A study of foam targets for direct-drive inertial confinement fusion



Robert W. Paddock

St Hilda's College

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Hilary Term 2023

Acknowledgements

I have worked with a large number of brilliant people over the last few years - far too many to list in full here. Many thanks to all of you: to everyone who I collaborated with, discussed work with, or bothered over email. It was great working with you all, and I hope to continue to do so going forward! Particular thanks to everyone who worked through the uncertainty of the pandemic and overcame constantly-changing travel restrictions to take part in the Vulcan experiment. Thanks also to the Norreys' group past and present (particularly Ben, Ramy and Marko), and to the wider ALP sub-department.

A special mention is of course due to Peter Norreys, who I had the good fortune to have as my supervisor. Thank you Peter: for all the guidance, for the constant ideas, and for introducing me to so many great people. Thanks for accommodating my interests and giving me the freedom to pursue them, whether that was finding a way to do some experimental work, or allowing me to disappear off to parliament for three months. Most of all, thanks for the constant support and encouragement! It's been a pleasure to be your student.

Finally, a huge thanks to my family and friends, for all the love and support throughout the years.

Funding

This work was funded by UKRI-EPSRC and AWE-plc through an industrial CASE scholarship. I would like to thank both institutions for their financial support.

Abstract

Recent experiments using DT-wetted-foam layers in inertial confinement fusion capsules have demonstrated that the convergence of the implosion can be controlled through varying the initial capsule temperature. These experiments also demonstrated low amounts of hydrodynamic instability growth at low convergence ratio. In this thesis, simulations and an experiment relevant to these capsules are presented.

Simulation campaigns are discussed which explore the fusion performance that could be achieved in experiments using wetted-foam capsules, in a newly identified regime (based around limits to convergence ratio and other implosion parameters) where instability growth is expected to be minimal. Performance is explored first for third-harmonic of Nd:glass laser drivers, and then for higher frequency ArF drivers and a novel 'two-colour' laser scheme. Further simulations are then performed which explore how such capsules could be heated using electron beams, and the impact this has on fusion performance is discussed. 'Hydrodynamic-equivalent' capsules are also presented, which are surrogate targets where the DT-wetted-foam layer is replaced instead with an equivalent density layer of dry foam. Such capsules have comparable hydrodynamic performance, and could facilitate room temperature experiments to validate the low instability regime.

The equation of state of these foams is not well characterised. An experiment is therefore also presented which measured principal Hugoniot data of TMPTA foam at 260 $\mathrm{mg\,cm^{-3}}$, equivalent to the foam layer in the hydrodynamic-equivalent capsules. VISAR was used to measure shock transit times through an α-quartz reference layer and a TMPTA foam layer in a multi-layer target and impedance matching performed to enable calculation of the foam shock state, while the foam shock temperature was estimated using streaked optical pyrometry. The results demonstrate that for the pressure range probed in the experiment the TMPTA foam was well described by existing QEOS and SESAME models, suggesting that this foam can be well approximated as a low density homogenous plastic.

© Crown Owned Copyright/AWE [2023]

Contents

1 Introduction **1**

1.1 Introduction 1
1.2 Structure of the thesis 4
1.3 Author's contributions 5
1.4 List of peer-reviewed publications 7
1.4.1 First author publications covering the work in this thesis . . 7
1.4.2 Other publications over this time period 7
1.5 List of conference presentations 8

2 Theory **9**

2.1 Introduction 10
2.2 Plasmas and laser-plasma interactions 10
2.2.1 The Debye length 11
2.2.2 Plasma oscillations 11
2.3 Mathematical formulations of plasmas 13
2.3.1 The plasma distribution function 13
2.3.2 The Boltzmann and Vlasov equations 14
2.3.3 The fluid equations 15
2.4 Magnetohydrodynamics 17
2.5 Additional physics in fluid codes 19
2.6 Basic description of laser-driven inertial confinement fusion 26
2.6.1 Alternative ignition schemes 29
2.6.2 Polar direct drive 30
2.7 Instability growth 30
2.7.1 Hydrodynamic instabilities 31
2.7.2 Parametric instabilities 32
2.8 Shocks and impedance matching 34
2.8.1 Shock physics 34
2.8.2 Rankine-Hugoniot relations in a moving frame 36
2.8.3 The Hugoniot 38
2.8.4 Material interfaces and impedance matching 39
2.8.5 Measuring shock states experimentally 42

3 Definitions of key terms **44**
3.1 Introduction . . . 44
3.2 Convergence ratio . . . 45
3.3 Hotspot and shell . . . 46
3.4 Hotspot / shell temperature . . . 47
3.5 Areal density . . . 47
3.6 In-flight aspect ratio . . . 48
3.7 Implosion velocity . . . 49
3.8 Laser intensity . . . 49
3.9 Yield, gain and bang time . . . 49
3.10 Adiabat . . . 50
3.11 Ignition and break-even . . . 51

4 Low convergence ratio implosions **53**
4.1 Introduction . . . 53
4.2 Previous research into wetted-foam capsules and low convergence ratio implosions . . . 55
4.2.1 Experimental results . . . 55
4.2.2 Simulation work . . . 56
4.2.3 Implications of these previous findings . . . 58
4.3 Defining a 'low instability' regime . . . 59
4.3.1 The four 'low instability' criteria . . . 63
4.4 Capsule design and laser pulse-profile . . . 64
4.5 Performing the simulation campaign . . . 65
4.5.1 The simulation code `HYADES` . . . 65
4.5.2 Capsule meshing . . . 66
4.5.3 Conducting the simulation campaign . . . 67
4.5.4 Hydroscaling . . . 69
4.6 Results and analysis . . . 69
4.7 Limitations and suggested future work . . . 73
4.8 Conclusion . . . 75

5 Further simulation campaigns in the low instability regime **76**
5.1 Introduction . . . 77
5.1.1 An alternative fusion metric: the burning plasma parameter 78
5.2 Estimating the impact of the wetted-foam EOS . . . 79
5.3 Hydrodynamic-equivalent capsules . . . 82
5.4 Alternative laser drivers . . . 84
5.4.1 Background to possible IFE laser drivers . . . 85
5.4.2 ArF lasers . . . 87

5.4.3 Two-colour implosions . . . 88
5.4.4 The burning plasma parameter for alternative laser drivers . 92
5.4.5 The zoomed focus technique . . . 93
5.5 Auxiliary heating using relativistic electron beams . . . 94
5.5.1 Principles of the auxiliary heating scheme . . . 95
5.5.2 Simulating auxiliary heating in `HYADES` . . . 96
5.5.3 Timing of electron energy deposition . . . 98
5.5.4 Magnitude of deposited electron energy . . . 100
5.5.5 The burning plasma parameter for heated capsules . . . 102
5.5.6 Estimating the gain of heated capsules . . . 102
5.5.7 Comparison with previous simulated and experimental results 105
5.5.8 Auxiliary heating of high gain capsules . . . 107
5.5.9 Optimisation and heating of high areal-density capsules . . . 109
5.5.10 Limitations and suggested future direction . . . 112
5.5.11 Subsequent developments of the auxiliary heating technique 113
5.6 Conclusions . . . 114

6 Measuring the principal Hugoniot of TMPTA 116
6.1 Introduction . . . 117
6.2 Experiment overview and diagnostics . . . 118
6.2.1 The Vulcan laser . . . 119
6.2.2 VISAR . . . 121
6.2.3 Streaked Optical Pyrometry . . . 130
6.2.4 X-ray pinhole camera . . . 133
6.2.5 Fiducial . . . 134
6.2.6 Calorimetry . . . 135
6.2.7 Pulse length . . . 135
6.3 Preparation for the experiment . . . 135
6.3.1 Development of the optical set-up . . . 135
6.3.2 Target design and hydrodynamic simulations . . . 140
6.3.3 Hydrodynamic simulations . . . 144
6.4 Conducting the experiment . . . 149
6.4.1 Changes from planned target . . . 150
6.4.2 Poor VISAR illumination and subsequent changes to optical set-up . . . 151
6.4.3 Later than expected shock breakout times . . . 152
6.4.4 Signal strength issues . . . 154
6.5 Conclusion . . . 155

7 Measuring the principal Hugoniot of TMPTA: analysis and interpretation 156
7.1 Introduction . . . 157
7.2 Analysis of VISAR data . . . 158
7.2.1 Identifying shock velocities from raw VISAR data . . . 158
7.2.2 Impedance matching calculation . . . 160
7.2.3 Uncertainty quantification - Monte Carlo analysis . . . 163
7.2.4 Achieved quartz shock states . . . 167
7.2.5 Calculation of velocity from VISAR fringe curvature . . . 168
7.3 Analysis of SOP data . . . 171
7.3.1 Calculating grey-body temperature from SOP data . . . 171
7.3.2 Comparing SOP timings to VISAR timings . . . 174
7.4 Analysis from additional diagnostics . . . 175
7.4.1 Using the fiducial to estimate laser turn-on time . . . 176
7.4.2 Using the pinhole camera to estimate spot size . . . 177
7.4.3 Measuring laser pulse length using the fiducial and photodiode traces . . . 180
7.4.4 Estimating pulse intensity and comparison with shock strength 181
7.5 Shot selection for final data set . . . 183
7.5.1 Possible evidence of second shocks . . . 183
7.5.2 Confidence ranking of shots . . . 184
7.6 Results and discussion . . . 185
7.7 Comparison with simulation . . . 189
7.7.1 Shock propagation time . . . 191
7.7.2 Shock stability . . . 192
7.7.3 Shock merger . . . 194
7.8 Investigation of the shock strength discrepancy . . . 196
7.8.1 Analysis of non-data targets . . . 196
7.8.2 Proposed explanation for the discrepancy . . . 197
7.8.3 Simulation of a gap between layers . . . 198
7.9 Suggested improvements . . . 199
7.9.1 Target improvements . . . 199
7.9.2 Equipment . . . 200
7.9.3 Experimental layout . . . 201
7.10 Conclusion . . . 201

8 Summary and future work 204
8.1 Summary . . . 204
8.2 Future work . . . 205

Appendices

A Benchmarking HYADES **208**
A.1 Direct-drive OMEGA shots 208
A.2 Polar direct-drive NIF shot 211
A.3 Conclusions 213

B Key code **215**
B.1 Capsule meshing 215
B.1.1 Theory 216
B.2 Ray tracing 220
B.3 VISAR analysis 221

C Experiment photos and equipment list **223**
C.1 Final experimental results 223
C.2 Equipment list 225
C.3 Experimental set-up images 226
C.4 Target images 233

References **235**

1 Introduction

Contents

1.1 Introduction **1**
1.2 Structure of the thesis **4**
1.3 Author's contributions **5**
1.4 List of peer-reviewed publications **7**
1.4.1 First author publications covering the work in this thesis 7
1.4.2 Other publications over this time period 7
1.5 List of conference presentations **8**

1.1 Introduction

Nuclear fusion is the process through which small nuclei (typically isotopes of hydrogen) fuse together to form larger nuclei, with an associated release of energy. This basic process is the energy source of the stars, and the focus of a significant research effort which hopes to harness it for electricity production. Yet despite decades of effort, the goal of a fusion reactor has so far remained out of reach.

There are a variety of approaches to fusion being studied in laboratories around the world, but these fall broadly into two main categories. In magnetic confinement fusion (MCF), a reactor (typically a 'tokomak', but alternative designs such as 'stellarators' are also used) uses strong magnetic fields to confine a plasma as it is

heated to extreme temperatures. Fusion reactions will occur in the hot plasma, and if it can be confined for long enough timescales these will occur in such numbers that the fusion energy produced is larger than the energy required to power the reactor, resulting in net energy production. The second approach, inertial confinement fusion (ICF), is the focus of this thesis. In ICF, a capsule containing hydrogen fuel is compressed to extreme temperatures and densities - the highest ever achieved on earth. Such conditions cannot be sustained; as the capsule implodes the pressure inside quickly rises, and this will eventually overcome the drive pressure and cause the capsule to explode outwards (the fuel is 'confined' only briefly, by its own inertia). However, if the temperature and density are high enough in the brief period where the fuel is maximally compressed, large numbers of fusion reactions can be generated and the energy produced by these reactions can be larger than that required to drive the implosion.

The compression in ICF can be achieved using a variety of approaches. Examples include 'Z-pinches' (where extreme currents passed through a plasma lead to high magnetic fields through the Lorentz force, which cause the plasma to contract), impact ignition (where a fast moving projectile impacts on the target), or laser ICF, where the compression is provided via MJ-level pulses from high power lasers. This thesis will investigate the laser approach to ICF, and unless otherwise specified it can be assumed that the use of 'ICF' throughout is referring to this technique. In 'direct-drive' (the main focus of this thesis) the laser pulses irradiate the capsule directly, while in 'indirect-drive' they are incident instead on the inside of a hohlraum, producing a bath of x-rays which drive the capsule compression. The end result is the same; the laser pulse/x-rays ablate material from the capsule, accelerating the rest of the capsule inwards and compressing the fuel inside. This drives high temperatures and densities, ionising the fuel and turning it into a plasma.

The highest energy facility performing ICF implosions using laser compression is the National Ignition Facility (NIF), at Lawrence Livermore National Laboratory in California. At the NIF, 192 beams provide up to 2.1 MJ of laser energy in a polar beam configuration designed for indirect-drive experiments. The NIF has

now been in operation for just over a decade, during which it has demonstrated a continual improvement in achieved fusion performance. This recently resulted in a shot in December 2022 where it achieved break-even (where the fusion energy output was larger than the laser energy input) for the first time. This is a momentous achievement, and represents the first time net energy production has been achieved in any laboratory fusion experiment (MCF, ICF, or otherwise). This has sparked fresh excitement about ICF, and makes it an exciting time to be working in the field.

Despite this, there remains significant work to be done on this topic. Firstly, for a successful inertial fusion energy (IFE) power plant, a much higher fusion output (a 'gain', or ratio of energy out to energy in, of 50-100) is required [1]; this is necessary to cover the energy use in generating the laser beams and inside the power plant, and to make the whole operation financially viable. The cost of such a facility would be large, and a range of analyses have been performed to estimate the cost per unit energy required to make such power stations competitive [2, 3]. The reactor would also need to operate at high repetition rate, on the order of multiple shots a second, as opposed to the roughly one shot a day that the NIF can currently provide. Along with improvements to the laser technology, and the design of a reactor that could deliver targets at this rate and deal with the high neutron flux, this would require a supply of targets that can be produced quickly and to the required levels of precision at low cost [4].

