.00	1.01	36.6
9	20.90	20.88	0.41	0.42	10
11	29.00	29.36	0.70	0.70	25
12	30.50	31.06	0.80	0.80	30
13	30.60	31.13	0.00	0.00	0
16	24.90	25.48	1.80	1.80	60
				Average:	**30.7**

Table 7.4: Comparison between experimentally measured quantities and matched radiation-hydrodynamic simulations for each shot. Simulations are mapped to the experiment by matching the quartz shock velocity and the shock coalescence time.

7.6.3 Vacuum Gap Simulations

To quantify the influence of vacuum gaps on the measured foam shock velocities, one-dimensional radiation-hydrodynamic simulations were performed using the `HYADES` code [98]. The simulations modelled the full multi-layer target configuration used in the experiment, consisting of the CH pusher, Al shielding layer, α-quartz

reference, and a foam layer. The foam thickness was reduced to 105 μm relative to the experiment in order to isolate the effects of shock transmission across the quartz-foam interface.

Laser pulses were modelled with a 1 ns duration, comprised of a 0.1 ns linear rise and fall time, and a flat 0.8 ns peak. The laser intensity was varied from 40 to 200 TW/cm^2 to generate a suite of simulations. For each intensity, the quartz shock velocity at the quartz-foam interface was recorded, establishing a mapping between laser intensity and quartz shock velocity.

This mapping enables a direct correspondence between simulation and experiment based solely on the measured quartz shock velocity, without invoking any assumptions about the foam response. Because the drive calibration depends only on well-characterised materials (CH, Al, and α-quartz), it provides a robust means of matching simulated and experimental drive conditions. For each experimental shot, an effective simulation laser intensity was therefore selected such that the simulated quartz shock velocity matched the measured value to within $\pm$0.6 km/s, which are shown in Table 7.4. This procedure yielded a calibrated drive for each shot that reproduces the experimentally observed drive up to the quartz-foam interface.

Once the drive conditions were matched, additional simulations were performed that included a vacuum gap at the quartz-foam interface. A comparison between simulations with and without a vacuum gap for Shot 8 is shown in Figure 7.8. In the presence of a vacuum gap, shock transmission into the foam is delayed as material released from the rear surface of the quartz streams across the gap and accumulates within the foam. This process produces a phase of ramped compression that subsequently steepens and coalesces into a fully developed shock, as described earlier.

Figure 7.8 also illustrates the resulting difference in shock-front velocities around the quartz/vacuum - foam interface. In the simulation including a vacuum gap, the ramped-compression phase is characterised by an accelerating shock front, followed by a transition to a quasi-steady velocity once shock coalescence is complete. Importantly, the experimental VISAR data does not exhibit any acceleration phase. The experimentally measured duration therefore corresponds to the interval between

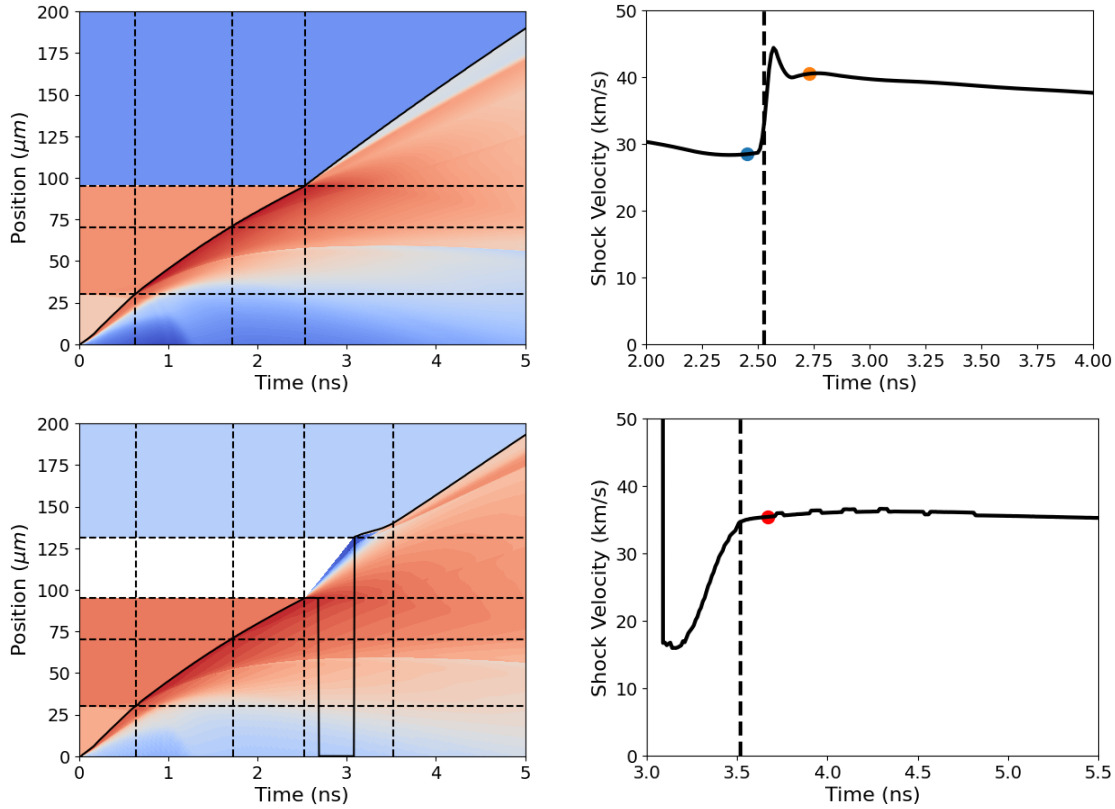


Figure 7.8: Space-time density contour plots (left) and corresponding shock velocity histories around the quartz/vacuum-foam interface (right) from 1D `Hyades` simulations (corresponding to Shot 8) of shock propagation through the multilayer target with (bottom) and without (top) a 36.6 μm vacuum gap between the quartz and foam. Horizontal dashed lines mark initial material boundaries, and the solid black line indicates the shock front. Vertical dashed lines indicate the time at which the shock reaches each interface. The presence of the vacuum gap results in a measurable reduction in the foam shock velocity.

shock breakout from the rear surface of the quartz and the formation of a fully coalesced shock in the foam, rather than the time of first material contact at the foam surface. The same definition is therefore applied in the simulations, such that the simulated and experimental durations correspond to the same physical process and can be directly compared. In the simulations, the shock coalescence time is defined as the point at which the foam shock velocity reaches a quasi-steady value, marked by the vertical dotted line in Figure 7.8.

The velocity profiles further demonstrate that the presence of a vacuum gap leads to a measurable reduction in the transmitted foam shock velocity. The simulated foam shock velocity was measured a short duration after the shock

velocity reached a quasi-steady value, as indicated by the red points in Figure 7.8. This reduction observed in the 1D simulations is reproduced in an independent 2D radiation-hydrodynamic simulation using the `FLASH` [252] code, performed by Piotr Rączka, to verify the 1D results.

To quantify the size of the vacuum gap for each shot, the gap width was treated as a free parameter in the simulations. The temporal interval between shock breakout from the rear surface of the quartz and the formation of a fully coalesced shock in the foam was extracted. An effective vacuum-gap width was then identified by matching this simulated coalescence time to the experimentally measured duration between quartz breakout and the onset of a steady foam shock.

This procedure yields simulations in which both the drive conditions (through quartz velocity matching) and the vacuum gap width (through coalescence duration matching) are consistent with experimental observations. The resulting effective gap widths inferred from this matching are listed in Table 7.4. These values represent gap widths that reproduce the experimentally observed delay to shock coalescence. The inferred gap widths are quite substantial, with a mean value of 30.7 μm, and exhibit significant shot-to-shot variation.

7.6.4 Updated Hugoniot

To correct for the reduction in the experimentally measured foam shock velocity induced by the vacuum gaps, the simulation results were used to compute shot-specific correction factors. To minimise sensitivity to the assumed equation of state (EOS), the analysis relies only on the relative change in foam shock velocity between simulations with and without a vacuum gap. This relative correction acts to quantify the reduction in transmitted shock velocity associated with shock propagation across the gap. For each shot, the foam shock velocity was extracted from the two matched simulations: one with and one without the vacuum gap. The ratio of these velocities quantified the velocity reduction caused by the gap:

$$U_{\text{foam, corrected}} = U_{\text{foam, measured}} \times \left(\frac{U_{\text{sim, no gap}}}{U_{\text{sim, with gap}}} \right). \tag{7.14}$$

The corrected foam shock velocities were then input into the impedance matching analysis to recalculate the shock pressure and particle velocity.

Although the correction factor is not strictly EOS-independent, additional simulations using QEOS models spanning a range of aerogel densities (50–200 mg/cm^3) show that it varies only weakly across this range, despite the large spread in the corresponding Hugoniots. This indicates that the correction factor is relatively insensitive to the precise EOS within the physically relevant aerogel regime, and that the derived correction factors are therefore robust, provided the assumed EOS lies within this large range.

Two additional uncertainties are included in the Monte Carlo impedance matching process to account for the uncertainties in the correction process. These come from the measurements of the foam shock velocities in the simulations with and without a vacuum gap. The measurement uncertainty of $U_{\text{sim, no gap}}$ was ± 1 km/s, and the uncertainty in $U_{\text{sim, with gap}}$ was ± 3 km/s. The larger uncertainty in the simulation including a vacuum gap encompasses the uncertainty associated with the size of the vacuum gap, and the larger uncertainty in the foam measurement itself.

The resulting corrected Hugoniot is shown in Figure 7.9 and displayed numerically in Table 7.5. After applying the gap-based corrections, the experimental data now aligns closely with the QEOS model, eliminating the earlier systematic offset. The corrected data also exhibits the same behaviour as the higher density silica aerogel foams measured by Knudson et al., sitting slightly below the QEOS model. This indicates that vacuum gaps are likely the dominant source of error in the uncorrected measurements, and that the corrected results are consistent with QEOS predictions within the present QEOS-model-assisted framework.

An alternative approach commonly used to account for small vacuum gaps is backwards propagation, in which a linear or exponential fit is applied to the measured velocity profile in the foam and extrapolated back to the α-quartz interface. This method was applied to each shot using both linear and exponential fits for comparison. For gap widths $\lesssim 20$ μm, the inferred shock velocities agree with the simulation-based corrections to within ~ 3 %, consistent with typical experimental

Shot	U_s^{Qz} (km/s)	u_p^{Qz} (km/s)	P^{Qz} (GPa)	Γ_{eff}	U_s^{foam} (km/s)	u_p^{foam} (km/s)	P^{foam} (GPa)	CF (%)
4	28.75 ± 1.00	18.75 ± 0.80	1430 ± 111	0.579 ± 0.036	41.02 ± 3.99	33.90 ± 1.7	124 ± 12.0	9.56
5	25.4 ± 1.00	16.07 ± 0.80	1084 ± 96.2	0.577 ± 0.036	35.7 ± 4.08	29.04 ± 1.68	92.2 ± 10.5	11.4
7	27.81 ± 1.00	18.00 ± 0.80	1328 ± 107	0.579 ± 0.036	38.09 ± 4.34	32.68 ± 1.72	111 ± 12.2	21.3
8	27.78 ± 1.00	17.97 ± 0.80	1325 ± 107	0.579 ± 0.036	38.62 ± 4.14	32.59 ± 1.70	112 ± 11.8	13.9
9	20.9 ± 1.00	12.53 ± 0.78	696 ± 76.3	0.549 ± 0.036	27.44 ± 3.95	22.64 ± 1.61	55.1 ± 7.87	7.65
11	29 ± 1.00	18.95 ± 0.80	1459 ± 112	0.579 ± 0.036	40.88 ± 3.99	34.33 ± 1.71	125 ± 12.1	10.2
12	30.5 ± 1.00	20.15 ± 0.81	1631 ± 119	0.579 ± 0.036	44.65 ± 4.18	36.37 ± 1.7	144 ± 12.4	11.8
13	30.6 ± 1.00	20.23 ± 0.80	1643 ± 119	0.579 ± 0.036	43.81 ± 1.99	36.59 ± 1.66	143 ± 8.61	0
16	24.9 ± 1.00	15.67 ± 0.80	1036 ± 94.3	0.576 ± 0.036	35.14 ± 4.43	28.30 ± 1.68	88.3 ± 10.9	21.6

Table 7.5: Updated version of Table 7.3 after correcting the foam shock velocity based on hydrodynamic simulation corrections.