Conventional ICF targets (including those used in the recent ignition shot) tend to use a frozen layer of deuterium and tritium (DT) as the main bulk of the fuel. This DT ice surrounds a central region containing DT vapour, and is in turn surrounded by a shell of ablator material. Such designs enabled impressive performance, but each target is expensive and time-consuming to produce [5]. 'Wetted-foam' or 'liquid-layer' targets are being explored as a potential alternative. Rather than use an ice-layer, these targets contain a low density carbon foam saturated with liquid DT, where the liquid is held in the foam by capillary forces. Such targets could solve some of the manufacturing challenges associated with ice-layers, and can be 3D printed with the potential for rapid mass-production [6].

The interest in wetted-foam capsules is not solely based on these practical reasons, as they also have a range of potential advantages for fusion performance. Of particular interest for this thesis is the opportunity they provide to control capsule convergence. Ice layer targets must be fielded at temperatures lower than 19 K, in order to ensure that the DT remains frozen. Wetted-foam capsules however are fielded at higher temperatures between 20 - 26 K (corresponding to the range of temperatures over which DT is in liquid form) [7]. Increasing this temperature leads to a higher density of the DT vapour in the central void of the capsule (between 0.6 and 4.0 $\mathrm{mg\,cm^{-3}}$), and this higher density leads to reduced compression and convergence. By varying the initial temperature of wetted-foam capsules, the degree of convergence experienced in the implosion can be controlled and explored; this is of interest, since low convergence implosions are known to be more robust to hydrodynamic instability growth.

1.2 Structure of the thesis

This thesis investigates the potential performance of direct-drive ICF implosions of wetted-foam capsules, along with studying the compression behaviour of foams relevant to these designs. It is split into the following chapters:

- Chapter 2 will describe the key theory of plasmas, ICF, and shock propagation through materials.
- Chapter 3 will define a number of key terms and implosion parameters used throughout this thesis, and discuss how exactly they are evaluated for the simulations presented in later chapters.
- Chapter 4 will discuss the role of low convergence ratio in minimising instability growth, and present a simulation campaign investigating the performance of wetted-foam capsules in an expected low instability regime.
- Chapter 5 will present further simulations extending the results of the previous chapter. It will investigate the impact of the equation of state model used,

designs for surrogate targets to enable room temperature experiments based on these designs, and then include further simulation campaigns investigating the performance that can be obtained when alternative laser drivers or auxiliary heating techniques are applied within the low instability regime.

- The next two chapters will focus on experimental work investigating the equation of state of a trimethylolpropane triacrylate (TMPTA) plastic foam relevant to these previous designs. Chapter 6 will describe the principles, diagnostics, and preparation work performed for the experiment, while Chapter 7 will focus on the data analysis and interpretation of the results.

- Chapter 8 will then summarise the thesis. There are also three appendices, which will focus on the operation of key code for this work, benchmarking of the simulations, and further details of the experimental set-up.

1.3 Author's contributions

The work in this thesis is largely my own, but there are places where I have included the work of collaborators, or the results of students who I supervised. I have indicated this wherever it has occurred, but also provide a brief summary here of these contributions.

Throughout the thesis, I relied upon the guidance and assistance of Dr Robbie Scott when running simulations in `HYADES/h2D`. In particular, the physics models used (such as for ionisation, radiation transport etc.) were based on his recommendations. The benchmarking of `HYADES` presented in Appendix A was performed using data provided by Dr Alex Zylstra and Dr Brian Haines, who also provided advice during this process.

While performing this work I provided day-to-day supervision of a number of summer and master's students. These students helped to perform a number of the simulation campaigns discussed in this thesis, using the code and methodology I had created for this purpose and under my instruction; this allowed us to explore a wider range of implosions than would otherwise have been possible. They included

Heath Martin and Rusko Ruskov (who contributed to the simulations of alternative laser drivers), along with Tat Li and Eugene Kim (who applied auxiliary heating to some of these alternative laser driver results, and in Tat's case also optimised two implosions for maximum areal density).

At the end of the section on auxiliary heating, I have provided a brief summary of work performed on this topic since the simulations presented in that section were performed. While I contributed to this (and was co-author on the resulting paper), the main contributions were made by Jordan Lee, Heath Martin and Rusko Ruskov. A brief discussion of these results is included for the sake of completeness, and in order to provide an up-to-date and accurate description of the heating process.

The experiment was a large collaboration, and many groups and individuals contributed to its development, operation and analysis. I performed the work presented unless stated otherwise, but was supported throughout by input from our collaborators during the proposal writing, through regular planning meetings in the year before the experiment, and in analysis meetings afterwards. A large number of collaborators attended the experiment itself (including Matthew Oliver, who helped out a great deal in his role as link scientist, and Paul Mabey, who acted as target area operator for two weeks). I have also included a few results generated by these collaborators in the thesis. The VISAR was produced by Professor Daniel Eakins, and he provided the analysis code which was used to calculate shock velocity from VISAR fringe motion. Michael Woodward produced all the CAD models presented throughout the experiment chapters. Rati Goshadze, Dr Valentin Karasiev, and Professor Suxing Hu performed the DFT simulations calculating the foam reflectivity at a range of pressures, which allowed me to calculate the grey-body foam shock temperatures. Finally, Professor Piotr Rączka, Dr Paul Neumayer, and Dr Artem Martynenko performed pre- and post-experiment simulations to support my own simulations; the results of some of these are included in Figure 7.24, and some of Piotr's 2D `FLASH` simulations are also presented in Figures 7.23 and 7.27

1.4 List of peer-reviewed publications

1.4.1 First author publications covering the work in this thesis

- 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).
- R. W. Paddock et al., "Pathways towards break even for low convergence ratio direct-drive inertial confinement fusion", Journal of Plasma Physics, **88**, 905880314 (2022). **(Featured article)**
- R. W. Paddock et al., "Measuring the principal Hugoniot of inertial-confinement-fusion-relevant TMPTA plastic foams", Physical Review E, **107**, 025206 (2023).

1.4.2 Other publications over this time period

- J. J. Lee et al., "Towards more robust ignition of inertial fusion targets", Physics of Plasmas, **30**, 022702 (2023).
- M. W. Von der Leyen, et al., "Observation of monoenergetic electrons from two-pulse ionization injection in quasilinear laser wakefields", Physical Review Letters, **130**, 105002 (2023).
- J. J. Lee et al., "Multi-Scale Simulations of Hot Electron Transport in Overdense Plasma", Scientific Reports, Under review.
- R. Timmis et al., "Attosecond duration nano-Coulomb electron bunches from intense laser-plasma interactions", Physical Review E., Under review.
- I. Ouatu et al., "Ionization states for the multipetawatt laser-QED regime", Physical Review E., **106**, 015205 (2022).
- Q. S. Feng et al., "Suprathermal electrons from the anti-Stokes Langmuir decay instability cascade", Physical Review E., **105**, 045208 (2022).
- P. A. Norreys et al., "Preparations for a European R&D roadmap for an inertial fusion demo reactor", Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences, **379**, 20200005 (2021).

- 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**, 20200159 (2021).
- B. T. Spiers et al., "Methods for extremely sparse-angle proton tomography", Physical Review E, **104**, 045201 (2021).
- M. W. Mayr et al. "Nonlinear wakefields and electron injection in cluster plasma", Physical Review Accelerators and Beams, **23**, 093501 (2020).

1.5 List of conference presentations

- Poster at 'Prospects for high gain inertial fusion energy' discussion meeting at the Royal Society, 02 - 03 March 2020.
- Invited AWE plasma physics colloquium seminar, 07 January 2021.
- Oral presentation at AWE Physics Student Conference 2021, 14-15 April 2021.
- Poster at EPS Conference on Plasma Physics, 21-25 June 2021.
- Oral presentation at APS DPP 2021, 8-12 November 2021.
- Oral presentation at AI4ICF, Oxford, 6-7 December 2021.
- Oral presentation at OxCHEDS, 21-22 March 2022
- Oral presentation at AWE Materials and Analytical Student Conference 2022, 4th-5th April 2022.
- Oral presentation at DDFIW 2022, 3-5 May 2022.
- Oral presentation at ECLIM 2022, 19-23 September 2022.
- Invited ALP seminar at Oxford, 10 October 2022.
- Oral presentation at APS DPP 2022, 17-21 October 2022.
- Invited CFSA seminar at Warwick, 1 November 2022.
- Poster and oral presentation at HPL 2022, 19-21 December 2022.
- Poster at NIF user group 2023, 21-23 February 2023.
- Oral presentation at OxCHEDS, 20-21 March 2023.
- Oral presentation at IOP Annual Plasma Physics Conference, 27-30 March 2023.
- Oral presentation at AWE Physics Student Conference 2023, 17-18 April 2023.

2 Theory

Contents

2.1 Introduction 10
2.2 Plasmas and laser-plasma interactions 10
2.2.1 The Debye length 11
2.2.2 Plasma oscillations 11
2.3 Mathematical formulations of plasmas 13
2.3.1 The plasma distribution function 13
2.3.2 The Boltzmann and Vlasov equations 14
2.3.3 The fluid equations 15
2.4 Magnetohydrodynamics 17
2.5 Additional physics in fluid codes 19
2.6 Basic description of laser-driven inertial confinement fusion 26
2.6.1 Alternative ignition schemes 29
2.6.2 Polar direct drive 30
2.7 Instability growth 30
2.7.1 Hydrodynamic instabilities 31
2.7.2 Parametric instabilities 32
2.8 Shocks and impedance matching 34
2.8.1 Shock physics 34
2.8.2 Rankine-Hugoniot relations in a moving frame 36
2.8.3 The Hugoniot 38
2.8.4 Material interfaces and impedance matching 39
2.8.5 Measuring shock states experimentally 42

2.1 Introduction

This section briefly outlines the key theory relevant to the later chapters of this thesis. First, the concept of 'plasma' is defined, and some of the key behaviours of plasmas are described. The kinetic description of plasmas is then presented, along with the fluid approximations used to make long time-scale simulations viable, followed by further discussion of the additional physics required to simulate plasmas and ICF implosions effectively. A brief overview of inertial confinement fusion is then provided, along with a description of the instabilities that are experienced during such experiments. Finally, the last section of this chapter discusses shock waves and how they interact with material interfaces (which is central to the experiment presented in the latter half of this thesis).

2.2 Plasmas and laser-plasma interactions

The definition of a plasma, as provided in the textbook 'Introduction to Plasma Physics and Controlled Fusion' by Francis Chen [8], is 'a quasi-neutral gas of charged and neutral particles that exhibit collective behaviour'. This definition can be split into three key parts. The first, from the mention of 'charged particles', is that ionisation is essential to plasma physics; however, this definition also makes it clear that ionisation alone is not sufficient for the creation of a plasma. The second component of this definition is the concept of 'quasi-neutrality'. This means that concentrations of charge or external electric fields are effectively shielded by the electrons in a plasma, such that it can be considered electrically neutral over large length-scales. The final component of this definition is the concept of 'collective behaviour', which highlights that motion in plasmas is not simply local. This is due to the fact that the plasma consists of charged particles, the motion of which leads to electromagnetic forces which result in motions throughout the entire plasma. This section discusses some key fundamental concepts of plasma physics, which in turn explores the concepts of quasi-neutrality and collective

behaviour in further detail. This discussion follows that found in standard plasma physics textbooks, such as [8] and [9].

2.2.1 The Debye length

The Debye length λ_D is a characteristic length scale in a plasma, defined by

$$\lambda_D = \sqrt{\frac{\epsilon_0 k_B T_e}{n e^2}} \,, \tag{2.1}$$

where T_e is the electron temperature, n is the number density of ions/electrons, e is the electron charge, ϵ_0 is the permittivity of free space, and k_B is Boltzmann's constant. The Debye length defines the (exponential decay) length-scale over which charges are 'shielded' by electrons within the plasma. Given that it is this electron shielding of charges which is responsible for quasi-neutrality, an approximate mathematical condition for a plasma can be written as

$$N_D = n \cdot \frac{4}{3}\pi\lambda_D^3 \gg 1 \,, \tag{2.2}$$

where N_D is the number of particles in a 'Debye sphere', i.e. the number of electrons that would be within the Debye length of a charge, and thus available to shield it. Clearly for the charge to be effectively shielded and thus quasi-neutrality to be established, this number should be much larger than one.

2.2.2 Plasma oscillations

If the electrons in the plasma are displaced slightly from the ions[1], there will be a restoring force generated. This will accelerate the electrons towards the ions, which they will then overshoot; this leads to oscillations of the electrons around the ion background. Such oscillations can be derived using the continuity equation and equation of motion for the electrons[2] and Poisson's equation,

$$m n_e [\frac{\partial \vec{v_e}}{\partial t} + (\vec{v_e} \cdot \boldsymbol{\nabla})\vec{v_e}] = -e n_e \vec{E} \,, \tag{2.3}$$

[1]Given the mismatch between the ion and electron mass in a typical plasma, when considering short time scales the ions are often considered to be stationary compared to the highly mobile electrons.

[2]These are derived from kinetic theory in Section 2.3.3, or can be derived by considering fluid motion as in [8] and [10].

$$\frac{\partial n_e}{\partial t} + \boldsymbol{\nabla} \cdot (n_e \vec{v_e}) = 0 \ , \tag{2.4}$$

$$\epsilon_0 \boldsymbol{\nabla} \cdot \vec{E} = e(n_i - n_e) \ , \tag{2.5}$$

where $\vec{v}$ is the velocity, m is the mass, $\vec{E}$ is the electric field, and the subscripts e and i refers to the electrons and ions respectively. This set of equations is solved through linearisation, and results in sinusoidal oscillations of the electrons at the plasma frequency,

$$\omega_p = \sqrt{\frac{n_0 e^2}{\epsilon_0 m}} \ . \tag{2.6}$$

If a pressure term ($\boldsymbol{\nabla} \cdot P_e = 3k_B T_e \boldsymbol{\nabla} \cdot n_e$) is added to the electron equation of motion (equation 2.3), then thermal motion of the electrons causes the plasma oscillation to propagate as a plasma wave, with a dispersion relation described by

$$\omega^2 = \omega_p^2 + \frac{3}{2} k^2 v_{\text{th}}^2 \ , \tag{2.7}$$

where k is the wave vector and $v_{\text{th}} = \frac{2k_B T_e}{m}$.

Such waves are referred to as plasma waves, electron waves, Langmuir waves, or plasmons. The existence of such waves serves to illustrate the collective behaviour that defines plasmas. Another condition for a plasma can be defined through $\omega_p \tau \gg 1$, where τ is the collision time between charged and neutral ions - such a condition essentially expresses that plasma oscillations must occur faster than collisions with neutral atoms, such that electromagnetic forces govern the motion rather than normal hydrodynamic forces.

Similar analyses are performed to derive a large range of possible oscillations in plasmas, in response to applied fields[3] and electromagnetic waves. Two significant results are highlighted. First, when the derivation is performed for ions (rather than electrons), the 'ion-acoustic wave' is obtained. This is equivalent to a sound wave in a plasma, and has the dispersion relation

$$\omega = k v_s \ , \tag{2.8}$$

[3]While plasmas effectively shield electric fields, externally applied magnetic fields are not shielded and lead to a large number of potential plasma modes.

where v_s is the plasma sound speed. Secondly, in the absence of external fields the dispersion relation for an electromagnetic wave in a plasma is given by

$$\omega^2 = \omega_p^2 + c^2 k^2 \ . \tag{2.9}$$

The form of this relation is such that in this case, EM waves can only propagate if their frequency is higher than the plasma frequency. As $\omega_p \propto \sqrt{n}$, an EM wave travelling through plasma of increasing density eventually encounters a 'critical density' past which it cannot propagate, thus resulting in an evanescent wave that decays exponentially in strength past this point over the 'skin depth' of the plasma.

2.3 Mathematical formulations of plasmas

A detailed mathematical description of plasmas is given by kinetic theory. In this section, a basic introduction to this kinetic description is provided, before showing how simpler fluid approximations can be made. This eventually results in the equations for ideal magnetohydrodynamics, which allow plasmas to be simulated as fluids over large time and length scales. The discussion in this section follows those presented in [8, 10, 11].

2.3.1 The plasma distribution function

The state of a plasma can be completely described using its *velocity distribution function* f. This distribution function describes the particle populations within the plasma in a 6-dimensional phase space as a function of time,

$$f = f(\vec{r}, \vec{v}, t) \ , \tag{2.10}$$

where $\vec{r}$ and $\vec{v}$ are the position and velocity of each particle.

The number of particles at a position $\vec{r}$ and with velocities in the infinitesimally narrow velocity range $d\vec{v}$ is thus described by $f(\vec{r}, \vec{v}, t)d\vec{v}$, and the number density $n(\vec{r}, t)$ of the plasma can be obtained via integrating over the full velocity phase space,

$$n(\vec{r}, t) = \int_{-\infty}^{\infty} f(\vec{r}, \vec{v}, t) \, \mathrm{d}\vec{v} = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f(\vec{r}, \vec{v}, t) \, \mathrm{d}v_x \mathrm{d}v_y \mathrm{d}v_z \ . \tag{2.11}$$

2.3.2 The Boltzmann and Vlasov equations

The evolution of the distribution function with time is now considered. The 'rate of change' of a plasma is described from two perspectives; either through measuring the change in the plasma at a fixed point, or through measuring the change experienced by a 'fluid element'. This is expressed through the *convective derivative*,

$$\left(\frac{d}{dt}\right) = \left(\frac{\partial}{\partial t} + \vec{v}\cdot\boldsymbol{\nabla}\right) \ . \tag{2.12}$$

Equations describing changes to the plasma are expressed either in Lagrangian form (considering changes to a moving plasma 'packet', and so using the single derivative $\frac{d}{dt}$) or Eulerian form (considering a stationary point, and using $\frac{\partial}{\partial t} + \vec{v}\cdot\boldsymbol{\nabla}$ for the derivative). In this thesis, the Eulerian forms of the equations are used.

The evolution of the distribution function in the presence of an external force F is governed by the *Boltzmann equation*,

$$\frac{\partial f}{\partial t} + \vec{v}\cdot\boldsymbol{\nabla}_x f + \frac{\vec{F}}{m}\cdot\boldsymbol{\nabla}_v f = 0 \ , \tag{2.13}$$

where Newton's second law $\frac{\vec{F}}{m} = \frac{d\vec{v}}{dt}$ provides the acceleration of the particles due to this force. The Boltzmann equation in this form does not include collisions.

Next, consider the case where the only forces acting are electromagnetic. Then, the Boltzmann equation becomes known as the *Vlasov* equation, and takes the form

$$\frac{\partial f}{\partial t} + \vec{v}\cdot\boldsymbol{\nabla}_x f + \frac{q}{m}(\vec{E} + \vec{v}\times\vec{B})\cdot\boldsymbol{\nabla}_v f = 0 \ , \tag{2.14}$$

where the force is now the Lorentz force, and $\vec{B}$ is the magnetic field.

Collisions are included in both the Boltzmann and Vlasov equations through the addition of a collision operator $\left(\frac{\partial f}{\partial t}\right)_c$. Many forms for such an operator exist; a common choice is the 'Fokker-Planck' operator, which describes binary Coulomb collisions. When this operator is added to the Vlasov equation, it forms the *Vlasov-Fokker-Planck (VFP) equation*,

$$\frac{\partial f}{\partial t} + \vec{v}\cdot\boldsymbol{\nabla}_x f + \frac{q}{m}(\vec{E} + \vec{v}\times\vec{B})\cdot\boldsymbol{\nabla}_v f = \left(\frac{\partial f}{\partial t}\right)_{\text{Fokker}-\text{Planck}} \ . \tag{2.15}$$

The Vlasov equation describes the full kinetic evolution of the plasma, and can be solved in Vlasov codes to return the distribution function as a function of time. This is extremely useful for simulating kinetic effects, but is computationally expensive. It is not feasible to perform such simulations over the length and time scales required for a full ICF implosion. To simulate such implosions, it is therefore necessary to make approximations to reduce the computational cost.

2.3.3 The fluid equations

Fluid variables, which vary as a function of $\vec{r}$ and t only, can be obtained by taking moments of velocity of the distribution function (i.e. performing the integral $\int_{-\infty}^{\infty} \vec{v}^i f d\vec{v}$, where i represents the moment being taken). The first three moments return the particle density n, fluid velocity $\vec{u}$ (or the momentum $\vec{p}$), and pressure tensor $\underline{\underline{P}}$:

$$n = \int_{-\infty}^{\infty} f \, \mathrm{d}\vec{v} \,, \tag{2.16}$$

$$\vec{p} = mn\vec{u} = m \int_{-\infty}^{\infty} \vec{v} f \, \mathrm{d}\vec{v} \,, \tag{2.17}$$

$$\underline{\underline{P}} = m \int_{-\infty}^{\infty} (\vec{v} - \vec{u})(\vec{v} - \vec{u}) f \, \mathrm{d}\vec{v} \,. \tag{2.18}$$

Fluid equations governing the evolution of each of these variables are also derived[4] by taking the equivalent moment of the Vlasov equation. Under the fluid approximation, these fluid equations are solved instead of the Vlasov equation to describe the plasma evolution; the reduced dimensionality reduces the computational cost, but means kinetic effects where the velocity distribution is important (such as the 'bump-on-tail' instability and Landau damping) are not included.

The zeroth moment of the Vlasov equation is obtained by integrating each term (technically, this is the integration of the product of each term with $v^0 = 1$). Doing this integration returns the *continuity equation*,

$$\frac{\partial n}{\partial t} + \boldsymbol{\nabla} \cdot (n\vec{u}) = 0 \,. \tag{2.19}$$

[4]This discussion focusses on deriving them from the kinetic description of plasmas, but they can be derived from first principles by considering the fluid evolution instead - see derivations in [8] and [10].

Next, the Vlasov equation is multiplied by the first order moment, $m\vec{u}$, before integrating (as performed in [8]). This results in the *equation of motion*,

$$mn(\frac{\partial \vec{u}}{\partial t} + (\vec{u} \cdot \boldsymbol{\nabla})\vec{u}) = qn(\vec{E} + \vec{u} \times \vec{B}) - \boldsymbol{\nabla} \cdot \underline{\underline{P}} + \underline{\underline{P}}_{ie} \ , \tag{2.20}$$

where $\underline{\underline{P}}$ is the pressure tensor arising from inter-species particle reactions, and $\underline{\underline{P}}_{ij}$ is a tensor describing momentum exchange between species (for simplicity, both of these terms are often assumed to be isotropic and thus replaced with scalars).

Finally, multiplying by the second order moment, $mv^2/2$, returns the energy equation,

$$\frac{\partial}{\partial t}(\rho\epsilon) = \boldsymbol{\nabla} \cdot (\rho\vec{u}\epsilon) - \boldsymbol{\nabla} \cdot (P\vec{u}) + \phi - \boldsymbol{\nabla} \cdot \vec{F} + Q \ . \tag{2.21}$$

There are multiple ways to present this equation; the one presented here (based on that presented in [9]) has the kinetic energy subtracted from both sides, and groups terms into a dissipation function ϕ (containing material viscosities). Q is the energy exchange from collisions, F is the heat flux, and ϵ is the internal energy of the fluid.

The fluid equations are used to describe the evolution of the fluid variables directly, without consideration of the distribution function. However, note that the i'th order fluid equation always contains the (i+1)'th moment; for instance, the first-order equation of motion contains the pressure (the second moment of the distribution function), while the second order energy equation contains the heat flux. Such a system of equations is therefore incomplete, and an additional equation is required in order to close the set. If only the first two equations are considered (i.e. up to the equation of motion), this additional equation is the *equation of state*, which expresses the pressure (the higher order term) in terms of the temperature and density (lower order variables). If the energy equation is included, then the equation required for closure is the heat transfer equation (an expression for the higher order heat flux term).

2.4 Magnetohydrodynamics

The equations in the previous section are sometimes referred to as the 'two-fluid' equations, since the electrons and ions are treated as separate fluids (with mass m_s, number density n_s and charge q_s, where the species $s = i, e$ represents ions or electrons respectively), each described by its own continuity equation and equation of motion. A common approximation, magnetohydrodynamics, combines these to form a set of equations governing the overall behaviour of a single combined fluid. This fluid is described by its mass density ρ, charge density σ, mass-averaged velocity $\vec{U}$, current density $\vec{J}$, and total pressure $\underline{\underline{P}}$. These five variables are calculated from the two particle species as

$$\rho = n_i m_i + n_e m_e \ , \tag{2.22}$$

$$\vec{U} = \frac{n_i m_i \vec{u_i} + n_e m_e \vec{u_e}}{n_i m_i + n_e m_e} \ , \tag{2.23}$$

$$\underline{\underline{P}} = \underline{\underline{P}}^i + \underline{\underline{P}}^e \ , \tag{2.24}$$

$$\sigma = n_i q_i + n_e q_e \ , \tag{2.25}$$

$$\vec{J} = n_i q_i \vec{u_i} + n_e q_e \vec{u_e} \ . \tag{2.26}$$

This section briefly outlines how the ideal MHD equations (the equations that govern magnetohydrodynamics) are obtained. A simple version is described based on only the zeroth and first order moment equations (the continuity equation and equation of motion), but alternative formulations are derived in the literature which also include the energy equation.

A continuity equation for the single-fluid is obtained by summing the individual continuity equations for the ions and electrons. This gives the *conservation of mass* equation,

$$\frac{\partial \rho}{\partial t} + \boldsymbol{\nabla} \cdot \left(\rho \vec{U} \right) = 0 \ . \tag{2.27}$$

A *conservation of momentum* equation is likewise obtained, by summing the equations of motion for the two individual species,

$$\frac{\partial \rho \vec{U}}{\partial t} + \rho (\vec{U} \cdot \boldsymbol{\nabla}) \vec{U} = \sigma \vec{E} + \vec{J} \times \vec{B} - \boldsymbol{\nabla} \cdot \underline{\underline{P}} \ . \tag{2.28}$$

The interspecies momentum exchange terms cancel, and thus do not appear in this equation.

An equation for the current density is obtained by multiplying the ion and electron equations of motion by the relevant charge over mass q_s/m_s, and summing the two. The algebra for this is complicated, and to reach the final equation some of the key assumptions of ideal MHD are required:

- The plasma is quasi-neutral, such that σ is approximately zero (this has the effect that Gauss' law, $\boldsymbol{\nabla} \cdot \vec{E} = \sigma/\epsilon_0$, cannot be applied over large length scales in the plasma). This assumption also means the $\sigma\vec{E}$ term in Equation 2.28 can be neglected.

- The length scales and time scales of interest in the plasma are large (i.e. much longer than the period of a plasma oscillation, or a Debye length).

- The effects of electron inertia are negligible, i.e. $m_e \ll m_i$.