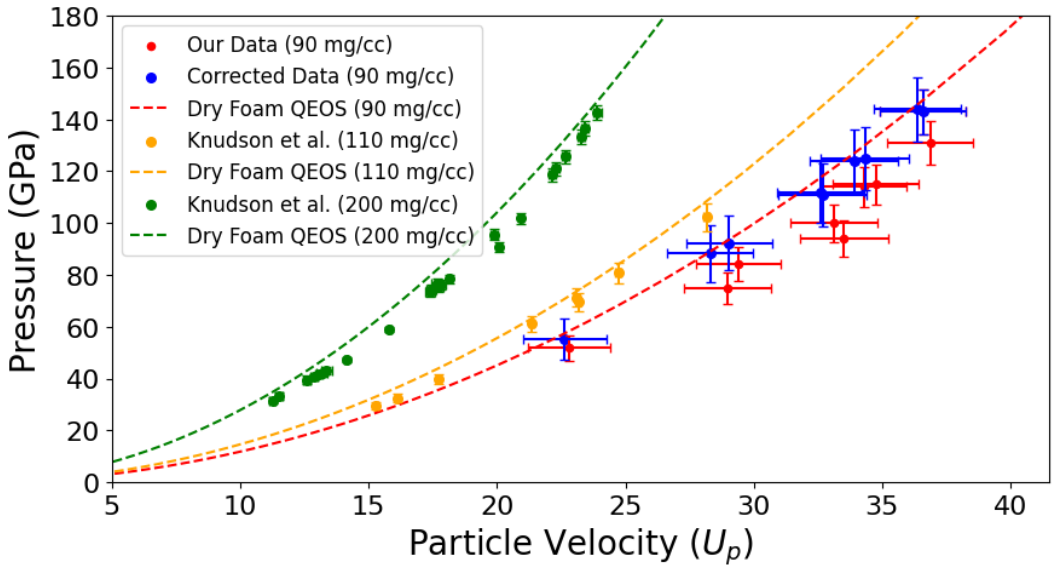


Figure 7.9: Updated Hugoniot for 90 mg/cm^3 aerogel foam after correcting experimental shock velocities based on vacuum-gap simulation results (blue). Foam shock velocities were scaled by the simulated velocity ratio (with vs. without gap) and reprocessed through impedance matching calculations. The corrected data is now closely aligned with the QEOS model (dashed line, red), supporting that the vacuum gap accounts for the original discrepancy.

uncertainties. However, the deviation between the two approaches increases with gap size, reflecting the growing influence of ramped compression and layed shock formation, which are not captured by the backwards propagation method.

7.7 Summary

The principal Hugoniot of 90 mg/cm^3 silica aerogel foam has been measured under laser-driven shock compression to pressures up to 160 GPa. These experiments, conducted at the GEKKO XII laser facility, used VISAR diagnostics to characterise the shock state via impedance matching against an α-quartz reference. This work extends previous studies by Knudson et al. to a lower-density silica regime, indicating that, at these densities, the aerogel behaves as a homogeneous material well described by standard QEOS modelling, providing new experimental constraints in a previously underexplored region of parameter space.

Initial results showed a systematic underestimation of foam shock velocities relative to the QEOS model. Analysis of the VISAR data revealed the consistent presence of vacuum gaps at the quartz-foam interface, which were characterised and quantified across shots. One-dimensional radiation-hydrodynamic simulations confirmed that such gaps reduced the transmitted shock velocities and shifted the apparent Hugoniot below its true value. After correcting for the vacuum gap using simulation-derived velocity scalings, the experimental Hugoniot closely matches QEOS predictions. This supports the use of QEOS for modelling low-density aerogel foams in high-pressure regimes associated with these experimental conditions.

These results contribute new experimental data to the high-energy-density EOS database for low-density foams - materials critical to inertial confinement fusion and other high-energy-density applications. More broadly, this work highlights the need for careful interpretation of VISAR data, as even small vacuum gaps at material interfaces can lead to underestimation of shock pressure.

8

Summary and Further Work

Contents

8.1 Summary 221
8.2 Further Work 223

8.1 Summary

This thesis has investigated new approaches to both the design and modelling of inertial confinement fusion (ICF) implosions, with the aim of improving performance, robustness, and predictive capability. The work has focused on three complementary areas: optimisation of implosion designs, exploration of an auxiliary heating mechanism, and experimental characterisation of materials relevant to ICF targets.

Modern ICF implosions are governed by high-dimensional parameter spaces, where each design evaluation requires computationally expensive radiation-hydrodynamics simulations. Consequently, there is a need for efficient optimisation frameworks capable of systematically exploring these design spaces. While several optimisation tools exist, many are either not tailored to ICF applications or remain inaccessible outside individual research laboratories. This thesis addressed this need through the development of complementary Bayesian optimisation and particle

swarm optimisation packages, specifically designed for ICF optimisation problems. The methods were benchmarked on wetted-foam implosion designs, demonstrating improved sampling efficiency and the ability to identify higher-performing implosion designs. By making these packages freely available to the academic community, this work provides practical optimisation tools that can be readily applied to future ICF design studies.

Novel auxiliary heating concepts have recently attracted interest as a possible route to improving hotspot performance and increasing the robustness of ignition. However, the practical applicability of such schemes under realistic ICF conditions remained uncertain. This thesis investigated the underlying beam-plasma coupling physics, identified the conditions required for efficient energy deposition, and evaluated the generation, transport, and coupling of relativistic electron beams in realistic implosion environments. Although efficient collisionless energy coupling was shown to be possible, electron beam transport, and in particular beam divergence, was identified as the dominant limitation preventing effective hotspot heating. These results therefore clarify the principal challenge that must be overcome if auxiliary heating is to become a viable approach for future ICF implosions.

Low-density foams are becoming increasingly important in advanced ICF target concepts, yet their equations of state remain poorly constrained under ICF-relevant conditions, limiting the predictive capability of radiation-hydrodynamic simulations. Addressing this limitation requires new experimental data with which equation-of-state models and tables can be validated and improved. This thesis contributes new principal Hugoniot measurements for 90 mg/cm^3 silica aerogel foams, extending the available experimental database for these materials. In addition, a methodology was developed to account for the presence of vacuum gaps in principal Hugoniot experiments, improving the interpretation of experimental measurements and providing a framework for future studies of similar materials.

Taken together, this thesis has addresses three complementary challenges limiting the development of laser-driven inertial confinement fusion: the efficient optimisation of implosion designs, the assessment of novel auxiliary heating concepts, and

the uncertainty associated with material modelling. The work provides practical optimisation tools for the wider ICF community, establishes the physical limitations governing a proposed auxiliary heating mechanism, and extends the experimental database constraining the equation-of-state of low-density foam materials.

8.2 Further Work

The work presented in this thesis opens several avenues for further research, spanning optimisation methodologies, auxiliary heating physics, and material characterisation.

For the optimisation frameworks, a natural extension is the incorporation of multi-fidelity modelling, where inexpensive lower-fidelity simulations are combined with high-fidelity radiation-hydrodynamic calculations to further improve sampling efficiency. Additionally, applying these optimisation approaches directly to experimental campaigns represents a promising direction, enabling adaptive design of experiments in real time. Outside of additional applications, further improvements to the optimisation algorithms themselves can be explored, in particular through the support of asynchronous evaluations, reducing the amount of downtime between function evaluations and enabling more efficient use of computational resources.

For the auxiliary heating scheme, the primary challenge identified in this work is the generation and transport of suitable electron beams. Future work should therefore focus on more detailed modelling of laser-plasma interactions and electron beam transport, including fully kinetic simulations that capture self-generated electromagnetic fields and collective effects. Experimental investigation of electron beam divergence and transport in ICF-relevant conditions would also be crucial for validating these models. In parallel, exploration of alternative beam generation mechanisms or collimation strategies, such as structured targets or externally applied fields, may provide pathways to overcoming the limitations identified in this thesis.

For the characterisation of low-density foams, further experimental studies across a wider range of foam densities, compositions, and microstructures would significantly improve the available EOS data. In particular, the development of deterministic or structured foams using advanced manufacturing techniques such

as two-photon polymerisation (TPP) [253] offers an opportunity to systematically investigate the role of microstructure on shock propagation and energy transport. Integrating experimental measurements directly into EOS model development and validation frameworks represents another important direction, enabling more accurate and predictive simulation capabilities.

More broadly, an important area for future work lies in the integration of these three strands. Combining improved optimisation techniques with novel implosion designs and experimentally validated material properties would enable more reliable and predictive design of ICF implosions. Such an integrated approach will be essential for advancing toward robust ignition and high-gain fusion performance.

Appendices

A

Theory Section Derivations

A.1 Maxwell-Boltzmann Distribution

On a microscopic level, or using a 'kinetic description' of the plasma, the system is described by a distribution function, $f(\mathbf{x}, \mathbf{p}, t)$, which describes the number density of particles at a given position $\mathbf{x}$, momentum $\mathbf{p}$, and time t. In thermodynamic equilibrium, the distribution function is time independent, so the rate of particles entering a state is equal to the rate of particles leaving that state. This idea of equivalent rates into and out of a state is referred to as detailed balance and can be used to derive the momentum distribution of particles in thermodynamic equilibrium.

Consider a simple collision between two particles, 1 and 2, with initial energies E_1 and E_2, which collide elastically (meaning energy is conserved) to energies E_1' and E_2'. Detailed balance requires that the probability of finding the two particles in the initial state is equal to the probability of finding the two particles in the final state, i.e.

$$f(\mathbf{v_1})f(\mathbf{v_2}) = f(\mathbf{v_1'})f(\mathbf{v_2'}). \tag{A.1}$$

The elastic nature of the collision also requires that momentum and energy are conserved, i.e.

$$m\mathbf{v_1} + m\mathbf{v_2} = m\mathbf{v_1'} + m\mathbf{v_2'}, \tag{A.2}$$

$$\frac{1}{2}m_1v_1^2 + \frac{1}{2}m_2v_2^2 = \frac{1}{2}m_1v_1'^2 + \frac{1}{2}m_2v_2'^2. \tag{A.3}$$

Taking the logarithm of the detailed balance equation (Equation A.1) gives

$$\ln f(\mathbf{v_1}) + \ln f(\mathbf{v_2}) = \ln f(\mathbf{v_1'}) + \ln f(\mathbf{v_2'}). \tag{A.4}$$

Letting $g(\mathbf{v}) = \ln f(\mathbf{v})$, and $G(\mathbf{v_1}, \mathbf{v_2}) = g(\mathbf{v_1}) + g(\mathbf{v_2})$, then $G(\mathbf{v_1}, \mathbf{v_2})$ must be constant on the set of all states with the same total momentum and energy:

$$g(\mathbf{v_1}) + g(\mathbf{v_2}) = G(\mathbf{P}, E) = G(\mathbf{v_1} + \mathbf{v_2}, v_1^2 + v_2^2). \tag{A.5}$$

However, the left-hand side is separable in $\mathbf{v}_1$ and $\mathbf{v}_2$, while the right-hand side depends only on the combinations $\mathbf{v}_1 + \mathbf{v}_2$ and $v_1^2 + v_2^2$. The only functions $g(\mathbf{v})$ for which $g(\mathbf{v}_1) + g(\mathbf{v}_2)$ depends only on these combinations are quadratic functions of $\mathbf{v}$, of the form

$$g(v) = \alpha - \beta v^2 + \boldsymbol{\gamma} \cdot \mathbf{v}, \tag{A.6}$$

where α, β, and $\boldsymbol{\gamma}$ are constants. A negative sign for β is chosen to ensure that $f(\mathbf{v}) \to 0$ as $v \to \infty$, otherwise the distribution would not be normalisable.

In equilibrium, the system is isotropic and has no net bulk flow. The absence of a preferred direction implies that the distribution cannot depend linearly on $\mathbf{v}$, so the coefficient of the $\mathbf{v}$ term must vanish. $\boldsymbol{\gamma}$ is therefore set to zero, giving

$$g(\mathbf{v}) = \alpha - \beta v^2. \tag{A.7}$$

Exponentiating both sides gives the form of the Maxwell-Boltzmann distribution:

$$f(\mathbf{v}) = \alpha e^{-\beta v^2} = \alpha e^{-\beta \frac{p^2}{m^2}}. \tag{A.8}$$

The constant β can be determined by considering the average kinetic energy of the particles in the system. The average kinetic energy of a particle in three dimensions is given by

$$\langle E_k \rangle = \frac{\int_{\mathbb{R}^3} \frac{p^2}{2m} f(p) d^3p}{\int_{\mathbb{R}^3} f(p) d^3p} = \frac{\int_{\mathbb{R}^3} \frac{p^2}{2m} \alpha e^{-\beta \frac{p^2}{m^2}} d^3p}{\int_{\mathbb{R}^3} \alpha e^{-\beta \frac{p^2}{m^2}} d^3p} \tag{A.9}$$

Using the standard equations for Gaussian integrals:

$$\int_{\mathbb{R}^3} p^2 e^{-ap^2} d^3p = \frac{3}{2}\left(\frac{\pi}{a}\right)^{\frac{3}{2}} \frac{1}{a}, \tag{A.10}$$

$$\int_{\mathbb{R}^3} e^{-ap^2} d^3p = \left(\frac{\pi}{a}\right)^{\frac{3}{2}}, \tag{A.11}$$

the average kinetic energy can be evaluated to be

$$\langle E_k \rangle = \frac{3m}{4\beta}. \tag{A.12}$$

Standard thermodynamics tells us that the average kinetic energy of a particle in three dimensions is related to the temperature of the system by

$$\langle E_k \rangle = \frac{3}{2} k_B T, \tag{A.13}$$

giving the value of β to be

$$\beta = \frac{m}{2k_B T}. \tag{A.14}$$

The form of the full Maxwell-Boltzmann distribution is therefore

$$f(\mathbf{p}) = \alpha e^{-\frac{p^2}{2mk_B T}}, \tag{A.15}$$

where α is a normalisation constant that can be determined by ensuring that the integral of the distribution function over all momenta is equal to the total number of particles in the system.