Applying these assumptions and completing the algebra allows the *Generalised Ohm's law* to be derived [10, 8],

$$\vec{E} = -\vec{U} \times \vec{B} + \frac{\vec{J} \times \vec{B}}{n_e} - \frac{\boldsymbol{\nabla} \cdot \underline{\underline{P^e}}}{n_e} + \eta \cdot \vec{J} \,. \tag{2.29}$$

This equation is further simplified in ideal MHD by assuming the pressure and $\vec{J} \times \vec{B}$ terms are negligible [8], and assuming no collisions (which leads to the disappearance of the $\eta \cdot \vec{J}$ term, where η is the resistivity tensor, which includes the interspecies collision terms). This leads to a simplified form of

$$\vec{E} = -\vec{U} \times \vec{B} \,. \tag{2.30}$$

The ideal MHD equations are then closed using Maxwell's equations (in particular, forms of Ampere's and Faraday's laws reduced by the MHD approximations) and the equation of state. The complete set of ideal single-fluid MHD equations consists of:

$$\frac{\partial \rho}{\partial t} + \boldsymbol{\nabla} \cdot \left(\rho \vec{U}\right) = 0 \,, \tag{2.31}$$

$$\frac{\partial \rho \vec{U}}{\partial t} + \rho(\vec{U} \cdot \boldsymbol{\nabla})\vec{U} = \vec{J} \times \vec{B} - \boldsymbol{\nabla} \cdot \underline{\underline{P}} \,, \tag{2.32}$$

$$\vec{E} = -\vec{U} \times \vec{B} \,, \tag{2.33}$$

$$\boldsymbol{\nabla} \times \vec{B} = \mu_0 \vec{J} \,, \tag{2.34}$$

$$\boldsymbol{\nabla} \times \vec{E} = -\frac{\partial \vec{B}}{\partial t} \,, \tag{2.35}$$

$$P = P(\rho) \,. \tag{2.36}$$

More accurate variations of this formalism include additional terms in the equations (accounting for collisions and resistivity), or include the energy equation as well (which requires a further heat transport equation for closure). Since the full distribution function does not need to be evaluated, fluid codes (which solve these equations to describe the fluid behaviour of plasma) are used to perform simulations over the time and length scales required to describe a full ICF implosion. However, it is also important to note that the MHD assumptions mean that a range of plasma effects are neglected. In particular, the assumption of long time and length scales means that 'fast' or high frequency phenomena, such as the plasma oscillations discussed in Section 2.2.2, cannot be described in such codes. It is important to bear such limitations in mind when performing simulations.

2.5 Additional physics in fluid codes

The previous section demonstrated how the moment equations are not complete, and require further models (the equation of state, or the heat transport equation) in order to be solved. In reality, simulation codes based on this approach typically involve additional models, which are used to improve accuracy and are included by adding additional terms to the MHD equations.

This section provides a brief overview of some of the other physics that is important in describing plasmas (such as laser absorption and ionisation) and how these are included in fluid simulations, as well as outlining some of the additional physics models that are added (such as heat and radiation transport, and the equation of state). As this thesis presents a large number of hydrodynamic

simulations using the 1D radiation-hydrodynamics code HYADES [12], the specific models that are used in this code are indicated where relevant. This section is intended to provide a brief overview of this physics, and a more complete description can be found in [9] (on which it is based).

Discretisation of the plasma

In Section 2.3.2, it was described how the Eulerian forms of plasma equations track changes at a fixed point in space, while Lagrangian forms follow a fluid 'packet' and track how it changes as it moves. This distinction is particularly evident in simulation codes.

In Eulerian codes (which solve the Eulerian form of the fluid equations), the fluid is described by a fixed mesh, where the fluid variables are described by a single value in each zone between mesh points. Following the Eulerian perspective of tracking changes at a fixed point in space, the mesh points are fixed throughout the simulation and the fluid is free to flow between zones. The Lagrangian form of the equations, however, describes how fixed 'packets' of fluid evolve with time. As such, the mesh points are not fixed in time but rather move with the plasma, and each zone describes a fixed and unchanging 'packet' of fluid. In Lagrangian codes the mass of each zone is therefore constant as a function of time (and the mesh boundaries move to maintain this), while in Eulerian codes the position of the mesh is fixed and the mass per zone changes as the fluid moves within it.

Both approaches have their advantages and disadvantages. The movement of the Lagrangian mesh means that interfaces can be tracked throughout the simulation, and the fact the mesh moves with the mass can be helpful for maintaining resolution if the plasma is compressed (as in ICF). However, in 2D and 3D codes, the movement of the mesh can result in tangling, which can cause the code to crash - something Eulerian codes do not suffer from. HYADES is a Lagrangian code (as it is 1D mesh tangling is not a concern, but this is a problem in its 2D equivalent h2D).

Laser propagation and absorption

The critical density associated with the laser frequency places a limit on how far into a plasma the light propagates before it is reflected[5]. The coronal plasma is partially or fully ionised, meaning that there are both bound and free electrons that the incoming light interacts with. At the lower end of 'high intensity' (high intensity defined here as around 10^{13} W cm^{-2} and above), the photon energy is low compared to the ionisation potential and so absorption by free electrons dominate. Inverse bremsstrahlung is thus the most significant absorption mechanism; this is a collisional process, where a free electron is accelerated by a photon in a process mediated by a nearby ion (the reverse of the more famous bremsstrahlung emission). Inverse bremsstrahlung scales as $T_e^{-3/2}$, and thus absorption through this process decreases as the temperature increases. This is the mechanism through which most energy is absorbed in ICF applications.

Other absorption processes also play a role. Resonant absorption occurs very close to the critical density surface, where the plasma frequency matches the frequency of the laser, and thus Langmuir waves can be excited resonantly[6]. This is a collisionless process, and becomes more significant at higher temperatures. Other mechanisms also become significant at high intensity, and see the incoming EM wave excite a range of plasma modes; these are termed 'parametric instabilities', and will be treated in a later section.

Fluid codes typically contain a 'ray-tracing' model, which calculates the path that an incoming ray of laser light takes. Based on the conditions along this path, an absorption/reflection probability can be calculated. `HYADES` models both inverse bremsstrahlung and resonant absorption (while the Langmuir waves driving this absorption are not included, the energy deposition at the critical surface that would result from this process is approximated by the code). Some alternative

[5]While the laser initially is incident on a solid surface, the laser will quickly ablate away some solid material. This establishes a 'coronal plasma', with a density gradient that increases towards the main fuel. The incoming laser now travels through this plasma, and the laser absorption processes occur within it.

[6]In fact, the relationship between density and refractive index means that the incoming ray is typically reflected slightly before the critical density surface.

codes also include models to estimate the effects of parametric instabilities (again, without simulating the plasma oscillations).

Ionisation

As a laser interacts with a material it will ionise it, through a variety of different mechanisms. Generally the photon energies are not high enough to ionise electrons directly, and thus multi-photon ionisation processes are required. Field ionisation becomes significant when the electric field associated with the laser becomes comparable to the Coulomb barrier trapping the electron; the barrier becomes perturbed such that at some distance from the electron the barrier's potential energy is lower than the electron's kinetic energy, and the electron is thus able to quantum-mechanically tunnel through the barrier and escape the potential well.

Modelling the ionisation state of the plasma (which will affect the charge and current of the fluid in the MHD equations) can be performed through a range of models. The Saha equation can be used to estimate the relative populations of ionisation levels for low density plasmas at high temperatures in local thermal equilibrium, or LTE; however, it performs less well when these conditions are not satisfied. The Thomas-Fermi ionisation model considers an electron gas consisting of free and bound electrons, and treats this gas using Fermi-Dirac statistics (as opposed to the Boltzmann statistics used by the Saha equation). The model used for the `HYADES` simulations in this thesis is an average atom LTE ionisation model. This model, rather than attempting to find the shell populations of many different atomic configurations, considers a single average-atom with a few discrete energy levels. The average-atom is allowed to have a non-integer number of bound electrons in each level, in order to best approximate the average over all the atomic configurations. It includes a model for electron screening, and LTE equations are used to calculate the average shell populations.

Radiation transport

'Radiation hydrodynamics' extends the kinetic description of the plasma by adding a radiation fluid, which exchanges energy and momentum with the ion and electron

fluids. The radiation is described through the specific intensity, $I(\vec{r}, \upsilon, \vec{\Omega}, t)$ where υ is the frequency and $\vec{\Omega}$ are two angular coordinates. This quantity is an energy flux, and is related to the photon distribution function through

$$I_\upsilon(\vec{r}, \vec{\Omega}, t)\, \mathrm{d}\upsilon \mathrm{d}\vec{\Omega} = h\upsilon c f(\vec{r}, \upsilon, \vec{\Omega}, t)\, \mathrm{d}\upsilon \mathrm{d}\vec{\Omega}\ , \tag{2.37}$$

where I_υ is the spectral intensity for a given frequency, h is Planck's constant and c is the speed of light.

This third fluid is coupled to the other fluids by adding a radiation flux term F_υ to their energy conservation equations, which then propagates through into the single-fluid equations. The inclusion of an extra variable means that an additional equation must be included in order to close the set. This new equation is the radiation transport equation - the equation of motion for the photons in the radiation fluid. This equation is expressed as

$$\frac{1}{c}\left(\frac{\partial I_\upsilon}{\partial t} + c\vec{\Omega} \cdot \boldsymbol{\nabla} I_\upsilon\right) = \rho \kappa'_\upsilon [B_\upsilon(T) - I_\upsilon]\ , \tag{2.38}$$

where $B_\upsilon(T)$ is the black-body intensity at frequency υ for a temperature T, and $\kappa'_\upsilon = \kappa_\upsilon(1 - \exp\{-h\upsilon/k_B T\})$ where κ_υ is the absorption coefficient.

This equation cannot be solved exactly, and so approximations are required. The most common is the diffusion approximation, which assumes that the total radiation flux is isotropic and thus the net flux is zero. This assumption is valid for optically thick plasmas, and allows the radiation transport equation to be converted into the diffusion equation, which is then solved by a simulation code (this derivation is not reproduced here). To solve the diffusion equation the emission, absorption and scattering coefficients must be calculated. These variables are continuous in frequency - this is resolved in the derivation by averaging the quantity over a frequency range[7]. These variables vary over many orders of magnitudes as a function of frequency, and thus a single average is often a poor approximation. As a result, radiation hydrodynamics codes such as `HYADES` use 'multi-group' radiation transport, where a number of frequency groups are specified, and a separate average absorption coefficient is calculated for each group.

[7]The 'Planck mean' is used for optically thin plasmas, while the 'Rosseland mean' is used when the plasma is optically thick.

Thermal transport

The second moment of the Boltzmann equation gives the energy transport equation (Equation 2.21). This contains the heat flux $\vec{F}$, which is described using Fourier's law

$$\vec{F} = -\chi \boldsymbol{\nabla} T \ , \tag{2.39}$$

where χ is the thermal conductivity constant. In order to close the set of MHD equations, an equation for this thermal conductivity is required.

The most common thermal transport model used to calculate the transport coefficients is the Spitzer-Härm model, where the Fokker-Planck equation is used to find the collision term and it is assumed that the velocity distribution is a Maxwellian with a small perturbation. This model finds that $\chi \propto \frac{T^{5/2}}{\ln \Lambda}$, where the denominator is the 'Coulomb logarithm'.

However, the Spitzer-Härm model overestimates electron heat transport into cold-dense material in the presence of large temperature gradients. This can result in electron fluxes above the 'free streaming limit' - the maximum possible flow of thermal energy. To account for this, hydrodynamic simulations such as `HYADES` typically apply a flux-limiter; a parameter set with a value typically around 0.05. The simulation code will then use the smaller value of either the Spitzer-Härm solution, or the free-streaming flux multiplied by this flux-limiter.

Equation of state

The equation of state (EOS) needs to be evaluated for the equations to be closed; this means being able to determine the material pressure as a function of its density and temperature. This is done through determining the partition function describing a material, using thermodynamic relations to determine the free energy, and then using this to calculate the pressure, internal energy and entropy as a function of the density and temperature.

There are a number of potential EOS models. A simple common model is the Thomas-Fermi equation of state; this treats all the electrons (free and bound) as a single fluid obeying Fermi-Dirac statistics, while all the ions are treated as

nuclei obeying Boltzmann statistics. This model is simple to evaluate, and is a reasonable description for high ionisation and high temperature plasmas. One of the key limitations is that it does not include coupling forces, and thus it is a poor description of low temperature solids. Another important EOS model is the Grüneisen equation of state, which describes solids where the atoms are held in a lattice structure. This formulation treats the pressure and energy in the solid as a sum of compressional and thermal components, and defines a Grüneisen parameter Γ which is used to define the material and is proportional to the ratio of its thermal pressure to its thermal energy. The Grüneisen EOS can be used to describe low temperature solids, but it assumes that the movement of the atoms around the equilibrium positions are small and thus breaks down as the thermal motion of the atoms increases.

It is clear that the equation of state is a complicated topic. More complicated models exist based on the quotidian Equation of State (QEOS) [13]. This uses a modified Thomas-Fermi model with a bonding correction term to describe the electrons, while combining a number of models (including the Grüneisen EOS) to describe the ions - producing an EOS that is multi-phase and thermodynamically consistent. Many variations of this model exist.

Hydrocodes such as `HYADES` provide a number of options for the equation of state, but there are two main ones used for ICF simulations. Firstly, many codes accept tabulated equation of state data such as the SESAME tables [14], which are maintained for a number of materials by Los Alamos National Laboratory. These EOS tables can be based on a range of models (including first-principle and QEOS models), and are widely tested against experiments. Alternatively, `HYADES` also provides an in-line QEOS model that can be used. Both approaches are used to describe different materials in the simulations presented in this thesis.

2.6 Basic description of laser-driven inertial confinement fusion

This section provides a brief overview of a laser-driven inertial confinement fusion (ICF) implosion; a more detailed discussion is available in [15]. A driver is used to compress a capsule containing deuterium-tritium fuel to high pressures, densities and temperatures. In direct-drive ICF, this driver is a laser pulse which irradiates the capsule directly, while in indirect-drive ICF the laser pulse is incident on the inside of a hohlraum, generating a bath of x-rays which then irradiate the capsule. In either case, the intense electromagnetic radiation ablates the outer surface. This ablation and expansion of material accelerates the remaining material inwards, compressing the capsule. The driver is typically designed to provide a pressure vs time profile such that a series of shocks are generated to compress the capsule at different times; the use of multiple shocks reduces the entropy injected into the fuel, allowing higher densities to be achieved than is possible with a single shock.

The capsule typically consists of a central DT vapour region. This is then surrounded by a layer of cold dense DT fuel (ice in many cases, or liquid DT-wetted-foams as is the focus of this thesis), which in turn is typically surrounded by an ablator layer. The high density ablator is used to absorb the radiation and accelerate the rest of the material inwards effectively. This thesis will focus on 'central hotspot ignition', where as the capsule implodes it forms a central 'hotspot' (consisting of high temperature but low density material) and a surrounding 'shell' consisting of high density but low temperature material.