$$\int_{\mathbb{R}^3} f(\mathbf{p}) d^3p = N \tag{A.16}$$

$$\int_{\mathbb{R}^3} \alpha e^{-\frac{p^2}{2mk_B T}} d^3p = N \tag{A.17}$$

$$\alpha \left(2\pi m k_B T\right)^{\frac{3}{2}} = N \tag{A.18}$$

$$\alpha = \frac{N}{\left(2\pi m k_B T\right)^{\frac{3}{2}}}. \tag{A.19}$$

A.2 Linearised Governing Equations (Plasma Oscillations)

The equations which govern the dynamics of electrons responding to perturbations are

$$m_e n_e \left[\frac{\partial \mathbf{u}_e}{\partial t} + (\mathbf{u}_e \cdot \nabla)\mathbf{u}_e\right] = -en_e\left(\mathbf{E} + \mathbf{u}_e \times \mathbf{B}\right), \tag{A.20}$$

$$\frac{\partial n_e}{\partial t} + \nabla \cdot (n_e \mathbf{u}_e) = 0, \tag{A.21}$$

$$-\nabla \cdot \mathbf{E} = -\frac{e}{\epsilon_0}(Zn_i - n_e). \tag{A.22}$$

First linearising the equation of motion:

$$m_e(n_0 + n_1)\left[(-i\omega)\mathbf{u}_1 + (\mathbf{u}_1 \cdot i\mathbf{k})\mathbf{u}_1\right] = -e(n_0 + n_1)\left[\mathbf{E}_1 + \mathbf{u}_1 \times (\mathbf{B}_0 + \mathbf{B}_1)\right], \tag{A.23}$$

$$m_e(n_0 + n_1)\left[(-i\omega)\mathbf{u}_1\right] = -en_0\mathbf{E}_1 - en_0\mathbf{u}_1 \times \mathbf{B}_0, \tag{A.24}$$

$$-i\omega m_e n_0 \mathbf{u}_1 = -en_0(\mathbf{E}_1 + \mathbf{u}_1 \times \mathbf{B}_0), \tag{A.25}$$

$$i\omega m_e \mathbf{u}_1 = e(\mathbf{E}_1 + \mathbf{u}_1 \times \mathbf{B}_0) \tag{A.26}$$

Next the continuity equation:

$$(-i\omega)n_1 + i\mathbf{k} \cdot (n_0 \mathbf{u}_1) = 0, \tag{A.27}$$

$$-i\omega n_1 + in_0 \mathbf{k} \cdot \mathbf{u}_1 = 0, \tag{A.28}$$

$$\omega n_1 = n_0 \mathbf{k} \cdot \mathbf{u}_1. \tag{A.29}$$

Finally Poisson's equation:

$$-i\mathbf{k} \cdot \mathbf{E_1} = -\frac{e}{\epsilon_0}(Zn_i - n_e), \tag{A.30}$$

$$-i\mathbf{k} \cdot \mathbf{E}_1 = -\frac{e}{\epsilon_0}(-n_1), \tag{A.31}$$

$$i\epsilon_0 \mathbf{k} \cdot \mathbf{E}_1 = -en_1. \tag{A.32}$$

A.3 Coulomb Scattering Angle

Consider two charged particles with charges q_1 and q_2, and masses m_1 and m_2, interacting via the Coulomb potential. In the centre-of-mass frame, the motion reduces to that of a single particle with reduced mass $\mu = \frac{m_1 m_2}{m_1 + m_2}$ and relative velocity v.

In the limit of small deflection (large impact parameter or high velocity), the trajectory may be approximated as a straight line. The separation between the two particles at time t is then

$$r(t) = \sqrt{b^2 + v^2 t^2}. \tag{A.33}$$

The transverse component of the Coulomb force is

$$F_\perp(t) = F(r)\frac{b}{r} = \frac{Kb}{r^3} = \frac{Kb}{(b^2 + v^2 t^2)^{3/2}}, \tag{A.34}$$

where $b/r = \sin(\theta)$ is the projection of the central force perpendicular to the initial velocity.

The total change in transverse momentum of the moving particle is given by integrating the transverse force over time:

$$\Delta p_\perp = \int_{-\infty}^{\infty} F_\perp(t)dt = \int_{-\infty}^{\infty} \frac{Kb}{(b^2 + v^2 t^2)^{3/2}}dt. \tag{A.35}$$

Making the substitution $u = vt$ ($dt = du/v$) gives

$$\Delta p_\perp = \int_{-\infty}^{\infty} \frac{Kb}{(b^2 + u^2)^{3/2}}\frac{du}{v} = \frac{Kb}{v}\int_{-\infty}^{\infty} \frac{du}{(b^2 + u^2)^{3/2}} = \frac{2K}{vb}. \tag{A.36}$$

For small angles, the scattering angle is the transverse change in momentum divided by the initial momentum, giving

$$\chi(b) \approx \frac{\Delta p_\perp}{p} = \frac{2K}{\mu v^2 b}, \tag{A.37}$$

A.4 Stimulated Emission

It was known that matter could absorb photons, as blackbody emission excited matter, and also emit photons, as atoms returned to lower energy states and emitted radiation. Consider a simple system which consists of an ensamble of atoms with two energy levels, $E_1 < E_2$, and let $\hbar\omega = E_2 - E_1$. Let the absorption rate from state 1 to state 2 be given by $W_{1\to 2} = B_{12}\rho(\omega)N_1$, where $\rho(\omega)$ is the number density of photons of frequency ω (i.e. the excitation rate is proportional to the number of photons present and the number of atoms in the de-excited state). Similarly, let

the emission rate from state 2 to 1 be $W_{2\to1} = A_{21}N_2$, where there is no $\rho(\omega)$ term as de-excitation occurs even without the presence of photons. Thermal equilibrium requires no net flow between the two energy levels, requiring

$$B_{12}\rho N_1 = A_{21}N_2. \tag{A.38}$$

Since the atoms are in thermal equilibrium, the ratio of excited atoms to de-excited atoms is

$$\frac{N_2}{N_1} = e^{-\hbar\omega/k_BT}. \tag{A.39}$$

This then gives

$$\rho = \frac{A_{21}}{B_{12}}\frac{N_2}{N_1} = \frac{A_{21}}{B_{12}}\frac{1}{e^{\hbar\omega/k_BT}} \tag{A.40}$$

where $\rho \propto e^{-\hbar\omega/k_BT}$, which is the Wien distribution, which matches the Plankian distribution only at large ω.

To arrive at the full Plankian distribution, a -1 in the denominator is required. To achieve this, the emission process must also depend on ρ, so that ρ appears on both sides of the balance equation. An additional emission process was therefore proposed that was dependent on the ρ, given by $W^{\text{stim}}_{2\to1} = B_{21}\rho N_2$. The balance equation now looks like

$$B_{12}\rho N_1 = A_{21}N_2 + B_{21}\rho N_2, \tag{A.41}$$

which can be rearranged to give

$$\rho = \frac{A_{21}}{B_{21}}\frac{1}{\frac{B_{12}}{B_{21}}e^{\hbar\omega/k_BT} - 1}. \tag{A.42}$$

and the Plankian distribution is obtained.

A.5 ICF Lawson Criterion

The n-T-τ triple product condition (Equation 2.70) can be further manipulated to form an ICF specific Lawson criterion. This is done by considering the confinement time, τ_E of the implosion. The confinement time is the time it takes for the

implosions to disassemble, which is approximately the time it takes for a ion travelling at the thermal velocity to cross the compressed region, i.e.

$$\tau_E \approx \frac{R}{v_{th}} \tag{A.43}$$

Substituting this expression for τ_E into the n-T-τ triple product condition, dividing by T, and using that $n = \rho/m$ gives

$$n\tau_E = \frac{\rho}{m}\frac{R}{v_{th}} \geq \frac{12k_BT}{E_\alpha\langle\sigma v\rangle}. \tag{A.44}$$

Rearranging gives the ICF specific Lawson criterion:

$$\rho R \geq m v_{th}\frac{12k_BT}{E_\alpha\langle\sigma v\rangle}. \tag{A.45}$$

Typical values for the parameters on the right hand side of this equation are:

- $k_BT \approx 5$ keV
- $E_\alpha = 3.5$ MeV
- $m = 2.5$ u
- $\langle\sigma v\rangle \approx 1.3 \times 10^{-23}$ m^3/s [66]
- $v_{th} = 6.2 \times 10^5$ m/s

While these values can vary depending on the exact conditions of the implosion, they give an estimate of $\rho R = 0.34$ for ignition. The constraint often varies, but are usually between around 0.1-1 g/cm^2.

B

Bayesian Optimisation Derivations

B.1 Correlated Multivariate Gaussian Distribution PDF

To derive the PDF of a correlated multivariate Gaussian distribution, we start with a standard normal distribution, which is a special case of the multivariate Gaussian distribution where the mean vector $\mu = 0$ and the covariance matrix Σ is the identity matrix I. The PDF of a standard normal distribution (Equation 3.8) over n variables is given by:

$$p_x(\mathbf{x}) = \frac{1}{(2\pi)^{n/2}} \exp\left(-\frac{1}{2}\mathbf{x}^T\mathbf{x}\right). \tag{B.1}$$

To produce a Guassian distribution with desired mean μ and covariance Σ, we can apply a linear transformation to the standard normal variable $\mathbf{x}$. We define this desired variable $\mathbf{z}$ as follows:

$$\mathbf{z} = \mu + L\mathbf{x}, \tag{B.2}$$

where L is the matrix such that $LL^T = \Sigma$. This is know as the Cholesky decomposition of the covariance matrix Σ, and discussed in the following Appendix section.

Before deriving the PDF of $\mathbf{z}$, it is first important to confirm that $\mathbf{z}$ indeed has the desired mean and covariance structure. The mean of $\mathbf{z}$ is given by:

$$\mathbb{E}[\mathbf{z}] = \mathbb{E}[\mu + L\mathbf{x}] = \mu + L\mathbb{E}[\mathbf{x}] = \mu, \tag{B.3}$$

since $\mathbb{E}[\mathbf{x}] = 0$ for a standard normal distribution. The covariance of $\mathbf{z}$ is given by:

$$\text{Cov}(\mathbf{z}) = \mathbb{E}[(\mathbf{z} - \mathbb{E}[\mathbf{z}])(\mathbf{z} - \mathbb{E}[\mathbf{z}])^T] = \mathbb{E}[(L\mathbf{x})(L\mathbf{x})^T] = L\mathbb{E}[\mathbf{x}\mathbf{x}^T]L^T = LIL^T = \Sigma, \tag{B.4}$$

where the fact that $\mathbb{E}[\mathbf{x}\mathbf{x}^T] = I$ for a standard normal distribution has been used. Thus, $\mathbf{z}$ has the desired mean μ and covariance Σ.

The PDF of $\mathbf{z}$ can now be derived by performing the change of variables from $\mathbf{x}$ to $\mathbf{z}$, $\mathbf{x} = L^{-1}(\mathbf{z} - \mu)$. The PDF is then given by:

$$p_z(\mathbf{z}) = p_x(L^{-1}(\mathbf{z} - \mu)) \left|\det\left(J_{L^{-1}}\right)\right|, \tag{B.5}$$

where $J_{L^{-1}}$ is the Jacobian of the transformation from $\mathbf{x}$ to $\mathbf{z}$. The Jacobian is given by:

$$\left|\det\left(J_{L^{-1}}\right)\right| = \left|\det\left(L^{-1}\right)\right| = 1/\left|\det\left(L\right)\right|. \tag{B.6}$$

The PDF is therefore given by:

$$p_z(\mathbf{z}) = \frac{1}{(2\pi)^{n/2}} \exp\left(-\frac{1}{2}(L^{-1}(\mathbf{z} - \mu))^T(L^{-1}(\mathbf{z} - \mu))\right) \cdot \frac{1}{\left|\det\left(L\right)\right|}. \tag{B.7}$$

The exponent can be expanded as follows:

$$(L^{-1}(\mathbf{z}-\mu))^T(L^{-1}(\mathbf{z}-\mu)) = (\mathbf{z}-\mu)\left(L^{-T}L^{-1}\right)(\mathbf{z}-\mu) = \mathbf{z}-\mu\Sigma^{-1}\left(\mathbf{z}-\mu\right), \tag{B.8}$$

using the fact that $L^{-T}L^{-1} = (LL^T)^{-1} = \Sigma^{-1}$.

The final piece of the puzzel is the determinant of L. Since $LL^T = \Sigma$, it follows that $\det(L)^2 = \det(\Sigma)$, so $\left|\det\left(L\right)\right| = |\Sigma|^{1/2}$. Thus, the PDF of $\mathbf{z}$ is given by:

$$p_z(\mathbf{z}) = \frac{1}{(2\pi)^{n/2}|\Sigma|^{1/2}} \exp\left(-\frac{1}{2}\left(\mathbf{z} - \mu\right)^T \Sigma^{-1} \left(\mathbf{z} - \mu\right)\right), \tag{B.9}$$

which is the desired PDF of a correlated multivariate Gaussian distribution with mean μ and covariance Σ.