As the capsule implodes, the pressure in the centre of the capsule increases. At some point this overcomes the drive pressure, causing deceleration of the shell, with the shell kinetic energy converted to thermal energy of the hotspot. Eventually the shell comes to a complete stop, at the time of maximum compression, before being accelerated outwards again. The fuel is 'confined' for a brief period by the inertia of the shell, before it rapidly begins to disassemble.

If the temperature and density within the hotspot are sufficiently high, fusion reactions begin to occur. When DT fuel is used as described here, the

dominant reaction is

$$D + T \rightarrow \alpha(3.5\text{MeV}) + n(14.1\text{MeV}) \ , \tag{2.40}$$

where D is a deuterium ion, T is a tritium ion, α is an alpha particle (helium ion) and n is a neutron [15]. This reaction generates 17.6 MeV which is distributed between the kinetic energy of the neutron and the alpha particle. DT fuel is used because this reaction has the largest cross section at low energies.

The neutron has a very large mean free path and therefore a very low chance of absorption, and thus escapes from the fuel assembly. The alpha particle however has a shorter mean free path and thus is more readily 'absorbed' by the plasma (i.e. undergoes collisions with the plasma electrons, such that its energy is absorbed). This returns 20% of the 17.6 MeV generated from the fusion reaction back into the plasma, providing further heating. A balance of energies shows that alpha particle 'self-heating' overtakes the losses in the hotspot (due to bremsstrahlung emission) at around 4.3 keV for a pure DT plasma[8]. This is referred to as ignition. In central hotspot ignition, the hotspot undergoes this process, and then generated alpha particles heat the cold-dense DT shell, inspiring further fusion reactions and generating a propagating burn wave that travels out from the centre of the capsule[9]

The plasma conditions required to achieve ignition were first quantified by Lawson [17], who estimated the temperatures and densities based on the power balance in a reactor. This is most commonly expressed through the famous 'triple product'

$$n\tau_E T \approx 3.3 \times 10^{15} \text{ cm}^{-3} \text{ s keV} \ , \tag{2.41}$$

for a temperature $8 < T < 25$ keV, and where n and τ_E are the density and plasma confinement time respectively [15]. However this expression is more suited

[8]Bremsstrahlung emission scales with atomic number Z as Z^2, and so the presence of heavier ions in the hotspot will lead to higher required temperatures. For wetted-foams, the approximately 2% atomic fraction of carbon in the plasma raises the required temperature to 5.2 keV [16].

[9]It is worth noting that as the temperature increases, the plasma becomes increasingly transparent to alpha particles and their stopping range also increases. This means that as the hotspot heats up, the generated alpha particles can propagate larger distances and thus heat the cold DT shell. This enables the generated energy to propagate outwards and the burn wave to be launched, rather than alpha-deposition being confined to and further heating the high temperature region only [15].

to magnetic confinement fusion. For ICF, ignition instead needs to be defined through the hotspot conditions - the hotspot must be large enough for self-heating to occur, and to produce enough energy to launch the propagating burn wave. This requires a large number of the alpha particles to be absorbed within the hotspot, leading to conditions on the areal density ρR (the integral of density with radius over a specified volume, and proportional to how much fuel an alpha particle would need to pass through to escape) and temperature T of the hotspot. A variety of measures of 'margin' (or how well a capsule is performing) have been developed and used for ICF, but a simple and often quoted one is that presented by Cheng *et al.*:

$$(\rho R)_{\mathrm{HS}} \geq 4\frac{\kappa_{\mathrm{C}}}{f_{\mathrm{T}}}\left(\frac{T}{\mathrm{keV}}\right)^{-2.5} \ \mathrm{g/cm^2}\ , \tag{2.42}$$

where $\kappa_{\mathrm{C}} \simeq 5.514$ is a constant depending on the power law approximation for DT reactivity, and $f_{\mathrm{T}} \propto \sqrt{\rho_{\mathrm{p}}/\rho_{\mathrm{HS}}}$ is a tamping factor (typically assumed to be 1 for basic estimates of ignition threshold) based on the density of the hotspot (ρ_{HS}) and the cold fuel (ρ_{p}) [18]. This quantifies ignition in terms of the temperature T and areal density ρR of the DT fuel within the hotspot. This formula suggests that for an ion temperature of $T_{hs} \approx 4$keV, a hotspot areal density of approximately $(\rho R)_{hs} \approx 0.3$ g/cm^2 is required.

This criteria determines whether ignition occurs, but ignition on its own is not sufficient to generate high gain. For this, the capsule shell must contain a large mass of compressed DT fuel, and a large fraction of this fuel must be burnt (undergo fusion) before the shell disassembles. Ignition sparks the burn wave through this dense fuel, but how efficiently it burns depends upon the shell conditions. This is quantified using the fuel burn-up fraction Φ, which compares the number of fusion reactions that occur to the total number of DT ion pairs in the capsule (i.e. the maximum number of potential reactions). This is typically estimated as

$$\Phi = \frac{\rho R}{6.3\ \mathrm{g/cm^2} + \rho R}\ , \tag{2.43}$$

where in this case ρR is the areal density of the full fuel assembly (i.e. hotspot and shell) [19]. Using this equation, it is evident that the total areal density must

be around 3 g/cm^2 in order to achieve a burn up factor of around 0.3. It is thus clear that achieving a high gain ICF implosion is therefore dependent on achieving both high temperatures and areal densities in the hotspot, as well as achieving large areal densities in the compressed capsule shell.

2.6.1 Alternative ignition schemes

In central hotspot ignition the driver fulfils the dual roles of fuel compression and hotspot formation simultaneously. This leads to stringent symmetry requirements. Alternative ignition designs have been proposed which separate these two processes. Two significant alternative approaches are discussed in this section.

In fast ignition [20], the driver is used to compress the fuel to high density. A lower laser power is used to compress the capsule at much lower velocity and form an isochoric fuel assembly, where the fuel is roughly constant density throughout[10]. The low laser power and more moderate accelerations serve to prevent significant hydrodynamic instability growth. A particle beam is then used to heat a small volume at the fuel surface to high temperature; this forms the hotspot, generates fusion reactions, and causes a burn-wave that propagates through the rest of the fuel [21, 22]. The particle beam could consist of either electrons or ions. Schemes have also been proposed using cones to allow the particle beam to irradiate fuel closer to the centre of the assembly (so called 'cone-in-shell' fast ignition) [23, 24]. A major challenge to achieving fast ignition is generating and transporting the particle beams, and using them to heat the fuel efficiently [25].

Another scheme that has generated significant interest is shock ignition [26]. Here, lower laser powers are again used to compress the fuel at lower velocities, reducing hydrodynamic instability growth - but unlike fast ignition, the compression results in an isobaric fuel assembly as in central hotspot ignition. A high intensity laser pulse is then applied late in the implosion to drive an extremely strong final shock through the fuel, which generates high temperatures and is responsible for forming the hotspot. The main concern for shock ignition is that the high intensity

[10]Central hotspot ignition is said to form an isobaric fuel assembly, where the pressure is instead constant throughout the hotspot [15].

of the late laser pulse can result in significant parametric instability growth, which could lead to hot electron generation and preheating of the fuel. An alternative scheme - shock augmented ignition - has recently been proposed to mitigate this, where a dip in the laser power is used before the final shock [27]. This dip means that the laser intensity is then increased to more moderate levels while still generating the strong shock (it is the increase in intensity that is important, rather than the absolute value), and thus high levels of parametric instability growth are avoided.

2.6.2 Polar direct drive

High energy direct-drive experiments are currently limited by a lack of experimental facilities. Direct-drive experiments are currently primarily conducted using OMEGA at the University of Rochester Laboratory for Laser Energetics; this facility has 60 beams focussed on the target, but has a maximum energy of 40 kJ. Laser Megajoule (LMJ) will also have the capability for direct-drive, but is still being developed.

Facilities such as the NIF are designed for indirect-drive, and the lasers are thus positioned at the two capsule poles rather than in a spherical configuration. A technique called polar-direct-drive (PDD) is used to perform direct-drive shots at such facilities [28]; the beam pointing and powers of the lasers are adjusted, in order to provide as close to spherical direct-drive compression as possible despite the polar laser arrangement [29]. This approach has obvious issues with drive uniformity[11], and is particularly susceptible to cross-beam energy transfer [30]. The NIF was previously limited to performing PDD shots at room temperature [31], preventing capsules with DT ice or wetted-foam layers from being used, but this has recently been addressed and cryogenic PDD shots are now in preparation.

2.7 Instability growth

The key challenge for achieving high performing ICF implosions is the prevention, mitigation and control of both hydrodynamic and parametric instabilities. Instabili-

[11] Capsule designs such as 'Saturn' capsules [29] have been proposed in attempts to mitigate this.

ties are effects beyond the pure compression and burn physics discussed so far, which typically result in reductions to fusion performance. Their growth is one of the key reasons why ICF experiments have underperformed compared to simulation [32, 33, 34]. The following sections explore the physics of these two types of instability.

2.7.1 Hydrodynamic instabilities

As the ICF capsule is compressed and implodes, small perturbations arise which grow and significantly affect the course of the implosion [15]. These hydrodynamic instabilities result in significant asymmetry of the implosion, leading to non-uniform (and reduced) compression, and also lead to the mixing of the different layers (such as of ablator material into the cold DT fuel layer, or of cold fuel into the hotspot). These effects significantly degrade fusion performance [35].

The most significant hydrodynamic instability for ICF is the Rayleigh-Taylor (RT) instability, which occurs when a lower density fluid is accelerated into a higher density one. Small perturbations undergo rapid growth, leading to 'spikes' at the interface of the two materials [15]. This leads to mixing, and in extreme cases complete break-up, of the different layers of material within the capsule. RT instability occurs at the outer edge of the shell throughout the implosion, and at the shell/hotspot boundary as the shell decelerates. RT growth at the capsule outer surface is reduced to some extent by ablative stabilisation, where the ablation of the surface reduces the perturbations as they grow[12]. The ablation-stabilized growth rate for the RT instability can be expressed as

$$\gamma = \alpha\sqrt{ak} - \beta k u_{\mathrm{a}} \,, \tag{2.44}$$

where a is acceleration, k is the wave number of the perturbation, α and β are numerically determined coefficients, and u_{a} is the ablation velocity [37, 38].

RT instability growth is 'seeded' by small initial perturbations, which can arise from a number of possible factors [39]. These include asymmetries in the laser radiation (or x-ray) drive, small perturbations to the capsule (such as pits/voids),

[12]Some ablative stabilisation is also provided at the hotspot/shell boundary during the deceleration phase due to alpha heating [36]

the capsule fill-tube (a small tube used to fill the capsule with the DT fuel), and the capsule support tent/stalk (i.e. the system used to hold the capsule in place). Efforts are made to minimise and control all possible RT seeds as much as possible to prevent and minimise instability growth.

Other hydrodynamic instabilities also affect ICF capsules, albeit to a lesser extent [15]. The Richtmyer-Meshkov instability occurs when a shock wave crosses between two materials and the interface is not-perfectly flat; the velocity of the surface is thus non-uniform, and the perturbations grow. These perturbations then in-turn seed RT instability growth. Another example is the Kelvin-Helmholtz instability, which occurs when there is shear flow between two different fluids. This occurs as RT 'spikes' develop, and thus plays a role in the development of the later stages of this growth.

2.7.2 Parametric instabilities

Parametric instabilities are non-linear instabilities arising from laser-plasma interactions. In such instabilities, the wave-wave interactions leads to the amplification of 'signal' and 'idler' waves, at the expense of an incoming 'pump' wave. For this to occur, the following wave-matching conditions must be satisfied:

$$\vec{k_0} = \vec{k_1} + \vec{k_2} \ , \tag{2.45}$$

$$\omega_0 = \omega_1 + \omega_2 \ . \tag{2.46}$$

$\vec{k}$ is the wave vector[13] and ω is the angular frequency, while the subscript 0 represents the initial (pump) wave and 1 and 2 represent the excited waves (the signal and idler) [8].

A simple analysis of the frequency wave-matching conditions demonstrate how the resonant excitation of a plasma mode can occur. The incoming pump is an EM wave, which (in the simplest cases of SRS and SBS, defined and described in the following subsections) scatters into another EM wave. The third wave is a

[13] $\vec{k}$ is expressed here in vector form since the relative propagation directions for the waves is significant for the wave matching. The vector nature of this was neglected when discussing the 1D dispersion relations for simplicity, and it was instead expressed as the scalar k.

plasma mode (such as those described in Section 2.2.2 - the type of mode defines the exact nature of the instability). The scattered wave (the idler, with frequency ω_2) will beat with the large amplitude pump wave (with frequency ω_0). Where the wave-matching conditions are satisfied, this beat frequency will match the frequency of the plasma mode (the signal), $\omega_0 - \omega_2 = \omega_1$, and thus this mode undergoes resonant excitation [40].

In ICF, parametric instabilities occur when the high intensity laser light passes through the ablated/coronal plasma. As energy is transferred from the incoming beams to plasma modes/scattered waves, this results in reduced absorption of laser light through inverse bremsstrahlung (the desired from of absorption for driving capsule compression), less symmetry in the drive (which can thus lead to less uniform compression, and seed hydrodynamic instabilities), and preheating of the fuel (again reducing performance). This section discusses the key types of parametric instability for ICF, following the descriptions in [8] and [41].

Stimulated Raman scattering (SRS)

In stimulated Raman scattering, the incoming electromagnetic wave scatters from a Langmuir wave (or plasmon), producing a scattered EM wave. By substituting into Equation 2.46 the dispersion relations for the EM waves ($\omega_0^2 = \omega_p^2 + c^2k_0^2$ for the incident wave, and $\omega_2^2 = \omega_p^2 + c^2k_2^2$ for the scattered wave) and the Langmuir wave ($\omega_1^2 = \omega_p^2 + \frac{3}{2}k_1^2v_{\text{th}}^2$), it is clear that $\omega_o > 2\omega_p$. The SRS instability therefore only occurs for densities below the quarter critical density, $n < n_c/4$. The resulting Langmuir wave can lead to fast moving 'hot' electrons being generated [42]. These electrons travel significant distances before being absorbed and thus result in preheating of the cold DT fuel within the capsule, making it harder to compress and thereby degrading the overall fusion performance.

Two-plasmon decay (TPD)

In two-plasmon decay, the scattered EM wave in SRS is replaced by a second plasmon. This requires that $\omega_0 \approx 2\omega_p$, which means that the instability is only possible over a

small density range around the quarter critical density. As with SRS, this instability leads to preheating of the cold DT fuel through the generation of hot electrons [43].

Stimulated Brillouin scattering (SBS)

In stimulated Brillouin scattering, the incoming EM wave scatters into an ion-acoustic wave and a scattered EM wave. Since the frequency of an ion-acoustic wave is relatively small ($\omega_1 = k_1 v_s$), the frequency of the scattered EM wave is only slightly lower than that of the incident wave, and the instability occurs throughout the entire coronal plasma. SBS is a significant instability for ICF due to its role in cross-beam energy transfer (CBET), where energy is transferred out from the incoming beams[14] - an effect which can reduce the absorption of laser light by 10-20% [44]. CBET can be useful in indirect-drive however, where wavelength detuning is used to control the energy transfer between beams in order to tune the implosion symmetry [45].

Mitigation of parametric instabilities

Parametric instabilities are generated when high intensities of EM radiation pass through a plasma. There is thus a threshold above which these processes begin to occur in significant quantities, and they are mitigated by keeping the intensity below this value. However, this can be challenging to balance with the high intensity demands of ICF. Alternatively, parametric instabilities are mitigated by using high bandwidth lasers, so that the intensity is spread over a wider range of frequencies [46]. This keeps the intensity of individual frequencies under the threshold, and is seen as a key instability mitigation strategy for direct-drive in the future [47].

2.8 Shocks and impedance matching

2.8.1 Shock physics

Shocks, referring to a sudden change in plasma properties over a thin transition region, are ubiquitous in plasmas and inertial confinement fusion. In an ICF

[14] In direct-drive, this typically occurs when rays from the edge of one beam refract in the coronal plasma and are turned outward, and then overlap and undergo CBET with central rays from the other incoming beams.

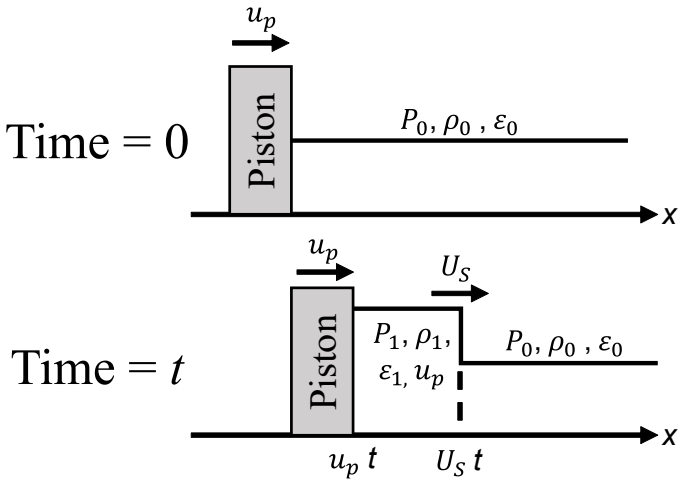


Figure 2.1: A piston moving at velocity u_p drives a shock through a medium. The shock moves at velocity U_S, leaving material behind it in a new shock state denoted by the subscript 1. Based on a diagram in [48].

implosion, shocks are generated when there is a change in the applied pressure (such as an increase in the laser intensity), and are used to accelerate the shell inwards. Shock physics is of great importance to Chapters 6 and 7 of this thesis, which will describe an experiment where the shock compression of a foam material is measured to help quantify its equation of state. This section therefore provides a brief overview of the key principles of shock physics, based on the explanation in [48].

Consider a material that is initially at rest, before being compressed by a piston moving at a velocity u_p. This scenario is displayed in Figure 2.1. After a time t the piston itself has moved by a distance $u_p t$, while the disturbance (shock) it has generated has propagated with a 'shock velocity' U_S through the material to a position $U_S t$. Behind the shock the material is at the piston pressure and velocity P and u_p, and at a new higher density ρ_1. Ahead of the disturbance, the material is unaffected. As the newly shocked material moves at velocity u_p, this quantity is known as the 'particle velocity'.

Three equations are derived based on the conservation of mass, momentum, and energy across the disturbance. Considering that the mass behind the shock front cannot have changed, the expression

$$\rho_0 U_S = \rho_1 (U_S - u_p) \tag{2.47}$$

is obtained[15]. Considering the momentum of the material (which changes according to the impulse, arising from the difference in pressure either side of the boundary), the following equation is obtained:

$$\rho_0 U_S u_p = P_1 - P_0 \ . \tag{2.48}$$

Finally, considering the change in total energy (both internal energy ϵ and kinetic energy), which must equal the work applied by the piston as it moves the distance ut, enables a third equation to be derived:

$$\rho_0 U_S(\epsilon_1 - \epsilon_0 + \frac{u_p^2}{2}) = P_1 u_p \ . \tag{2.49}$$

These equations relate the four shock variables U_S, u_p, P and ϵ to one another. If the state of the material ahead of the shock is known, then finding any two of these variables allows the others to be calculated. They are referred to collectively as the *Rankine-Hugoniot* relations. They were derived here in a stationary frame (where the shock moves with velocity U_S), and are typically expressed as

$$\rho_1 = \frac{\rho_0 U_S}{U_S - u_p} \ , \tag{2.50}$$

$$P_1 - P_0 = \rho_0 U_S u_p \ , \tag{2.51}$$

$$\epsilon_1 - \epsilon_0 = \frac{P_1 u_p}{\rho_0 U_S} - \frac{{u_p}^2}{2} \ . \tag{2.52}$$

2.8.2 Rankine-Hugoniot relations in a moving frame

Alternatively, the same equations can be considered in a moving frame, where the shock is stationary. The frame now moves at velocity U_S, as displayed in Figure 2.2. This means the unshocked material moves into the stationary shock front at a velocity $u_0 = -U_S$, while the material behind the shock moves away at the slower velocity $U_1 = -(U_S - u_p)$. This transformation is then applied to each of Equations 2.50, 2.51 and 2.52 to give the Rankine-Hugoniot relations in an alternative form,

$$\rho_1 u_1 = \rho_0 u_0 \ , \tag{2.53}$$

[15]Here, as in the following equations, t and the area A cancel from both sides of the equation during the derivation.

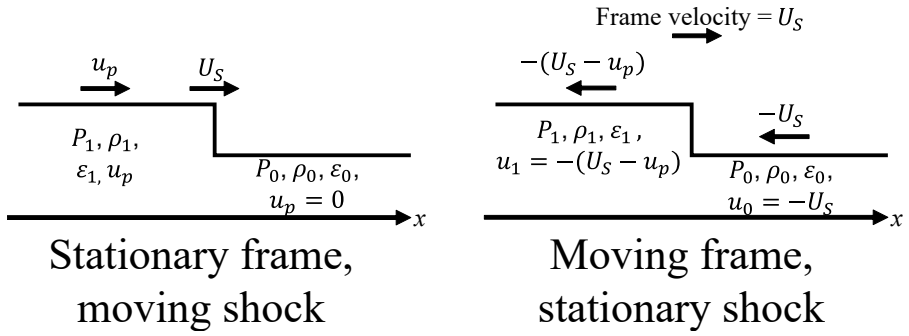


Figure 2.2: The shock as viewed in the two frames. In the stationary frame, the shock moves with velocity U_S, the shocked material with velocity u_p, and the unshocked material is stationary. In the moving frame, the shock is stationary, the shocked material has velocity u_1, and the unshocked material has velocity u_0. The other variables are unchanged.

$$P_1 + \rho_1 u_1^2 = P_0 + \rho_0 u_0^2 \,, \tag{2.54}$$

$$\epsilon_1 + \frac{P_1}{\rho_1} + \frac{u_1^2}{2} = \epsilon_0 + \frac{P_0}{\rho_0} + \frac{u_0^2}{2} \,. \tag{2.55}$$

It is also possible to derive these relations by integrating over the differential equations describing the conservation equations given in Section 2.3.3).

These equations can be further manipulated to derive some other useful expressions. Equations 2.53 and 2.54 give new expressions for the velocities u_1 and u_0,

$$u_1^2 = \nu_1^2 \frac{P_1 - P_0}{\nu_0 - \nu_1} \,, \tag{2.56}$$

$$u_0^2 = \nu_0^2 \frac{P_1 - P_0}{\nu_0 - \nu_1} \,, \tag{2.57}$$

where ν is the specific volume, $\nu = \frac{1}{\rho}$. Substituting Equations 2.56 and 2.57 into Equation 2.55 then gives

$$\epsilon_1 - \epsilon_0 = \frac{1}{2}(P_0 + P_1)(\nu_0 - \nu_1) \,. \tag{2.58}$$

Equation 2.58 is known as the Hugoniot relation. The Hugoniot is also defined,

$$P_1 = H(P_0, \nu_0, \nu_1) \,, \tag{2.59}$$

which depends on the specific energy function $\epsilon(P, \nu)$. The form of the Hugoniot is notable as it is a function of the initial state of the material (P_0, ν_0). The significance of this will be discussed in the next section.

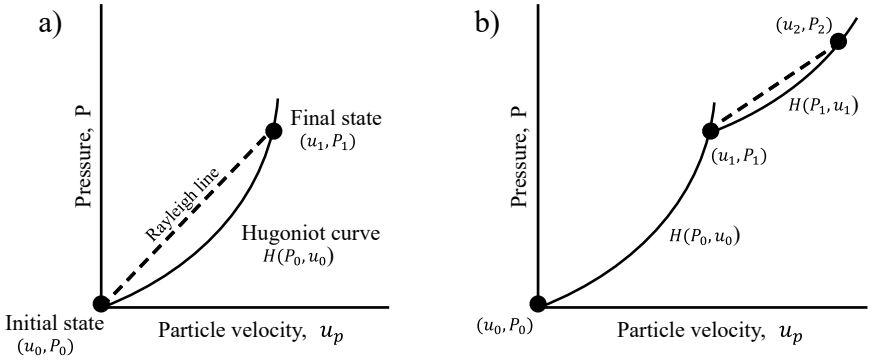


Figure 2.3: a) The principal Hugoniot in (u_p, P) space. A shock compresses material in the initial state along a compression path described by the Rayleigh line to a shock state on the Hugoniot curve. The curve represents the range of possible shock states that could be achieved for a given initial state. b) If a subsequent shock passes through the material, the shock state it will generate lies somewhere on a secondary Hugoniot, which is a unique curve generated for the new initial condition (the shock state achieved by the first shock).

2.8.3 The Hugoniot

The Hugoniot curve describes the locus of possible shock states that are achievable when a single shock is driven through a target with the specified initial state. It is formally defined as a function of P and ν but the Rankine-Hugoniot relations mean it can be equivalently expressed as a function of any two of the shock variables, and it is often more convenient to discuss it in (ρ, P) or (u_p, P) space. Further discussion of this topic is provided in [49].

The dependence of the Hugoniot on the initial state is significant, as is highlighted in Figure 2.3. For the initial state (u_0, P_0), the Hugoniot curve is plotted; this describes the states accessible by a shock from that initial point. Driving a shock with pressure P_1 through the target allows the shock state (u_1, P_1) to be achieved. However, following this with a higher shock does not allow a new state on this curve to be accessed, due to the fact that the material ahead of this new shock is no longer in state (u_0, P_0), is but instead (u_1, P_1). To find the state achieved by this second shock requires a new Hugoniot curve to be plotted, $P_2 = H(P_1, \nu_1)$; this curve will pass through the state (u_1, P_1), and describe the locus of states that the new shock is able to reach.

The Hugoniot curve describing the states that are accessible from an initially unshocked material, $P_1 = H(P_0 = 0, \nu = \nu_0)$, is particularly significant for two obvious reasons; most materials begin experiments unshocked (and so the first shock state achieved will be on this curve), and the states on this curve are the easiest to measure and access. This curve is known as the 'principal Hugoniot', and is determined experimentally by measuring the state achieved when shocks of different strengths are passed through initially unshocked material.

It is also worth noting that the Hugoniot curve is not the path that the material takes through (P, ν) space as it is compressed (rather, it is the possible end states of such a compression). The compression path is given by a different curve called the 'Rayleigh line' [49]. The Rayleigh line is described by Equation 2.51, repeated below:

$$P_1 - P_0 = \rho_0 U_S u_p \, . \tag{2.60}$$

Note that in (u_p, P) space, the Rayleigh line is a straight line with a gradient of $Z = \rho_0 U_S$. If a shock travels through the material with a shock velocity U_S, then the final shock state will be the intercept between the Rayleigh line and the Hugoniot. This gradient Z is called the shock impedance, and is significant when considering shocks between different materials. It's also important to note that using Equation 2.51 requires all velocities to be in the stationary frame, where the material ahead of the shock is stationary[16].

2.8.4 Material interfaces and impedance matching

The previous sections considered shocks moving through a single material; in this section, the interface between materials is considered following the descriptions given in [49] and [50]. Imagine a shock S_1 travelling through material A, which then reaches an interface with material B. Some fraction of the shock energy will

[16] In the case shown in Figure 2.3 (b) where this material ahead of the shock has been pre-shocked, the 'stationary frame' is now moving so that the material ahead of the shock appears stationary. If the material has been pre-shocked to a pressure P_a and velocity u_a, and the new shock moving at U_S drives the material to a pressure P_b and velocity v_b, then the Rayleigh line would be given by $P_b - P_a = \rho_0(U_S - u_a)(u_b - u_a)$.

be transmitted across the interface into material B, while some will be reflected into material A. The pressure behind the reflected shock P_R is given [9] by

$$\frac{P_R}{P_1} = \frac{Z_2 - Z_1}{Z_2 + Z_1} \,. \tag{2.61}$$

Three illustrative cases are useful to consider. If the impedances of the two materials are very similar $Z_2 \approx Z_1$, the reflected shock pressure will be effectively zero and almost all the shock energy propagates into the new material. If there is a high impedance mismatch between the materials, then a large fraction of the shock energy is reflected. If the new material is higher impedance $Z_2 > Z_1$, then the reflected pressure is positive and a shock is reflected. If the new material is instead low impedance $Z_2 < Z_1$ then a rarefaction/relaxation wave is instead reflected.