B.2 Cholesky Decomposition

A covariance matrix, Σ, can be written as the product of a lower triangular matrix, L, and its transpose, L^T, such that $\Sigma = LL^T$. This is known as the Cholesky decomposition of the covariance matrix. The Cholesky decomposition is a useful tool for generating correlated multivariate Gaussian random variables from uncorrelated standard normal random variables, as shown in the previous section.

The process of Cholesky decomposition requires that the covariance matrix Σ be positive definite. This condition is satisfied if the matrix is symmetric and all its eigenvalues are positive. For a covariance matrix, the first of these conditions is always satisfied since covariance matrices are symmetric by definition. The second condition, positive definiteness, is not guaranteed, however positive semi-definiteness is, as eigenvalues of the system may be zero. To make the matrix positive definite, a small jitter term can be added to the diagonal of the covariance matrix, $\Sigma' = \Sigma + \epsilon I$, where ϵ is a small positive constant and I is the identity matrix. This ensures that all eigenvalues of Σ' are positive, making it positive definite.

To then perform the Cholesky decomposition it is best to go column by column (or equivalently row by row). The matrix multiplication LL^T in 3 dimensions is given by:

$$LL^T = \begin{pmatrix} L_{11} & 0 & 0 \\ L_{21} & L_{22} & 0 \\ L_{31} & L_{32} & L_{33} \end{pmatrix} \begin{pmatrix} L_{11} & L_{21} & L_{31} \\ 0 & L_{22} & L_{32} \\ 0 & 0 & L_{33} \end{pmatrix} \tag{B.10}$$

$$= \begin{pmatrix} L_{11}^2 & L_{11}L_{21} & L_{11}L_{31} \\ L_{11}L_{21} & L_{21}^2 + L_{22}^2 & L_{21}L_{31} + L_{22}L_{32} \\ L_{11}L_{31} & L_{21}L_{31} + L_{22}L_{32} & L_{31}^2 + L_{32}^2 + L_{33}^2 \end{pmatrix} = \begin{pmatrix} \Sigma_{11} & \Sigma_{21} & \Sigma_{31} \\ \Sigma_{21} & \Sigma_{22} & \Sigma_{32} \\ \Sigma_{31} & \Sigma_{32} & \Sigma_{33} \end{pmatrix} = \Sigma. \tag{B.11}$$

From the first column, $L_{11}^2 = \Sigma_{11}$ can be solved to give L_{11}. This can then be used to solve the following two equations in the first column to give L_{21} and L_{31}. This then cascades in a similar way down the subsequent columns. The general formula for the elements of L is given by:

$$L_{jj} = \begin{cases} \sqrt{\Sigma_{jj} - \sum_{k=1}^{j-1} L_{jk}^2} & \text{if } i = j \\ \frac{1}{L_{jj}} \left(\Sigma_{ij} - \sum_{k=1}^{j-1} L_{ik} L_{jk} \right) & \text{if } i > j \\ 0 & \text{if } i < j \end{cases} \tag{B.12}$$

C

Low-CR Optimisaiton Details

C.1 Low-CR Optimisation Hyper-Parameters

C.1.1 Bayesian Hyper-Parameters

All enhancements discussed in Section 3.4 were employed in the Low-CR optimisation campaigns. The hyper-parameters used for the optimisations are given in Table C.1.

C.1.2 Particle Swarm Optimisation Hyper-Parameters

Although Section 3.6 describes several enhancements to the particle swarm optimisation (PSO) package, only the multi-species capability was employed in the Low-CR optimisation studies. In all cases, the swarm consisted of two species: a larger, more explorative population responsible for broad parameter space coverage, and a smaller population focused on rapid refinement of promising optima. For the 3-pulse optimisations, the swarm was comprised of 64 particles, increasing to 96 particles in the 4-pulse cases to reflect the higher dimensionality of the search space. The swarm was divided in an 80/20 ratio between the two species. Species 1 used hyper-parameters $[0.72, 1, 1]$, while Species 2 used $[0.6, 0.6, 1.1]$, corresponding respectively to the inertia, local attraction, and global attraction coefficients.

C.2 Low-CR Optimisation Results

A summary of the optimised implosions for each optimisation approach (sequential linear benchmark, Bayesian optimisation, and PSO), pulse configuration (3-pulse and 4-pulse), and capsule size (0.5, 0.65, 0.85, 1.0, and 1.1) is provided in Tables C.2 (benchmark results), C.3 (Bayesian results) and C.4 (PSO results).

C.3 Fast Ignition Optimisation Hyper-Parameters

A swarm of 999 particles, divided between two species, was used to explore the 16-dimensional parameter space. The two species employed hyper-parameter sets $[0.5, 0.4, 1.1]$ and $[0.7, 0.7, 1.2]$, assigned with probabilities of 0.3 and 0.7 respectively. Velocity boosting was enabled during this optimisation. Particles were considered stagnant when their velocity fell below $1/e^2$ of the initial velocity for 10 consecutive iterations, at which point their velocity was increased to $1/e$ of the initial value. After each boost, the stagnation threshold was reduced multiplicatively to $1/e^3$, and the process repeated if necessary. These parameters were determined through benchmarking on test functions to ensure recovery from stagnation while preserving convergence stability in the high-dimensional search space.

	3-pulse					4-pulse				
Capsule size	0.5	0.65	0.85	1.0	1.1	0.5	0.65	0.85	1.0	1.1
Random batch size	126	126	128	128	128	128	128	128	128	128
Batch size	24	24	24	32	32	24	24	24	24	24
Sub-batch size	6	6	6	8	8	6	6	6	6	6
Max κ	25	25	25	50	50	25	25	25	25	25
Min κ	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01
First bounds reduction	15	15	18	18	18	20	20	20	25	25
Iterations between reduction	4	4	4	5	4	5	5	5	7	7
Reduction factor	0.8	0.8	0.75	0.8	0.7	0.75	0.8	0.8	0.8	0.8
Noise	10^{-4}	10^{-4}	10^{-3}	10^{-4}	10^{-4}	10^{-4}	10^{-4}	10^{-4}	10^{-3}	10^{-3}

Table C.1: Bayesian optimisation hyper-parameters used for the low-convergence-ratio re-optimisation studies. Values are shown for each capsule size and pulse configuration.

	3-pulse					4-pulse				
Capsule size	0.5	0.65	0.85	1.0	1.1	0.5	0.65	0.85	1.0	1.1
Result	0.13	0.46	8.88	35.2	46.5	0.21	1.0	33.8	44.5	69.5
Gain	0.13	0.46	8.9	35.2	49.6	0.24	1.0	34.1	53.9	69.5
Laser Energy (MJ)	0.80	1.8	4.0	6.6	8.9	0.77	1.7	4.2	6.7	8.5
Inner CR	15.30	15.47	15.88	15.96	15.97	15.99	15.82	15.93	15.95	15.71
Outer CR	12.10	12.22	12.93	13.04	13.26	13.03	12.68	13.06	13.48	12.85
IFAR	20.30	18.66	15.17	14.29	12.69	29.74	25.11	21.19	18.38	15.69
Implosion Velocity (km/s)	400.4	394.3	379.8	359.8	338.2	406.6	397.6	354.9	342.1	320.5
$I\lambda^2$	0.83	0.83	0.83	0.83	0.83	0.83	0.83	0.83	0.83	0.83
Pulse 1 Power (TW)	5.97	10.1	17.2	23.9	28.9	0.50	0.845	1.45	2.00	2.42
Pulse 1 End Time (ns)	3.65	3.80	5.00	6.20	7.00	2.60	3.60	3.70	5.20	5.70
Pulse 2 Power (TW)	53.7	90.8	155	215	260	3.50	5.92	10.1	14.0	16.9
Pulse 2 End Time (ns)	4.15	4.90	6.30	7.60	8.80	5.60	7.80	9.40	12.8	15.0
Pulse 3 Power (TW)	173	292	500	692	838	24.5	41.4	70.8	98.0	119
Pulse 3 End Time (ns)	8.5	10.5	13.8	16.5	18.6	6.80	9.50	11.4	15.3	18.0
Pulse 4 Power (TW)	-	-	-	-	-	173	292	500	692	838
Pulse 4 End Time (ns)	-	-	-	-	-	11.0	15.0	19.4	24.5	27.5

Table C.2: Results from the optimised Low-CR implosions obtained using the sequential linear scan benchmark method for all capsule sizes and pulse configurations.

	3-pulse					4-pulse				
Capsule size	0.5	0.65	0.85	1.0	1.1	0.5	0.65	0.85	1.0	1.1
Result	0.18	0.89	17.1	44.7	56.6	0.23	2.3	40.1	58.4	83.9
Gain	0.19	1.06	19.0	45.0	56.6	0.25	2.70	41.8	58.4	83.9
Laser Energy (MJ)	0.80	1.8	4.7	6.3	8.4	0.74	1.8	4.2	6.1	7.9
Inner CR	15.90	16.22	16.54	16.66	16.61	16.22	16.10	16.66	16.34	16.66
Outer CR	12.75	12.90	12.93	13.04	12.98	13.18	12.75	13.20	13.00	13.00
IFAR	19.88	20.37	19.28	18.37	18.53	29.74	28.55	22.09	20.91	17.27
Implosion Velocity (km/s)	403.4	404.2	388.9	367.6	337.1	401.8	404.4	354.6	316.4	290.3
$I\lambda^2$	0.83	0.86	0.86	0.83	0.78	0.81	0.86	0.83	0.68	0.65
Pulse 1 Power (TW)	4.47	12.1	22.5	9.88	5.75	0.378	0.475	4.36	1.56	2.06
Pulse 1 End Time (ns)	3.77	3.33	3.95	5.25	6.90	3.69	6.20	4.34	7.86	8.18
Pulse 2 Power (TW)	44.3	49.2	74.1	60.9	58.2	4.32	7.60	25.6	19.1	21.1
Pulse 2 End Time (ns)	4.59	4.84	6.09	9.5	13.1	7.07	10.4	7.93	14.5	17.7
Pulse 3 Power (TW)	174	302	517	689	784	47.5	48.1	96.5	107	155
Pulse 3 End Time (ns)	8.82	10.5	20.7	18.2	23.3	7.81	12.1	9.90	17.0	20.4
Pulse 4 Power (TW)	-	-	-	-	-	169	301	501	570	653
Pulse 4 End Time (ns)	-	-	-	-	-	11.9	17.6	17.7	26.9	31.5

Table C.3: Results from the optimised Low-CR implosions obtained using the Bayesian optimisation method for all capsule sizes and pulse configurations.

	3-pulse					4-pulse				
Capsule size	0.5	0.65	0.85	1.0	1.1	0.5	0.65	0.85	1.0	1.1
Result	0.17	0.82	17.5	45.1	56.7	0.23	1.81	40.4	56.8	83.2
Gain	0.18	0.95	19.0	45.3	56.7	0.25	2.25	41.8	57.7	83.2
Laser Energy (MJ)	0.79	1.8	4.7	6.2	8.4	0.75	1.8	4.1	6.0	7.8
Inner CR	15.80	16.17	16.31	16.39	16.59	16.26	15.97	16.73	16.40	16.92
Outer CR	12.65	12.85	12.84	13.03	12.95	13.18	12.91	13.17	13.09	13.00
IFAR	18.82	19.13	17.62	19.29	18.68	29.12	26.33	21.7	19.56	18.30
Implosion Velocity (km/s)	403.1	403.3	389.9	369.4	337.4	402.6	403.7	354.9	312.6	281.3
$I\lambda^2$	0.83	0.86	0.85	0.84	0.78	0.82	0.86	0.84	0.68	0.62
Pulse 1 Power (TW)	4.18	9.45	19.0	6.93	5.57	0.462	0.515	4.58	1.82	3.34
Pulse 1 End Time (ns)	3.42	3.53	4.06	4.98	6.90	3.41	5.95	4.62	5.99	6.99
Pulse 2 Power (TW)	16.7	44.0	67.5	44.8	57.3	4.56	9.49	25.4	14.8	18.7
Pulse 2 End Time (ns)	4.62	5.14	6.09	10.2	13.1	6.67	10.8	8.04	13.4	16.8
Pulse 3 Power (TW)	173	302	514	696	784	48.1	94.2	100	110	142
Pulse 3 End Time (ns)	8.99	10.8	15.8	18.8	23.3	7.43	11.8	9.99	16.0	19.5
Pulse 4 Power (TW)	-	-	-	-	-	171	304	503	568	627
Pulse 4 End Time (ns)	-	-	-	-	-	11.5	17.2	17.6	25.8	31.0

Table C.4: Results from the optimised Low-CR implosions obtained using the particle swarm optimisation method for all capsule sizes and pulse configurations.

D

Beam-Plasma System

D.1 Linearisaing the Lorentz gamma factor

Linearisation of the Lorentz gamma factor can be done as shown below via Taylor expansion:

$$\gamma(\beta) = \left(1-\beta^2\right)^{-\frac{1}{2}}, \tag{D.1}$$

$$\approx \gamma(\boldsymbol{\beta_0}) + \left.\frac{d\gamma(\beta)}{d\beta}\right|_{\beta=\boldsymbol{\beta_0}} \cdot \boldsymbol{\beta_1} \tag{D.2}$$

$$\approx \left(1-\beta_0^2\right)^{-\frac{1}{2}} + \frac{d}{d\beta}\left[\left(1-\beta^2\right)^{-\frac{1}{2}}\right]\Big|_{\beta=\boldsymbol{\beta_0}} \cdot \beta_1, \tag{D.3}$$

$$\approx \left(1-\beta_0^2\right)^{-\frac{1}{2}} - \frac{1}{2}(1-\beta_0^2)^{-\frac{3}{2}}(-2\beta_0)\cdot\beta_1 \tag{D.4}$$

$$\approx \left(1-\beta_0^2\right)^{-\frac{1}{2}} + (\beta_0\cdot\beta_1)(1-\beta_0^2)^{-\frac{3}{2}}, \tag{D.5}$$

$$\approx \gamma_0 + \gamma_1, \tag{D.6}$$

$$\approx \gamma_0 + (\beta_0\cdot\beta_1)\gamma_0^3 \tag{D.7}$$

D.2 Matrix describing the linearised system, M

The full 18×18 matrix describing the multi-fluid linearised system is split into four blocks, A, B, C and D, with the matrix $M = (A, B, C, D)$. The matrix

uses the following notation:

$$\omega'_\alpha = -\omega + k_x\beta_{x,\alpha} + k_y\beta_{y,\alpha} + k_z\beta_{z,\alpha}, \tag{D.8}$$

$$\gamma_\alpha^3 = (1 - \beta_{x,\alpha}^2 - \beta_{y,\alpha}^2 - \beta_{z,\alpha}^2)^{3/2} \tag{D.9}$$

where α refers to the species. These are all unperturbed quantities since the perturbed quantities have been extracted to the $\mathbf{u}$ vector ($\mathbf{M} \cdot \mathbf{u} = 0$). As an example, $\beta_{x,1}$ is the x-component of species 1's initial velocity, as given as the 6^{th} element of $\mathbf{U_0}$ (Equation 5.1).

$$A = \begin{pmatrix}
i\omega'_0 & ik_x n_0 & ik_y n_0 & ik_z n_0 \\
0 & \frac{i\omega'_0}{\gamma_0^3}(1 - \beta_{y,0}^2 - \beta_{z,0}^2) & \frac{i\omega'_0}{\gamma_0^3}(\beta_{y,0} + \beta_{x,0}) & \frac{i\omega'_0}{\gamma_0^3}(\beta_{z,0} + \beta_{x,0}) \\
0 & \frac{i\omega'_0}{\gamma_0^3}(\beta_{x,0} + \beta_{y,0}) & \frac{i\omega'_0}{\gamma_0^3}(1 - \beta_{x,0}^2 + \beta_{z,0}^2) & \frac{i\omega'_0}{\gamma_0^3}(\beta_{z,0} + \beta_{y,0}) \\
0 & \frac{i\omega'_0}{\gamma_0^3}(\beta_{x,0} + \beta_{z,0}) & \frac{i\omega'_0}{\gamma_0^3}(\beta_{y,0} + \beta_{z,0}) & \frac{i\omega'_0}{\gamma_0^3}(1 - \beta_{x,0}^2 + \beta_{y,0}^2) \\
0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 \\
\beta_{x0} & n_0 & 0 & 0 \\
\beta_{y0} & 0 & n_0 & 0 \\
\beta_{z0} & 0 & 0 & n_0 \\
0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0
\end{pmatrix} \tag{D.10}$$

$$B = \begin{pmatrix}
0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 \\
i\omega_1' & ik_x n_1 & ik_y n_1 & ik_z n_1 \\
0 & \frac{i\omega_1'}{\gamma_1^3}(1-\beta_{y,1}^2-\beta_{z,1}^2) & \frac{i\omega_1'}{\gamma_1^3}(\beta_{y,1}+\beta_{x,1}) & \frac{i\omega_1'}{\gamma_1^3}(\beta_{z,1}+\beta_{x,1}) \\
0 & \frac{i\omega_1'}{\gamma_1^3}(\beta_{x,1}+\beta_{y,1}) & \frac{i\omega_1'}{\gamma_1^3}(1-\beta_{x,1}^2+\beta_{z,1}^2) & \frac{i\omega_1'}{\gamma_1^3}(\beta_{z,1}+\beta_{y,1}) \\
0 & \frac{i\omega_1'}{\gamma_1^3}(\beta_{x,1}+\beta_{z,1}) & \frac{i\omega_1'}{\gamma_1^3}(\beta_{y,1}+\beta_{z,1}) & \frac{i\omega_1'}{\gamma_1^3}(1-\beta_{x,1}^2+\beta_{y,1}^2) \\
0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 \\
\beta_{x1} & n_1 & 0 & 0 \\
\beta_{y1} & 0 & n_1 & 0 \\
\beta_{z1} & 0 & 0 & n_1 \\
0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0
\end{pmatrix} \tag{D.11}$$

$$
C = \begin{pmatrix}
0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 \\
i\omega_2' & ik_x n_2 & ik_y n_2 & ik_z n_2 \\
0 & \frac{i\omega_2'}{\gamma_2^3}(1-\beta_{y,2}^2-\beta_{z,2}^2) & \frac{i\omega_2'}{\gamma_2^3}(\beta_{y,2}+\beta_{x,2}) & \frac{i\omega_2'}{\gamma_2^3}(\beta_{z,2}+\beta_{x,2}) \\
0 & \frac{i\omega_2'}{\gamma_2^3}(\beta_{x,2}+\beta_{y,2}) & \frac{i\omega_2'}{\gamma_2^3}(1-\beta_{x,2}^2+\beta_{z,2}^2) & \frac{i\omega_2'}{\gamma_2^3}(\beta_{z,2}+\beta_{y,2}) \\
0 & \frac{i\omega_2'}{\gamma_2^3}(\beta_{x,2}+\beta_{z,2}) & \frac{i\omega_2'}{\gamma_2^3}(\beta_{y,2}+\beta_{z,2}) & \frac{i\omega_2'}{\gamma_2^3}(1-\beta_{x,2}^2+\beta_{y,2}^2) \\
\beta_{x2} & n_2 & 0 & 0 \\
\beta_{y2} & 0 & n_2 & 0 \\
\beta_{z2} & 0 & 0 & n_2 \\
0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0
\end{pmatrix} \tag{D.12}
$$

$$
D = \begin{pmatrix}
0 & 0 & 0 & 0 & 0 & 0 \\
1 & 0 & 0 & 0 & -\beta_{z0} & \beta_{y0} \\
0 & 1 & 0 & \beta_{z0} & 0 & -\beta_{x0} \\
0 & 0 & 1 & -\beta_{y0} & \beta_{x0} & 0 \\
0 & 0 & 0 & 0 & 0 & 0 \\
1 & 0 & 0 & 0 & -\beta_{z1} & \beta_{y1} \\
0 & 1 & 0 & \beta_{z1} & 0 & -\beta_{x1} \\
0 & 0 & 1 & -\beta_{y1} & \beta_{x1} & 0 \\
0 & 0 & 0 & 0 & 0 & 0 \\
1 & 0 & 0 & 0 & -\beta_{z2} & \beta_{y2} \\
0 & 1 & 0 & \beta_{z2} & 0 & -\beta_{x2} \\
0 & 0 & 1 & -\beta_{y2} & \beta_{x2} & 0 \\
i\omega & 0 & 0 & 0 & -ik_z & ik_y \\
0 & i\omega & 0 & ik_z & 0 & -ik_x \\
0 & 0 & i\omega & -ik_y & ik_x & 0 \\
0 & -ik_z & ik_y & -i\omega & 0 & 0 \\
ik_z & 0 & -ik_x & 0 & -i\omega & 0 \\
-ik_y & ik_x & 0 & 0 & 0 & -i\omega
\end{pmatrix} \tag{D.13}
$$

References

[1] A. Einstein. "Zur Elektrodynamik bewegter Körper". *Annalen der Physik* 322.10 (1905), pp. 891–921.

[2] Joseph John Thomson. *Rays of positive electricity and their application to chemical analyses*. Longmans, Green and Company, 1921.

[3] F.W. Aston. "LXXIV. A positive ray spectrograph". *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* 38.228 (1919), pp. 707–714.

[4] F.W. Aston. "LIX. The mass-spectra of chemical elements". *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* 39.233 (1920), pp. 611–625.

[5] A. S. Eddington. "The Internal Constitution of the Stars". *The Scientific Monthly* 11.4 (1920), pp. 297–303.

[6] George Gamow. "The Quantum Theory of Nuclear Disintegration". *Nature* 122 (1928), pp. 805–806.

[7] E. Rutherford. "LXXIX. The scattering of α and β particles by matter and the structure of the atom". *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* 21.125 (1911), pp. 669–688.

[8] J. D. Cockcroft and E. T. S. Walton. "Artificial Production of Fast Protons". *Nature* 129.3250 (1932), pp. 242–242.

[9] J. D. Cockcroft and E. T. S. Walton. "Disintegration of Lithium by Swift Protons". *Nature* 129.3261 (1932), pp. 649–649.

[10] James Chadwick. "Possible Existence of a Neutron". *Nature* 129.3252 (1932), pp. 312–312.

[11] James Chadwick. "The existence of a neutron". *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character* 136.830 (1932), pp. 692–708.

[12] M. L. Oliphant, P. Harteck, and E. Rutherford. "Transmutation Effects Observed with Heavy Hydrogen". *Nature* 133.3359 (1934), pp. 413–413.

[13] E. Bretscher and A. P. French. "Low Energy Cross Section of the $D-T$ Reaction and Angular Distribution of the Alpha-Particles Emitted". *Physical Review* 75 (1949), pp. 1154–1160.

[14] W. R. Arnold et al. "Cross Sections for the Reactions $\mathrm{D}(d,p)\mathrm{T}$, $\mathrm{D}(d,n)\mathrm{He}^3$, $\mathrm{T}(d,n)\mathrm{He}^4$, and $\mathrm{He}^3(d,p)\mathrm{He}^4$ below 120 kev". *Physical Review* 93 (1954), pp. 483–497.

[15] V.P. Smirnov. "Tokamak foundation in USSR/Russia 1950-1990". *Nuclear Fusion* 50.1 (2009), p. 014003.

[16] M G Haines. “Fifty years of controlled fusion research”. *Plasma Physics and Controlled Fusion* 38.5 (1996), p. 643.

[17] L M Goldman and L Jr Spitzer. *PRELIMINARY EXPERIMENTAL RESULTS WITH THE MODEL A STELLARATOR*. Tech. rep. Princeton Univ., N.J. Project Matterhorn, May 1953.

[18] J. Nuckolls et al. “Laser Compression of Matter to Super-High Densities: Thermonuclear (CTR) Applications”. *Nature* 239.5368 (1972), pp. 139–142.

[19] Lawrence Livermore National Lab. *Early steps toward inertial fusion energy (IFE) (1952 to 1962)*. Tech. rep. Lawrence Livermore National Lab., CA (United States), June 1998.

[20] T. H. Maiman. “Stimulated Optical Radiation in Ruby”. *Nature* 187.4736 (1960), pp. 493–494.

[21] Shutaro Takeda, Alexander Ryota Keeley, and Shunsuke Managi. “How Many Years Away is Fusion Energy? A Review”. *Journal of Fusion Energy* 42.1 (2023), p. 16.

[22] Jeremy P Chittenden. “The Z-pinch: approach to fusion”. *Physics World* 13.5 (2000), p. 39.

[23] K. Boyer et al. “Studies of Plasma Heated in a Fast-Rising Axial Magnetic Field (Scylla)”. *Physical Review* 119 (1960), pp. 831–843.

[24] J. J. Lee et al. “Toward more robust ignition of inertial fusion targets”. *Physics of Plasmas* 30.2 (2023), p. 022702.

[25] Jordan Lee et al. “Measuring the principle Hugoniot of low-density silica aerogel foam at pressures up to 160 GPa”. *Physical Review E* 113 (2026), p. 055210.

[26] J. J. Lee et al. “Bayesian and particle swarm approaches to inertial confinement fusion optimization”. *Physics of Plasmas* 33.6 (2026), p. 062702.

[27] R. W. Paddock et al. “Energy gain of wetted-foam implosions with auxiliary heating for inertial fusion studies”. *Plasma Physics and Controlled Fusion* 66.2 (2024).

[28] R. J. L. Timmis et al. “Attosecond and nano-Coulomb electron bunches via the Zero Vector Potential mechanism”. *Scientific Reports* 14 (2024), p. 10805.

[29] Eduard Atonga et al. “Gravitational waves from high-power twisted light”. *Physical Review D* 110 (2024), p. 044023.

[30] H. Martin et al. “Electrothermal filamentation of igniting plasmas”. *Physical Review E* 110 (2024), p. 035205.

[31] Robin J. L. Timmis et al. “Efficiency-optimized relativistic plasma harmonics for extreme fields”. *Nature* 652.8112 (2026), pp. 1153–1158.

[32] Irving Langmuir. “Oscillations in Ionized Gases”. *Proceedings of the National Academy of Sciences* 14.8 (1928), pp. 627–637.

[33] Francis F. Chen. *Introduction to plasma physics and controlled fusion*. 2018.

[34] Goldston R J and Paul H Rutherford. “Introduction to plasmas”. *Introduction to Plasma Physics*. CRC Press, 1995, pp. 1–18.