After the shock has crossed the interface, the material on either side is in the state produced by the transmitted/reflected waves S_2 and S_3. For the two materials to remain in contact, it is required that the pressure P and particle velocity u_p of the material either side of the boundary must be equal - if this were not the case, the boundary would not be in equilibrium, or the two layers would separate. This enables impedance matching calculations, where the state in each material following the shock crossing the interface is determined.

This is shown graphically for the case where the second material is lower impedance in Figure 2.4. Two materials, A and B, are in contact with an interface between them. They are both initially unshocked, in states described as (0) and (0') respectively. The pressure and particle velocity of state (0) are described as (u_0, P_0) - as both (0) and (0') are unshocked and stationary, both P and u are in fact 0. Note that the velocity u_n in this section is used to refer to the velocity of the state labelled n, rather than suggesting that a particular reference frame is being used - in fact, it is simplest to do the calculation in the stationary frame.

An initial shock S_1 travels through material A shocking it to the state (1) with (u_1, P_1), which must by definition lie on the principal Hugoniot for material A. S_1 reaches the interface between the two materials, and this generates a rarefaction wave S_3 which travels back into material A, and a shock wave S_2 which propagates

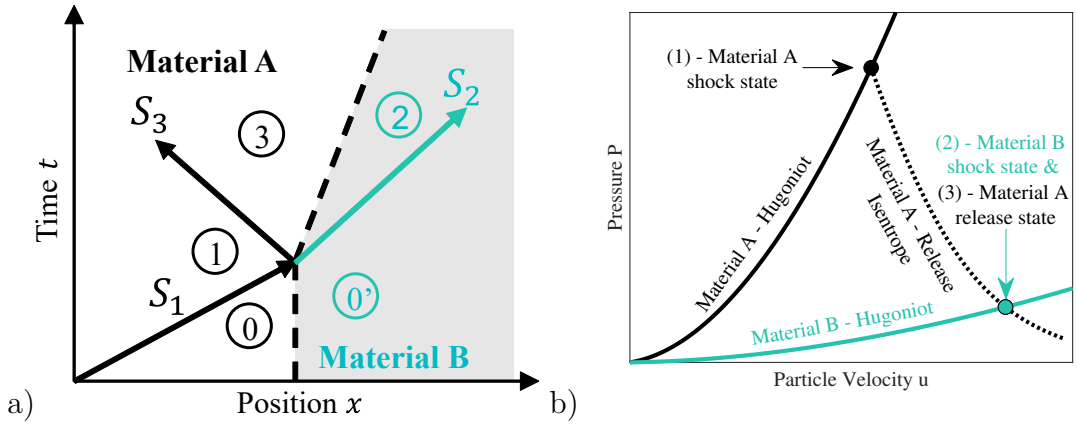


Figure 2.4: a) x-t diagram showing a shock transiting across the interface from a high impedance material into a low impedance one. The dashed line shows the material interface, while the arrows show the shocks. The shock/rarefaction waves are referred to as S_n, while the shock states are refereed to by the numbers. b) (u, P) plot showing the impedance matching graphically. S_1 shocks material in state (0) to a state (1) on the Hugoniot for material A. Rarefaction wave S_3 causes A to relax from (1) to a state (3) on the release isentrope for A. Shock S_2 shocks B from state (0') to state (2) on the Hugoniot of B. The impedance matching criteria means that (2) and (3) are equivalent in (u, P) space, and thus this state is the intercept of these two curves.

through material B. As material B is unshocked, the new state (u_2, P_2) lies on the principal Hugoniot for material B. A rarefaction wave S_3 travels through the pre-shocked state (u_1, P_1), and causes it to relax to a new state (u_3, P_3). This relaxation occurs along another curve known as the release isentrope, which describes the locus of states that the material could relax to given its shocked initial state of (u_1, P_1).

It is clear from Figure 2.4 (a) that after the shock has crossed into material B, the states either side of the interface are the shocked material B in state (2) and which lies on the B Hugoniot, and the relaxed material A in state (3) which lies on the A release isentrope. The impedance matching criteria require that both pressure and particle velocity must be constant across the interface, and as such $P_3 = P_2$ and $u_3 = u_2$. This condition is satisfied at the intercept of these two curves in (u, P) space, as shown in Figure 2.4 (b). Therefore, the intercept of these two curves[17] defines both states (2) and (3).

[17]There are two things to note here. 1) There is only one state which satisfies this condition. However, if the shock S_1 strength is changed, the state (u_1, P_1) is changed, and thus the release isentrope (which depends on this initial state) also changes - meaning that the impedance matching criteria where the Hugoniot B and release isentrope are satisfied now also occurs at a different

The general principle is the same if B has a higher impedance than A, but with a few changes. Here, the wave S_2 is now a shock wave rather than a rarefaction wave. This means that the state (3) is no longer a relaxed state on the release isentrope, but a new shock state which lies somewhere on a secondary Hugoniot of A (i.e. a Hugoniot generated for the initial conditions of state (1)). This also results in a pressure $P_3 = P_2$ which is greater than P_1. However, the impedance matching condition $P_3 = P_2$ and $u_3 = u_2$ still applies, and thus in this case the states (2) and (3) are instead the intercept between the principal Hugoniot of material B with this new secondary Hugoniot of material A (replacing the release isentrope).

2.8.5 Measuring shock states experimentally

It was shown earlier in this section that the Rankine-Hugoniot relations allow all the shock variables associated with a shock state to be calculated, if at least two of the variables (and the state of the unshocked material) are already known. 'Absolute' shock experiments aim to measure two such variables directly. However, impedance matching experiments make use of the impedance matching criteria to enable the shock state in a material to be determined by measuring only one variable (typically the shock velocity) in two different materials: the material of study, and a reference material. This is often simpler to perform experimentally.

Consider the situation described in the previous section, but with material A now being a known reference (which means that the relevant curves, such as the Hugoniots and the release isentrope, are already known) and material B being the material of study. Imagine that an experiment is conducted so that the shock velocity of S_1 in material A and of S_2 in material B are measured.

Measuring the shock velocity U_{S_1} allows the Rayleigh line for this shock to be plotted. As discussed in Section 2.8.3, the achieved shock state in material A will be the intercept of the Rayleigh line and the material Hugoniot. As A is a reference material where the Hugoniot is already known, this intercept can therefore

point. 2) While the (u, P) state of the two materials is equivalent, the Rankine-Hugoniot equations used to calculate the other variables reference the initial state of the material being considered - meaning that the other shock variables for these materials will not be equal.

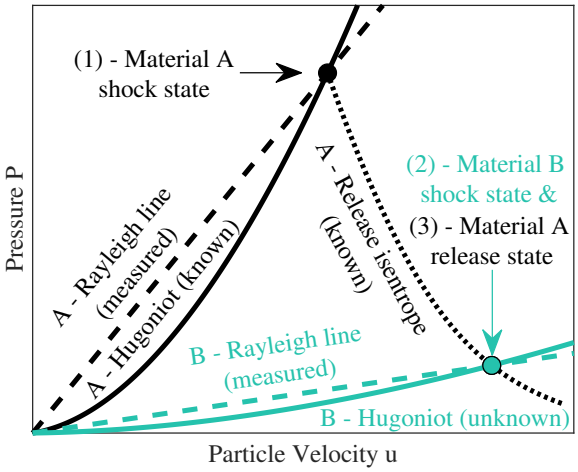


Figure 2.5: A (u, P) plot as in Figure 2.4, showing how impedance matching is used experimentally. Material A is a reference, where the Hugoniot (solid black line) and release isentrope (dotted black line) are known. Measuring the shock velocity of S_1 in A allows the Rayleigh line to be plotted (dashed black line) - the intercept of this with the known Hugoniot for A allows state (1) to be determined. Measuring the shock velocity of S_2 in B also allows a Rayleigh line to be plotted for this shock (dashed teal line). The intercept of this line with the known release isentrope of A allows the state (2), a shock state of B which lies on the unknown Hugoniot of B (solid teal line), to be calculated.

be identified, and the shock state (1) determined. This is shown in Figure 2.5. This shock state (1) can then be used to calculate the release isentrope[18] for material A. Again, as A is a reference, this isentrope is already known.

Measuring the shock velocity U_{S_2} enables the Rayleigh line for the shock S_2 in material B to be plotted, and the intercept with the Hugoniot for material B would define the shock state for that material. As material B is the material for which the shock behaviour is being investigated, this Hugoniot is not known. However, the impedance matching criteria mean that this is the same (u, P) state achieved in the relaxed (or re-shocked) material A, and thus also lies on the release isentrope for this material. The shock state (2) of material B is therefore identified from the intercept of the release isentrope for A and the Rayleigh line for B. This shock state is thus determined through only measurements of the two shock velocities, and knowledge of the reference material.

[18]This example is given again for where B is lower impedance than A. If B is higher impedance, replace all mentions of the 'isentrope' with the 'secondary Hugoniot', and any mention of 'relaxed' material A to 're-shocked' material A.

3

Definitions of key terms

Contents

3.1 Introduction **44**
3.2 Convergence ratio **45**
3.3 Hotspot and shell **46**
3.4 Hotspot / shell temperature **47**
3.5 Areal density **47**
3.6 In-flight aspect ratio **48**
3.7 Implosion velocity **49**
3.8 Laser intensity **49**
3.9 Yield, gain and bang time **49**
3.10 Adiabat **50**
3.11 Ignition and break-even **51**

3.1 Introduction

Many of the key terms that are ubiquitous in ICF, such as 'hotspot' and 'ignition', have no single agreed upon definition. This can make it difficult to compare results between different published works, and means that the use of such terms is ambiguous. This chapter will discuss the definitions for a range of terms and quantities that are used throughout this thesis, along with the details of exactly how these quantities were evaluated for the simulations presented in later chapters.

3.2 Convergence ratio

The definition of convergence ratio used in this work is

$$\mathrm{CR} = \frac{R_\mathrm{I}}{R_\mathrm{HS}} \,, \tag{3.1}$$

where R_I is the initial interior radius of the DT layer, and R_HS is the radius of the hotspot. This is the same definition as used in the works by Olson *et al.* [7] and Zylstra *et al.* [51], and was selected to ensure compatibility with their findings. This parameter is also referred to as the 'hot-spot convergence ratio', and remains a common definition of convergence ratio for wetted-foam research [6].

Other definitions for convergence ratio also exist. Some papers define the initial radius R_I as the 'initial inner radius of the shell' [41] or the 'initial radius of the DT layer' [52]; these are likely the same definition as discussed previously, but the phrasing is potentially ambiguous and could also apply to the DT/ablator interface. Alternatively, a 'capsule convergence ratio' is also used where the radius R_I is instead defined as the outer radius of the capsule, including the CH ablator [53]. Further ambiguity is inserted through the R_HS term. This is often stated as simply the hotspot radius [7, 6], but has also been described as the 'minimum' hotspot radius, the hotspot radius at peak compression [41], or the stagnated outer radius of the hotspot [51]. These terms are all interpreted to mean the minimum achieved hotspot radius, but it is noted that this does not necessarily occur at the the time of stagnation (itself a poorly defined term) or peak compression. More significant is the ambiguity in how the hotspot radius is identified, as discussed in the following section.

The convergence ratio is often evaluated in simulations where alpha deposition has been turned off [41]. This is because alpha deposition leads to increased hotspot pressure, which in turn will decrease the convergence ratio of the capsule. However, if the experimental performance does not match that of the simulations, then this pressure is reduced and the capsule converges further before burn begins; and thus using the value derived without alpha deposition represents a worst-case scenario. In this thesis, convergence ratio was calculated with alpha-heating included. This was

done for convenience during the capsule optimisation in Chapter 4, since otherwise two simulations would have been required for each set of parameters (one with alpha-heating to enable calculation of the gain, and one without to enable calculation of the convergence ratio). The focus of this work was largely on achieving break-even, and at these gains the discrepancy between burn and no-burn parameters is minimal (this effect only starts to become significant for gains of around 10 and above [54]).

3.3 Hotspot and shell

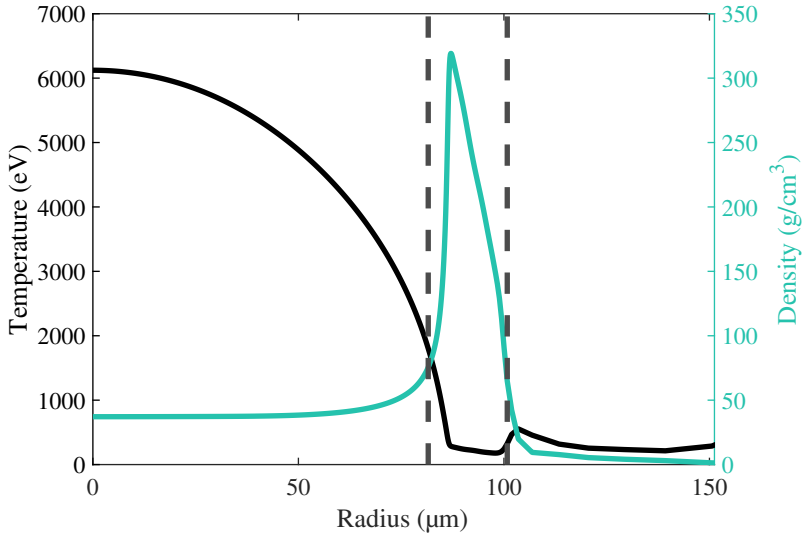


Figure 3.1: Ion temperature (black) and density (teal) at the time of maximum compression, for an implosion simulated in `HYADES`. The two dashed lines indicate the position of the hotspot and shell radii using the definitions defined in the text.