[35] Joachim Oxenius. “Thermal Equilibrium and Detailed Balance”. *Kinetic Theory of Particles and Photons: Theoretical Foundations of Non-LTE Plasma Spectroscopy*. Springer Berlin Heidelberg, 1986, pp. 1–34.

[36] Frederick Reif. *Fundamentals of statistical and thermal physics*. Levant Books, 2010.

[37] J. C. Maxwell. “V. Illustrations of the dynamical theory of gases.—Part I. On the motions and collisions of perfectly elastic spheres”. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* 19.124 (1860), pp. 19–32.

[38] LUDWIG BOLTZMANN. “Further Studies on the Thermal Equilibrium of Gas Molecules”. *The Kinetic Theory of Gases*, pp. 262–349.

[39] D. Bohm and E. P. Gross. “Theory of Plasma Oscillations. A. Origin of Medium-Like Behavior”. *Physical Review* 75 (1949), pp. 1851–1864.

[40] Albert Einstein. “Zur Quantentheorie der Strahlung”. German. *Physikalische Zeitschrift* 18 (1917), pp. 121–128.

[41] J. P. Gordon, H. J. Zeiger, and C. H. Townes. “Molecular Microwave Oscillator and New Hyperfine Structure in the Microwave Spectrum of NH_3”. *Physical Review* 95 (1954), pp. 282–284.

[42] Max Planck. “Ueber das Gesetz der Energieverteilung im Normalspectrum”. *Annalen der Physik* 309.3 (1901), pp. 553–563.

[43] J. E. Geusic, H. M. Marcos, and L. G. Van Uitert. “LASER OSCILLATIONS IN Nd-DOPED YTTRIUM ALUMINUM, YTTRIUM GALLIUM AND GADOLINIUM GARNETS”. *Applied Physics Letters* 4.10 (1964), pp. 182–184.

[44] W. Sun et al. “High Energy Flash-Lamp Pumped Nd:YAG Laser Emitting at 1073.8 nm”. *Laser Physics* 20 (2010), pp. 1324–1328.

[45] R. Burnham and N. Djeu. “Ultraviolet-preionized discharge-pumped lasers in XeF, KrF, and ArF”. *Applied Physics Letters* 29.11 (1976), pp. 707–709.

[46] T. R. Loree, K. B. Butterfield, and D. L. Barker. “Spectral tuning of ArF and KrF discharge lasers”. *Applied Physics Letters* 32.3 (1978), pp. 171–173.

[47] E. I. Moses. *The National Ignition Facility: Status and Plans for Laser Fusion and High-Energy-Density Experimental Studies*. 2001.

[48] Donna Strickland and Gerard Mourou. “Compression of amplified chirped optical pulses”. *Optics Communications* 55.6 (1985), pp. 447–449.

[49] Thomas H. Stix. *Waves in plasmas*. American Institute of Physics, 1992.

[50] William L. Kruer. *The physics of laser plasma interactions*. Frontiers in physics. Westview, 1988.

[51] L V Keldysh. “IONIZATION IN THE FIELD OF A STRONG ELECTROMAGNETIC WAVE”. *Zh. Eksperim. i Teor. Fiz.* Vol: 47 (Nov. 1964).

[52] W. Becker et al. “Above-Threshold Ionization: From Classical Features to Quantum Effects”. Vol. 48. Advances In Atomic, Molecular, and Optical Physics. Academic Press, 2002, pp. 35–98.

[53] H. A. Kramers. “XCIII. On the theory of X-ray absorption and of the continuous X-ray spectrum”. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* 46.275 (1923), pp. 836–871.

[54] "BREMSSTRAHLUNG". *Radiative Processes in Astrophysics*. John Wiley & Sons, Ltd, 1985. Chap. 5, pp. 155–166.

[55] Jeff Colvin and Jon Larsen. *Extreme physics : properties and behavior of matter at extreme conditions*. Cambridge University Press, 2014.

[56] Pawan Kumar and V. K. Tripathi. "Vacuum electron heating by surface plasma wave". *Applied Physics Letters* 99.2 (2011), p. 021502.

[57] Hong-bo Cai et al. "Short-pulse laser absorption via JxB heating in ultrahigh intensity laser plasma interaction". *Physics of Plasmas* 13.11 (2006), p. 113105.

[58] L. D. Landau and E. M. Lifshits. *Fluid mechanics*. Course of theoretical physics ; Volume 6. Pergamon Press, 1987.

[59] Stefano Atzeni and Jürgen Meyer-Ter-Vehn. "Inertial confinement by spherical implosion". *The Physics of Inertial Fusion: BeamPlasma Interaction, Hydrodynamics, Hot Dense Matter*. Oxford University Press, 2004.

[60] Shigeo Kawata. "Direct-drive heavy ion beam inertial confinement fusion: a review, toward our future energy source". *Advances in Physics: X* 6.1 (2021), p. 1873860.

[61] Michael Schirber. "Inertial-Confinement Fusion without Lasers". *Physics* 17 (2024), p. 22.

[62] R. Betti and O. A. Hurricane. "Inertial-confinement fusion with lasers". *Nature Physics* 12.5 (2016), pp. 435–448.

[63] A. L. Kritcher et al. "Design of the first fusion experiment to achieve target energy gain G > 1". *Physical Review E* 109 (2024), p. 025204.

[64] R. S. Craxton et al. "Direct-drive inertial confinement fusion: A review". *Physics of Plasmas* 22.11 (2015), p. 110501.

[65] John Lindl. "Development of the indirect-drive approach to inertial confinement fusion and the target physics basis for ignition and gain". *Physics of Plasmas* 2.11 (1995), pp. 3933–4024.

[66] J. D. Huba. *NRL Plasma Formulary*. Naval Research Laboratory, 2023.

[67] K. Batani. "Perspectives on research on laser driven proton-boron fusion and applications". *Journal of Instrumentation* 18.09 (2023), p. C09012.

[68] Tetsuo Tanabe. *Tritium : Fuel of Fusion Reactors*. Springer Japan, 2017.

[69] O. A. Hurricane et al. "Physics principles of inertial confinement fusion and U.S. program overview". *Reviews of Modern Physics* 95 (2023), p. 025005.

[70] Max Tabak et al. "Ignition and high gain with ultrapowerful lasers*". *Physics of Plasmas* 1.5 (1994), pp. 1626–1634.

[71] J D Lawson. "Some Criteria for a Power Producing Thermonuclear Reactor". *Proceedings of the Physical Society. Section B* 70.1 (1957), p. 6.

[72] A. B. Zylstra, D. A. Callahan, O. A. Hurricane, et al. "Burning plasma achieved in inertial fusion". *Nature* 601 (2022), pp. 542–548.

[73] A. B. Zylstra et al. "Experimental achievement and signatures of ignition at the National Ignition Facility". *Physical Review E* 106 (2022), p. 025202.

[74] H. Abu-Shawareb et al. "Achievement of Target Gain Larger than Unity in an Inertial Fusion Experiment". *Physical Review Letters* 132 (2024), p. 065102.

[75] Lawrence Livermore National Laboratory. *Achieving Fusion Ignition*. National Ignition Facility & Photon Science. 2025.

[76] Stefano Atzeni and Max Tabak. "Overview of ignition conditions and gain curves for the fast ignitor". *Plasma Physics and Controlled Fusion* 47.12B (2005), B769.

[77] M. Gatu Johnson et al. "Impact of asymmetries on fuel performance in inertial confinement fusion". *Physical Review E* 98 (2018), p. 051201.

[78] B. Bachmann et al. "Measuring and simulating ice-ablator mix in inertial confinement fusion". *Physics of Plasmas* 30.5 (2023), p. 052704.

[79] R. E. Olson et al. "Preheat Effects on Shock Propagation in Indirect-Drive Inertial Confinement Fusion Ablator Materials". *Physical Review Letters* 91 (2003), p. 235002.

[80] A. G. Seaton et al. "Cross-beam energy transfer in direct-drive ICF. II. Theory and simulation of mitigation through increased laser bandwidth". *Physics of Plasmas* 29.4 (2022), p. 042707.

[81] Stefano Atzeni and Jürgen Meyer-Ter-Vehn. "Hydrodynamic stability". *The Physics of Inertial Fusion: BeamPlasma Interaction, Hydrodynamics, Hot Dense Matter*. Oxford University Press, 2004.

[82] H. Takabe et al. "Self-consistent growth rate of the Rayleigh-Taylor instability in an ablatively accelerating plasma". *The Physics of Fluids* 28.12 (1985), pp. 3676–3682.

[83] Hiroki Kato et al. "Synthesis and characterization of diamond capsules for direct-drive inertial confinement fusion". *Diamond and Related Materials* 86 (2018), pp. 15–19.

[84] Y. Kato et al. "Random Phasing of High-Power Lasers for Uniform Target Acceleration and Plasma-Instability Suppression". *Physical Review Letters* 53 (1984), pp. 1057–1060.

[85] S. Skupsky et al. "Improved laser-beam uniformity using the angular dispersion of frequency-modulated light". *Journal of Applied Physics* 66.8 (1989), pp. 3456–3462.

[86] V. N. Goncharov et al. "Improved performance of direct-drive inertial confinement fusion target designs with adiabat shaping using an intensity picket". *Physics of Plasmas* 10.5 (2003), pp. 1906–1918.

[87] K. Anderson and R. Betti. "Laser-induced adiabat shaping by relaxation in inertial fusion implosions". *Physics of Plasmas* 11.1 (2004), pp. 5–8.

[88] D.E. Hinkel et al. "Optimization of capsule dopant levels to improve fuel areal density*". *High Energy Density Physics* 37 (2020), p. 100884.

[89] S. X. Hu et al. "Mitigating laser-imprint effects in direct-drive inertial confinement fusion implosions with an above-critical-density foam layer". *Physics of Plasmas* 25.8 (2018), p. 082710.

[90] D. X. Liu et al. "Mitigating laser imprint with a foam overcoating". *Physics of Plasmas* 29.7 (2022), p. 072707.

[91] Y. Zhao, S. M. Weng, H. H. Ma, et al. “Mitigation of laser plasma parametric instabilities with broadband lasers”. *Reviews of Modern Plasma Physics* 7 (2023), p. 1.

[92] I. V. Igumenshchev et al. “Laser-Beam Zooming to Mitigate Crossed-Beam Energy Losses in Direct-Drive Implosions”. *Physical Review Letters* 110 (2013), p. 145001.

[93] M. Gatu Johnson et al. “Neutron spectrometry—An essential tool for diagnosing implosions at the National Ignition Facility (invited)a)”. *Review of Scientific Instruments* 83.10 (2012), p. 10D308.

[94] D. S. Clark et al. “Three-dimensional simulations of low foot and high foot implosion experiments on the National Ignition Facility”. *Physics of Plasmas* 23.5 (2016), p. 056302.

[95] Charles Lion. “The LMJ program: An overview”. *Journal of Physics: Conference Series* 244.1 (2010), p. 012003.

[96] D Mihalas and B W Mihalas. *Foundations of radiation hydrodynamics*. Oxford University Press,New York, NY, Jan. 1984.

[97] J. VonNeumann and R. D. Richtmyer. “A Method for the Numerical Calculation of Hydrodynamic Shocks”. *Journal of Applied Physics* 21.3 (1950), pp. 232–237.

[98] Jon T. Larsen and Stephen M. Lane. “HYADES—A plasma hydrodynamics code for dense plasma studies”. *Journal of Quantitative Spectroscopy and Radiative Transfer* 51.1 (1994), pp. 179–186.

[99] J.J. MacFarlane, I.E. Golovkin, and P.R. Woodruff. “HELIOS-CR - A 1-D radiation-magnetohydrodynamics code with inline atomic kinetics modeling”. *Journal of Quantitative Spectroscopy and Radiative Transfer* 99.1 (2006), pp. 381–397.

[100] H. Hora. “D. R. Nicholson, Introduction to Plasma Theory, John Wiley & Sons, New York, 1983, XII, 292 pp.” *Laser and Particle Beams* 2.1 (1984), pp. 127–127.

[101] Charles K. Birdsall, A. Bruce. Langdon, and Institute of Physics (Great Britain). *Plasma physics via computer simulation*. Series in plasma physics. Taylor& Francis, 2005.

[102] T D Arber et al. “Contemporary particle-in-cell approach to laser-plasma modelling”. *Plasma Physics and Controlled Fusion* 57.11 (2015), p. 113001.

[103] J. Derouillat et al. “Smilei : A collaborative, open-source, multi-purpose particle-in-cell code for plasma simulation”. *Computer Physics Communications* 222 (2018), pp. 351–373.

[104] C.Z Cheng and Georg Knorr. “The integration of the vlasov equation in configuration space”. *Journal of Computational Physics* 22.3 (1976), pp. 330–351.

[105] N.J. Sircombe and T.D. Arber. “VALIS: A split-conservative scheme for the relativistic 2D Vlasov-Maxwell system”. *Journal of Computational Physics* 228.13 (2009), pp. 4773–4788.

[106] L. Divol et al. “Thermonuclear performance variability near ignition at the National Ignition Facility”. *Physics of Plasmas* 31.10 (2024), p. 102703.

[107] M. Brönner et al. “Particle swarm optimization of 1D isochoric compression designs for fast ignition”. *Physics of Plasmas* 32.2 (2025), p. 022710.

[108] V. Gopalaswamy et al. "Automated and highly parallelized Bayesian optimization scheme for direct drive fusion experiments on OMEGA". *Physical Review Research* 7 (2025), p. 013009.

[109] A. J. Crilly et al. "Automated simulation-based design via multi-fidelity active learning and optimization for laser direct drive implosions". *Physics of Plasmas* 33.1 (2026), p. 012702.

[110] P. W. Hatfield, S. J. Rose, and R. H. H. Scott. "The blind implosion-maker: Automated inertial confinement fusion experiment design". *Physics of Plasmas* 26.6 (2019), p. 062706.

[111] M. Pokornik et al. "Investigating performance and variability of NIF ICF experiments with deep learning". *Nuclear Fusion* 65.7 (2025), p. 076016.

[112] Roman Garnett. *Bayesian Optimization*. Cambridge University Press, 2023.

[113] M. D. McKay, R. J. Beckman, and W. J. Conover. "A Comparison of Three Methods for Selecting Values of Input Variables in the Analysis of Output from a Computer Code". *Technometrics* 21.2 (1979), pp. 239–245.

[114] Douglas C. Montgomery, Elizabeth A. Peck, and G Geoffrey. Vining. *Introduction to Linear Regression Analysis*. Wiley Series in Probability and Statistics Ser. John Wiley & Sons, Incorporated, 2012.

[115] Jorge Nocedal, Stephen Wright, and Stephen Wright. *Numerical Optimization*. Springer Series in Operations Research and Financial Engineering. Springer Nature, 2006.

[116] Christoph Witzgall. "Practical Methods of Optimization". *Mathematics of computation* 53.188 (1989), pp. 768–769.

[117] Donald R. Jones, Matthias Schonlau, and William J. Welch. "Efficient Global Optimization of Expensive Black-Box Functions". *Journal of Global Optimization* 13.4 (1998), pp. 455–492.

[118] Charles Audet and Warren. Hare. *Derivative-Free and Blackbox Optimization*. Springer Series in Operations Research and Financial Engineering. 2017.

[119] J. Kennedy. "The particle swarm: social adaptation of knowledge". *Proceedings of 1997 IEEE International Conference on Evolutionary Computation (ICEC '97)*. 1997, pp. 303–308.

[120] Carl Edward. Rasmussen and Christopher K. I. Williams. *Gaussian processes for machine learning*. Adaptive computation and machine learning. MIT Press, 2006.

[121] Christopher M. Bishop. *Pattern recognition and machine learning*. Information Science and Statistics. Springer, 2009 - 2006.

[122] Bobak Shahriari et al. "Taking the Human Out of the Loop: A Review of Bayesian Optimization". *Proceedings of the IEEE* 104.1 (2016), pp. 148–175.

[123] Alessandro De Palma et al. *Sampling Acquisition Functions for Batch Bayesian Optimization*. 2019.

[124] Jasper Snoek, Hugo Larochelle, and Ryan P Adams. "Practical Bayesian Optimization of Machine Learning Algorithms". *Advances in Neural Information Processing Systems*. Vol. 25. Curran Associates, Inc., 2012.

[125] David Eriksson et al. *Scalable Global Optimization via Local Bayesian Optimization*. 2020.

[126] Carl Edward Rasmussen. "Gaussian Processes in Machine Learning". *Advanced Lectures on Machine Learning: ML Summer Schools 2003, Canberra, Australia, February 2 - 14, 2003, Tübingen, Germany, August 4 - 16, 2003, Revised Lectures*. Springer Berlin Heidelberg, 2004, pp. 63–71.

[127] Riccardo Poli, James Kennedy, and Tim Blackwell. "Particle swarm optimization: An overview". *Swarm intelligence* 1.1 (2007), pp. 33–57.

[128] M. Clerc and J. Kennedy. "The particle swarm - explosion, stability, and convergence in a multidimensional complex space". *IEEE Transactions on Evolutionary Computation* 6.1 (2002), pp. 58–73.

[129] A. Ratnaweera, S.K. Halgamuge, and H.C. Watson. "Self-organizing hierarchical particle swarm optimizer with time-varying acceleration coefficients". *IEEE Transactions on Evolutionary Computation* 8.3 (2004), pp. 240–255.

[130] Weilin Du and Bin Li. "Multi-strategy ensemble particle swarm optimization for dynamic optimization". *Information Sciences* 178.15 (2008), pp. 3096–3109.

[131] Jacques Riget and Jakob Vesterstrom. "A Diversity-Guided Particle Swarm Optimizer - the ARPSO". 2002.

[132] Xiao-Feng Xie, Wen-Jun Zhang, and Zhi-Lian Yang. *A dissipative particle swarm optimization*. 2005.

[133] A. B. Zylstra et al. "Variable convergence liquid layer implosions on the National Ignition Facility". *Physics of Plasmas* 25.5 (2018), p. 056304.

[134] Brian M. Haines et al. "The effects of convergence ratio on the implosion behavior of DT layered inertial confinement fusion capsules". *Physics of Plasmas* 24.7 (2017), p. 072709.

[135] R. E. Olson and R. J. Leeper. "Alternative hot spot formation techniques using liquid deuterium-tritium layer inertial confinement fusion capsules". *Physics of Plasmas* 20.9 (2013), p. 092705.

[136] G. E. Kemp et al. "Exploration of polar direct drive wetted foam concepts for neutron sources on the National Ignition Facility laser". *Physics of Plasmas* 32.2 (2025), p. 022702.

[137] R. W. Paddock et al. "One-dimensional hydrodynamic simulations of low convergence ratio direct-drive inertial confinement fusion implosions". *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 379.20200224 (2021).

[138] A. Caruso and V.A. Pais. "The ignition of dense DT fuel by injected triggers". *Nuclear Fusion* 36.6 (1996), p. 745.

[139] M. Roth et al. "Fast Ignition by Intense Laser-Accelerated Proton Beams". *Physical Review Letters* 86 (2001), pp. 436–439.

[140] R. Kodama et al. "Fast heating scalable to laser fusion ignition". *Nature* 418.6901 (2002), pp. 933–934.

[141] J Meyer-ter-Vehn. "Fast ignition of ICF targets: an overview". *Plasma Physics and Controlled Fusion* 43.12A (2001), A113.

[142] S. P. Obenschain et al. "Direct drive with the argon fluoride laser as a path to high fusion gain with sub-megajoule laser energy". *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 378.2184 (2020), p. 20200031.

[143] J. Biener et al. "A new approach to foam-lined indirect-drive NIF ignition targets". *Nuclear Fusion* 52.6 (2012), p. 062001.

[144] Curtis Walters et al. "D2 and D-T Liquid-Layer Target Shots at the National Ignition Facility". *Fusion Science and Technology* 73.3 (2018), pp. 305–314.

[145] B. J. Kozioziemski et al. "Deuterium-Tritium Fuel Layer Formation for the National Ignition Facility". *Fusion Science and Technology* 59.1 (2011), pp. 14–25.

[146] S. W. Haan et al. "Point design targets, specifications, and requirements for the 2010 ignition campaign on the National Ignition Facility". *Physics of Plasmas* 18.5 (2011), p. 051001.

[147] A. A. Chernov et al. "Single crystal growth and formation of defects in deuterium-tritium layers for inertial confinement nuclear fusion". *Applied Physics Letters* 94.6 (2009), p. 064105.

[148] James K. Hoffer and Larry R. Foreman. "Radioactively induced sublimation in solid tritium". *Physical Review Letters* 60 (1988), pp. 1310–1313.

[149] R. E. Olson et al. "First Liquid Layer Inertial Confinement Fusion Implosions at the National Ignition Facility". *Physical Review Letters* 117 (2016), p. 245001.

[150] R E Olson et al. "Wetted foam liquid fuel ICF target experiments". *Journal of Physics: Conference Series* 717.1 (2016), p. 012042.

[151] Michael J. Rosenberg et al. "Direct-drive wetted-foam implosions on the NIF and OMEGA". *Bulletin of the American Physical Society* 70 (2025), PO08.001.

[152] R.E. Olson et al. "Physics of laser direct-drive liquid deuterium-tritium wetted foam capsules for inertial confinement fusion". *Nuclear Fusion* 66.2 (2025), p. 026002.

[153] S. Skupsky et al. "Polar direct drive on the National Ignition Facility". *Physics of Plasmas* 11.5 (2004), pp. 2763–2770.

[154] R. H. H. Scott et al. "Shock-Augmented Ignition Approach to Laser Inertial Fusion". *Physical Review Letters* 129 (2022), p. 195001.

[155] K. L. Baker et al. "Hydroscaling indirect-drive implosions on the National Ignition Facility". *Physics of Plasmas* 29.6 (2022), p. 062705.

[156] M H Key et al. "On point designs for high gain fast ignition". *Journal of Physics: Conference Series* 112.2 (2008), p. 022056.

[157] S. Atzeni, M. Temporal, and J.J. Honrubia. "A first analysis of fast ignition of precompressed ICF fuel by laser-accelerated protons". *Nuclear Fusion* 42.3 (2002), p. L1.

[158] S. Atzeni et al. "Fast ignitor target studies for the HiPER projecta)". *Physics of Plasmas* 15.5 (2008), p. 056311.

[159] R. Betti et al. "Gain curves for direct-drive fast ignition at densities around 300g/cc". *Physics of Plasmas* 13.10 (2006), p. 100703.

[160] Stefano Atzeni. “Inertial fusion fast ignitor: Igniting pulse parameter window vs the penetration depth of the heating particles and the density of the precompressed fuel”. *Physics of Plasmas* 6.8 (1999), pp. 3316–3326.

[161] D.S. Clark and M. Tabak. “A self-similar isochoric implosion for fast ignition”. *Nuclear Fusion* 47.9 (2007), p. 1147.

[162] John Nuckolls. “Grand Challenges of Inertial Fusion Energy”. *Journal of Physics: Conference Series* 244 (2010), p. 012007.

[163] Tomoyuki JOHZAKI, Kunioki MIMA, and Yasuyuki NAKAO. “Simulation Study of Ignition and Burn Characteristics of Fast Ignition DT Targets”. *Plasma and Fusion Research* 2 (2007), pp. 041–041.

[164] H. Nagatomo et al. “Formation of high areal density core using an efficient and robust implosion method for fast ignition”. *Nuclear Fusion* 64.10 (2024), p. 106019.

[165] Stephen Obenschain et al. “High-energy krypton fluoride lasers for inertial fusion”. *Appl. Opt.* 54.31 (2015), F103–F122.

[166] M. J. Rosenberg et al. “Stimulated Raman scattering mechanisms and scaling behavior in planar direct-drive experiments at the National Ignition Facility”. *Physics of Plasmas* 27.4 (2020), p. 042705.

[167] J. Myatt et al. “Modeling stimulated Brillouin scattering in the underdense corona of a direct drive inertial confinement fusion target”. *Physics of Plasmas* 11.7 (2004), pp. 3394–3403.

[168] W. Seka et al. “Two-plasmon-decay instability in direct-drive inertial confinement fusion experiments”. *Physics of Plasmas* 16.5 (2009), p. 052701.

[169] Konstantinos Papadopoulos, Melvyn L. Goldstein, and Robert A. Smith. “Stabilization of Electron Streams in Type III Solar Radio Bursts”. *The Astrophysical Journal* 190 (May 1974), pp. 175–186.

[170] Philip Chang et al. “The Effect of Nonlinear Landau Damping on Ultrarelativistic Beam Plasma Instabilities”. *The Astrophysical Journal* 797.2 (Dec. 2014), p. 110.

[171] G. Sarri et al. “Generation of neutral and high-density electron-positron pair plasmas in the laboratory”. *Nature Communications* 6 (2015).

[172] C. D. Arrowsmith et al. “Laboratory realization of relativistic pair-plasma beams”. *Nature Communications* 15.1 (2024), p. 5029.

[173] Erich S. Weibel. “Spontaneously Growing Transverse Waves in a Plasma Due to an Anisotropic Velocity Distribution”. *Physical Review Letters* 2.3 (Feb. 1959), pp. 83–84.

[174] S. Kathuria and R. Kalra. “Two-stream instability in a collisionless plasma”. *Physics of Fluids* 16.5 (1973), pp. 799–802.

[175] A. Bret, M. E. Dieckmann, and C. Deutsch. “Oblique electromagnetic instabilities for a hot relativistic beam interacting with a hot and magnetized plasma”. *Physics of Plasmas* 13.8 (Aug. 2006), p. 082109.

[176] Martin Lampe and P. Sprangle. “Saturation of the relativistic two-stream instability by electron trapping”. *The Physics of Fluids* 18.4 (Apr. 1975), pp. 475–481.

[177] M. J. Everett et al. “Coupling between electron plasma waves in laser-plasma interactions”. *Physics of Plasmas* 3.5 (May 1996), pp. 2041–2046.

[178] J. T. Mendonça et al. “Beam Instabilities in Laser-Plasma Interaction: Relevance to Preferential Ion Heating”. *Physical Review Letters* 94.24 (June 2005), p. 245002.

[179] N. Ratan et al. “Dense plasma heating by crossing relativistic electron beams”. *Physical Review E* 95.1 (Jan. 2017), p. 013211.

[180] O. Buneman. “Dissipation of Currents in Ionized Media”. *Physical Review* 115.3 (Aug. 1959), pp. 503–517.

[181] A. Bret, M.-C. Firpo, and C. Deutsch. “Characterization of the Initial Filamentation of a Relativistic Electron Beam Passing through a Plasma”. *Physical Review Letters* 94 (2005), p. 115002.

[182] L. D. Landau. “On the Vibrations of the Electronic Plasma”. *Journal of Physics (USSR)* 10 (1946), pp. 25–34.

[183] Ya. B. Fainberg, V. C. Shapiro, and V. I. Shevchenko. “Nonlinear Theory of Interaction between a “Monochromatic” Beam of Relativistic Electrons and a Plasma”. *Soviet Physics JETP* 30.3 (Mar. 1970). Translated from Zh. Eksp. Teor. Fiz. 57, 966-977 (1969), pp. 528–533.

[184] L. E. Thode and R. N. Sudan. “Two-Stream Instability Heating of Plasmas by Relativistic Electron Beams”. *Physical Review Letters* 30 (1973), pp. 732–735.

[185] Max Tabak et al. “Ignition and high gain with ultrapowerful lasers*”. *Physics of Plasmas* 1.5 (1994), pp. 1626–1634.

[186] M. Tabak et al. “Review of progress in Fast Ignition”. *Physics of Plasmas* 12.5 (2005), p. 057305.

[187] S. C. Wilks et al. “Absorption of ultra-intense laser pulses”. *Physical Review Letters* 69 (1992), pp. 1383–1386.

[188] F. Brunel. “Not-so-resonant, resonant absorption”. *Physical Review Letters* 59 (1987), pp. 52–55.

[189] F. N. Beg et al. “A study of picosecond laser–solid interactions up to 1019 W cm-2”. *Physics of Plasmas* 4.2 (1997), pp. 447–457.

[190] M. G. Haines et al. “Hot-Electron Temperature and Laser-Light Absorption in Fast Ignition”. *Physical Review Letters* 102 (2009), p. 045008.

[191] K. A. Tanaka et al. “Studies of ultra-intense laser plasma interactions for fast ignition”. *Physics of Plasmas* 7.5 (2000), pp. 2014–2022.

[192] R. Kodama et al. “Fast ignitor research at the Institute of Laser Engineering, Osaka University”. *Physics of Plasmas* 8.5 (2001), pp. 2268–2274.

[193] L. Volpe et al. “Controlling the fast electron divergence in a solid target with multiple laser pulses”. *Physical Review E* 90 (2014), p. 063108.

[194] J. S. Green et al. “Effect of Laser Intensity on Fast-Electron-Beam Divergence in Solid-Density Plasmas”. *Physical Review Letters* 100 (2008), p. 015003.

[195] R. A. Simpson et al. “Scaling of laser-driven electron and proton acceleration as a function of laser pulse duration, energy, and intensity in the multi-picosecond regime”. *Physics of Plasmas* 28.1 (2021), p. 013108.

[196] Kazuo A. Tanaka, Ryosuke Kodama, and Peter A. Norreys. “Integral Experiments for Fast Ignition Reserach”. *Fusion Science and Technology* 49.3 (2006), pp. 342–357.

[197] P. A. Norreys et al. “Experimental studies of the advanced fast ignitor scheme”. *Physics of Plasmas* 7.9 (2000), pp. 3721–3726.

[198] R. Kodama et al. “Fast heating of ultrahigh-density plasma as a step towards laser fusion ignition”. *Nature* 412.6849 (2001), pp. 798–802.

[199] P. A. Norreys et al. “Integrated implosion/heating studies for advanced fast ignition”. *Physics of Plasmas* 11.5 (2004), pp. 2746–2753.

[200] R. Kodama et al. “Study of Laser-Hole Boring into Overdense Plasmas”. *Physical Review Letters* 77 (1996), pp. 4906–4909.

[201] Marija Vranic, Ricardo A Fonseca, and Luis O Silva. “Extremely intense laser-based electron acceleration in a plasma channel”. *Plasma Physics and Controlled Fusion* 60.3 (2018), p. 034002.

[202] O. Culfa et al. “Hot electron production in laser solid interactions with a controlled pre-pulse”. *Physics of Plasmas* 21.4 (2014), p. 043106.

[203] Sophia Malko et al. “Enhanced relativistic-electron beam collimation using two consecutive laser pulses”. *Scientific Reports* 9 (2019).

[204] Robbie Scott et al. “Controlling fast-electron-beam divergence using two laser pulses.” *Physical review letters* 109 1 (2010), p. 015001.

[205] H Habara et al. “Efficient propagation of ultra-intense laser beam in dense plasma”. *Plasma Physics and Controlled Fusion* 57.6 (2015), p. 064005.

[206] B. T. Spiers et al. “Whole-beam self-focusing in fusion-relevant plasma”. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 379.2189 (2020), p. 20200159.

[207] Brian J. Duda and W. B. Mori. “Variational principle approach to short-pulse laser-plasma interactions in three dimensions”. *Physical Review E* 61 (2000), pp. 1925–1939.

[208] C. Ren and W. B. Mori. “Physical picture for the laser hosing instability in a plasma”. *Physics of Plasmas* 8.6 (2001), pp. 3118–3119.

[209] L Ceurvorst et al. “Mitigating the hosing instability in relativistic laser-plasma interactions”. *New Journal of Physics* 18.5 (2016), p. 053023.

[210] Nevill Francis Mott. “The scattering of fast electrons by atomic nuclei”. *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character* 124.794 (1929), pp. 425–442.

[211] Bernard Rosen. “Relativistic Coulomb Scattering”. *Journal of Mathematical Physics* 4.3 (1963), pp. 392–404.

[212] R. H. H. Scott et al. “A study of fast electron energy transport in relativistically intense laser-plasma interactions with large density scalelengths”. *Physics of Plasmas* 19.5 (2012), p. 053104.

[213] B. Ramakrishna et al. “Laser-Driven Fast Electron Collimation in Targets with Resistivity Boundary”. *Physical Review Letters* 105 (2010), p. 135001.

[214] D. J. Strozzi et al. "Fast-ignition transport studies: Realistic electron source, integrated particle-in-cell and hydrodynamic modeling, imposed magnetic fields". *Physics of Plasmas* 19.7 (2012), p. 072711.

[215] R.W. Paddock et al. "Pathways towards break even for low convergence ratio direct-drive inertial confinement fusion". *Journal of Plasma Physics* 88 (2022).

[216] R Paddock et al. "Energy gain of wetted-foam implosions with auxiliary heating for inertial fusion studies". *Plasma Physics and Controlled Fusion* 66 (2023).

[217] J. D. Lindl et al. "Progress toward a self-consistent set of 1D ignition capsule metrics in ICF". *Physics of Plasmas* 25.12 (2018), p. 122704.

[218] S. X. Hu et al. "Mitigating laser-imprint effects in direct-drive inertial confinement fusion implosions with an above-critical-density foam layer". *Physics of Plasmas* 25.8 (2018), p. 082710.

[219] H. Nishimura et al. "Indirect-direct hybrid target experiments with the GEKKO XII laser". *Nuclear Fusion* 40.3Y (2000), p. 547.

[220] B. Delorme et al. "Experimental demonstration of laser imprint reduction using underdense foams". *Physics of Plasmas* 23.4 (2016), p. 042701.

[221] S. Depierreux et al. "Laser Smoothing and Imprint Reduction with a Foam Layer in the Multikilojoule Regime". *Physical Review Letters* 102 (2009), p. 195005.

[222] A. S. Moore et al. "Foam-lined hohlraum, inertial confinement fusion experiments on the National Ignition Facility". *Physical Review E* 102 (2020), p. 051201.

[223] Ph. Nicolaï et al. "Experimental evidence of foam homogenization". *Physics of Plasmas* 19.11 (2012), p. 113105.

[224] Giora Hazak et al. "Shock propagation in a low-density foam filled with fluid". *Physics of Plasmas* 5.12 (1998), pp. 4357–4365.

[225] M. Cipriani et al. "Time-dependent measurement of high-power laser light reflection by low-Z foam plasma". *High Power Laser Science and Engineering* 9 (2021), e40.

[226] R. W. Paddock et al. "Measuring the principal Hugoniot of inertial-confinement-fusion-relevant TMPTA plastic foams". *Physical Review E* 107 (2023), p. 025206.

[227] K. Falk et al. "Temperature measurements of shocked silica aerogel foam". *Physical Review E* 90 (2014), p. 033107.

[228] Y. Aglitskiy et al. "Absolute Hugoniot measurements for CH foams in the 2–9 Mbar range". *Physics of Plasmas* 25.3 (2018), p. 032705.

[229] Michel Koenig et al. "Equation of state data experiments for plastic foams using smoothed laser beams". *Physics of Plasmas* 6.8 (1999), pp. 3296–3301.

[230] S. S. Kistler. "Coherent Expanded Aerogels and Jellies". *Nature* 127.3211 (1931), pp. 741–741.

[231] S. S. Kistler. "Coherent Expanded-Aerogels". *The Journal of Physical Chemistry* 36.1 (1932), pp. 52–64.

[232] L. M. Barker and R. E. Hollenbach. "Laser interferometer for measuring high velocities of any reflecting surface". *Journal of Applied Physics* 43.11 (1972), pp. 4669–4675.

[233] M. D. Knudson and M. P. Desjarlais. "Shock Compression of Quartz to 1.6 TPa: Redefining a Pressure Standard". *Physical Review Letters* 103 (2009), p. 225501.

[234] A. C. Mitchell and W. J. Nellis. "Shock compression of aluminum, copper, and tantalum". *Journal of Applied Physics* 52.5 (1981), pp. 3363–3374.

[235] D. G. Hicks et al. "Shock compression of quartz in the high-pressure fluid regime". *Physics of Plasmas* 12.8 (2005), p. 082702.

[236] R. M. More et al. "A new quotidian equation of state (QEOS) for hot dense matter". *The Physics of Fluids* 31.10 (1988), pp. 3059–3078.

[237] Sergei I Anisimov, A M Prokhorov, and Vladimir E Fortov. "Application of high-power lasers to study matter at ultrahigh pressures". *Soviet Physics Uspekhi* 27.3 (1984), p. 181.

[238] R. Cauble et al. "Absolute Equation-of-State Data in the 10–40 Mbar (1–4 TPa) Regime". *Physical Review Letters* 80 (1998), pp. 1248–1251.

[239] G. W. Collins et al. "Measurements of the Equation of State of Deuterium at the Fluid Insulator-Metal Transition". *Science* 281.5380 (1998), pp. 1178–1181.

[240] I K Krasyuk et al. "Study of extreme states of matter at high energy densities and high strain rates with powerful lasers". *Laser Physics* 26.9 (2016), p. 094001.

[241] S.A. Abrosimov et al. "Specific features of the behaviour of targets under negative pressures created by a picosecond laser pulse". *Quantum Electronics* 43.3 (2013), p. 246.

[242] Victor Malka et al. "Status and overview of research programs at LULI 2000". *ECLIM 2000: 26th European Conference on Laser Interaction with Matter*. Vol. 4424. International Society for Optics and Photonics. SPIE, 2001, pp. 70–77.

[243] Eugene. Hecht. *Optics*. Pearson, 2017.

[244] M. P. Desjarlais, M. D. Knudson, and K. R. Cochrane. "Extension of the Hugoniot and analytical release model of α-quartz to 0.2-3 TPa". *Journal of Applied Physics* 122.3 (2017), p. 035903.

[245] M. D. Knudson and M. P. Desjarlais. "Adiabatic release measurements in α-quartz between 300 and 1200 GPa: Characterization of α-quartz as a shock standard in the multimegabar regime". *Physical Review B* 88 (2013), p. 184107.

[246] Seth Root et al. "Shock Compression of a Fifth Period Element: Liquid Xenon to 840 GPa". *Physical Review Letters* 105 (2010), p. 085501.

[247] P. M. Celliers et al. "Line-imaging velocimeter for shock diagnostics at the OMEGA laser facility". *Review of Scientific Instruments* 75.11 (2004), pp. 4916–4929.

[248] D. G. Hicks et al. "Dissociation of Liquid Silica at High Pressures and Temperatures". *Physical Review Letters* 97 (2006), p. 025502.

[249] Makoto Tabata. "Transparent Silica Aerogel Blocks for High-Energy Physics Research". *Springer Handbook of Aerogels*. 2023, pp. 1333–1352.

[250] M. D. Knudson and R. W. Lemke. "Shock response of low-density silica aerogel in the multi-Mbar regime". *Journal of Applied Physics* 114.5 (2013), p. 053510.

[251] B. J. Henderson et al. "Shocked silica aerogel radiance transition". *Journal of Applied Physics* 132.9 (2022), p. 095902.

[252] B. Fryxell et al. "FLASH: An Adaptive Mesh Hydrodynamics Code for Modeling Astrophysical Thermonuclear Flashes". *The Astrophysical Journal Supplement Series* 131.1 (2000), p. 273.

[253] Matthew J. Herman et al. "Lithographic Printing Via Two-Photon Polymerization of Engineered Foams". *Fusion Science and Technology* 73.2 (2018), pp. 166–172.