The term 'hotspot' is used to refer to the low density, high temperature region at the centre of the ICF capsule, while the 'shell' is the cooler high density region surrounding it. A plot of a capsule at peak compression, such as that shown in Figure 3.1, clearly shows the relevance of this description. However, where the exact interface between the two lies is not clearly defined. This is particularly true if one wishes to extend the concept of the shell to earlier stages in the implosion before stagnation (for instance, in order to calculate the implosion velocity). There is no accepted definition for this interface, and the definition used in a particular paper is not usually provided; the only definition found was in [6], which defined

the hotspot as the central region where the temperature was above 5 keV. This is not a useful definition for small non-igniting capsules, which despite showing similar temperature and density behaviour to Figure 3.1 do not achieve such temperatures anywhere in the capsule. This definition also poorly describes the capsule in Figure 3.1, where quite a large region of low density material would be considered part of the shell. The definition of these parameters is particularly important, since many other terms (including the convergence ratio) are defined with reference to them.

As a result, a new definition for hotspot/shell was identified for this work. The hotspot/shell interface (or 'hotspot radius') is defined as the radius at which the density has increased by $1/e^2$ of the difference between the central hotspot density and the peak density in the shell. The shell outer radius is defined as the radius where the density has decreased to $1/e^2$ of the peak value. In Figure 3.1 this places the two boundaries at the dashed lines, which matches well with the qualitative description. As this definition is based on density, it works throughout the early stages of the implosion - defining a 'shell' which is then used for the calculation of other quantities.

3.4 Hotspot / shell temperature

The hotspot and shell temperatures referred to in this thesis are the mass-averaged ion temperatures over the hotspot and shell zones respectively. This (along with the definition for convergence ratio) is a clear example of how defining the hotspot and shell influences other important parameters.

3.5 Areal density

Despite areal density typically being referred to as ρR, which suggests that it is the product of the density and the radius at a given position, the more accepted and commonly used version [29, 55, 56] of this quantity is

$$\rho R = \int_{R_1}^{R_2} \rho \cdot \mathrm{d}R \,, \tag{3.2}$$

the integral of the density ρ with respect to radius R over the region being considered[1]. Definitions for areal density are often not provided, and some works do indeed discuss ρR as the product of these two quantities [15].

In this thesis, ρR refers to the integral version of this quantity. The major benefit of this definition is that the areal density of different regions sums linearly. This means that the total capsule areal density is the sum of the hotspot areal density and the shell areal density, which makes intuitive sense. Areal density can thus be thought of as a measure of how much DT fuel an alpha particle must move through, when generated at the centre of the capsule and travelling (in a straight line) outwards - which highlights why this quantity is important for alpha absorption, self-heating, and thus gain. It also correlates well to the experimentally determined 'neutron averaged' areal density, which is an average of the areal density across the neutron line of travel between the point of generation and the detector.

3.6 In-flight aspect ratio

The definition for the in-flight aspect ratio of the capsule used in this work is

$$\text{IFAR} = \frac{R_{\text{shell}}}{t_{\text{shell}}} , \tag{3.3}$$

where R_{shell} and t_{shell} are the shell radius and thickness. This quantity is evaluated when the capsule has been compressed to two-thirds of its initial radius, in accordance with the common definition provided in the review article by Craxton *et al.* [41] (other works evaluate the IFAR at different times, such as when the shell has converged by a factor of one third [57]). In the analysis performed in these simulations, the IFAR is evaluated at the point when the shell outer radius (as defined previously) reaches 2/3 of the initial inner radius of the wetted-foam layer. The shell thickness is defined as the shell radius minus the hotspot radius. It should be noted that some works define the shell radius as the position of peak shell density - the definition for R_{shell} used in this thesis is larger, and results in a slightly larger IFAR value than using this alternative definition.

[1]My first paper on this work defined ρR simply as the mass averaged product of these two quantities, which was replaced with the integral form for subsequent papers.

3.7 Implosion velocity

The term 'implosion velocity' is typically used to refer to the peak velocity achieved as the capsule implodes. The implosion velocity used in this work was a volume averaged velocity across the material used in the shell. This should more correctly be a mass-averaged implosion velocity, but the difference in these quantities is marginal.

3.8 Laser intensity

The laser intensity forms an important part of the criteria in Chapter 4 for the low instability regime. However, laser intensity is a function of the area over which the laser acts; if the same power is absorbed over a smaller surface area (i.e. as the capsule radius decreases) the intensity increases.

In this thesis, the intensity compared to the low instability criteria is calculated using the initial capsule radius. This was done for convenience; since the initial radius of the capsule is known in advance, the laser power can thus be chosen in order to ensure an appropriate intensity (when using this definition). However, it is also reasonable to assume that the time-dependent intensity at the critical density will not increase much beyond this value, as while the radius decreases the coronal plasma would be expected to decrease the power of the incoming laser through absorption and scattering processes.

3.9 Yield, gain and bang time

Yield (or 'fusion yield') throughout this thesis is used to refer to the neutron yield, or the energy carried by the neutrons that will escape the fuel assembly and could thus potentially be captured. `HYADES` outputs a 'total yield' which includes the energy of the alpha particles, and the neutron yield is 80% of this value. The number of fusion-generated neutrons can be obtained by dividing the neutron yield by the 14.1 MeV neutron energy [15].

The term 'gain' is used throughout this thesis to refer to the 'total' or 'target' gain, which is defined as

$$G = \frac{Y_{\text{tot}}}{E_{\text{las}}} \, , \tag{3.4}$$

where Y_{tot} is the neutron yield and E_{las} is the input laser energy. Alternative variants of gain include capsule gain ($G_{\text{cap}} = Y_{\text{tot}}/E_{\text{abs}}$, which compares the yield to only the energy absorbed by the capsule E_{abs}) and fuel gain ($G_{\text{fuel}} = Y_{\text{tot}}/E_{\text{fuel}}$, which compares yield instead to the energy absorbed by the DT fuel E_{fuel}), both of which are significantly larger [15, 56].

The bang time of the implosion is defined as the time of maximum neutron generation (i.e. where the peak number of fusion reactions occur) [41].

3.10 Adiabat

The adiabat (or isentrope parameter) of an ICF capsule is used as a measure of the entropy of the cold, highly compressed DT fuel. This is done through comparison with a cold gas with Fermi-degenerate electrons. It is defined as

$$\alpha = \frac{P(\rho, T)}{P_{\text{deg}}(\rho)} \, , \tag{3.5}$$

where $P(\rho, T)$ is the pressure of the DT fuel for which the adiabat is calculated, while $P_{\text{deg}}(\rho)$ is the pressure for the Fermi-degenerate DT at the equivalent density. This can be calculated using

$$P_{\text{deg}}(\rho) = A_{deg}\rho^{5/3} \, , \tag{3.6}$$

where ρ is the density in $\text{g}\,\text{cm}^{-3}$, and $A_{deg} = 2.17 \times 10^{12}$ $(\text{erg/g})/(\text{g/cm}^3)^{2/3}$ for equimolar DT [15].

There are multiple ways such a definition can be evaluated, and this is discussed in [58]. The adiabat is most commonly quoted at the time of peak implosion velocity [59], but it can also be evaluated post-stagnation. The quoted adiabat value is either a mass-weighted average, or the minimum adiabat in the DT fuel at that time. If it is an average, the range over which the average is taken also influences the results.

An alternative way of calculating the adiabat is also provided in [58]. Here, they instead calculate a mass-averaged entropy at the peak implosion velocity. The equation of state for the DT is then used to calculate the pressure for DT at 1000 $\mathrm{g\,cm^{-3}}$, on the isentrope defined by that particular entropy value. This calculated pressure is referred to as P_{cold}, and replaces the Fermi-degenerate pressure $P_{\mathrm{deg}}(\rho)$ in the Equation 3.5. This is because the cold DT fuel at typical densities achieved in the shell at peak implosion velocity is only partially ionised, and thus P_{cold} is lower than $P_{\mathrm{deg}}(\rho)$. The pressure definition of the adiabat does not well account for the density dependence of the DT entropy, and gives behaviour that is non-monotonic in entropy [60]. Another alternative definition is to again calculate the mass-averaged entropy of the DT fuel, and then use this value in a fitting formula to calculate the adiabat directly [58, 61].

Adiabat does not play a significant role in this thesis, and adiabat values are not reported. In part, this was due to difficulties in this calculation; the adiabat value calculated from Equation 3.5 was found to be dependent upon the selection of zones used. `HYADES` does not return the entropy of the fuel, which means that the alternative definitions could not be implemented.

3.11 Ignition and break-even

Of all the ICF parameters, 'ignition' is probably the most ambiguous. The National Research Council in 1977 defined ignition as the point at which the fusion energy produced exceeded the laser input [62]; this is easy to understand and describes an important milestone in fusion research, but is not linked to any major physics transition/threshold within the capsule. There are both alternative definitions (for example, it can be defined as the point at which a burn wave starts to propagate through the cold shell [63]), and multiple ways that such definitions can be quantified. When the NIF shot N210808 achieved a yield of 1.3 MJ, this ambiguity meant that it had to be evaluated against nine different criteria in order to claim 'ignition' had been achieved [56]. It was able to satisfy these criteria with high confidence despite not achieving the NIC criteria (which was finally satisfied by the December 2022 shot).

Due to this ambiguity and lack of single clear definition, the term 'ignition' has been avoided where possible in this thesis. Where it is used, it is done so causally and simply means to achieve a well performing capsule with high gain. The term 'break-even' is used instead, which has an agreed upon definition of a gain of unity (meaning that it is the same as the NIC ignition criteria, but without the ambiguity).

4

Low convergence ratio implosions

Contents

4.1 Introduction 53
4.2 Previous research into wetted-foam capsules and low convergence ratio implosions 55
4.2.1 Experimental results 55
4.2.2 Simulation work 56
4.2.3 Implications of these previous findings 58
4.3 Defining a 'low instability' regime 59
4.3.1 The four 'low instability' criteria 63
4.4 Capsule design and laser pulse-profile 64
4.5 Performing the simulation campaign 65
4.5.1 The simulation code `HYADES` 65
4.5.2 Capsule meshing 66
4.5.3 Conducting the simulation campaign 67
4.5.4 Hydroscaling 69
4.6 Results and analysis 69
4.7 Limitations and suggested future work 73
4.8 Conclusion 75

4.1 Introduction

Implosions using wetted-foam capsules provide an opportunity to explore how fusion performance varies with convergence. The use of liquid DT rather than DT ice means that capsules are fielded at higher temperatures (between 20 and 26 K), resulting in

higher densities in the central vapour region (between 0.6 and 4.0 $\mathrm{mg\,cm^{-3}}$) [7] and thus reduced capsule convergence. The convergence ratio of wetted-foam implosions is thereby controlled through the initial temperature of the capsule.

In an ideal ICF implosion, increasing the convergence ratio results in higher areal densities and increased gain. However, it is well established that increasing the convergence ratio also leads to increased hydrodynamic instabilities, which prevents these high gains from being achieved experimentally. At low convergence ratio the ideal performance is lower, but such instabilities are reduced: recent low convergence ratio wetted-foam shots on the NIF have achieved low instability growth, resulting in good agreement between experiment and simulation [7, 51].

In this chapter, one-dimensional (1D) simulations are performed to explore the possible fusion performance that could be obtained using low convergence ratio implosions of wetted-foam capsules. First, a range of relevant previous results are discussed in Section 4.2. An expected low instability regime is then identified (based on limits to convergence ratio and other key implosion parameters) so that 1D simulations are expected to be a reasonable estimate of experimental performance, in Section 4.3. A standard laser pulse sequence and capsule design are chosen and outlined in Section 4.4. These contain a number of 'free' parameters, which are optimised for a number of different capsule sizes during a large-scale simulation campaign. Section 4.5 outlines how this simulation/optimisation campaign was performed, and the resulting 'optimal' implosions over a range of different energies are presented and discussed in Section 4.6. Finally, the limitations of this approach are considered in Section 4.7, along with ways that it could be further developed in the future.

The purpose of this approach was to provide a first exploration of this parameter space, to determine whether IFE-relevant gains could be obtained with such implosions and to estimate the energy required for key milestones such as break-even or a reactor level gain. The use of 1D simulations means that the predictions made are not likely to be exact. The performance achieved in such a regime will not match the high gains that are routinely simulated for higher convergence ratio implosions; however, if significant performance is obtained in these simulations,

the low instability nature of the regime and the greater than typical expected agreement between simulation and experiment suggests an interesting possible direction for IFE research.

4.2 Previous research into wetted-foam capsules and low convergence ratio implosions

4.2.1 Experimental results

The first DT-wetted-foam shot on the NIF (designated N160421) was performed in 2016, and was reported by Olson *et al.* [7]. $(4.5 \pm 0.1) \times 10^{14}$ neutrons were generated, which suggested that a good level of performance was achieved with such implosions. More impressive however was the close agreement between the simulation and experiment. 2D simulations were able to match a wide range of implosion metrics (including temperatures, bang time, hotspot radius and inferred pressures) to within the experimental error. These metrics were also very well described even by a 1D `RAGE` simulation (with a mix model included), which gave a ratio between experimental and simulated yield of nearly 80%.

N160421, and subsequent wetted-foam shots N160626 and N161204, were then discussed by Zylstra *et al.* [51] in 2018. These three shots spanned a convergence ratio range of 12-20, and demonstrated the ability to control convergence ratio through initial capsule temperature. These shots were compared to simulations using the `HYDRA` and `xRAGE` codes. The agreement was very good at low convergence ratio (with yield and ion temperature agreeing to within 5% for N160421, which had a convergence ratio of 12), but as convergence increased the level of agreement was reduced, resulting in an experimental yield for the CR=20 shot that was only 30% of that predicted by the simulation. The intermediate shot, with a convergence ratio between 16 and 17, showed a good agreement of above 70%[1]. These NIF shots therefore showed that wetted-foam layers do indeed result in low convergence ratio implosions, and that the agreement with simulations at these low convergence ratios are relatively high.

[1]It should be noted that the exact reason for these discrepancies could not be identified in [51].

This was not the first time a link between low convergence ratio and agreement with simulation was observed; the same result is highlighted in the following (non-exhaustive) collection of references: [64, 65, 66, 67, 53, 68, 69, 52]. This agreement was identified as far back as 1980 [70], where it was observed that lower initial aspect ratios (thicker shells as a function of radius, resulting in lower implosion velocities and less convergence) resulted in performance that was more 1D-like. Two developments have led to increased recent interest in low convergence ratio implosions. Firstly, higher laser energies make it more feasible to achieve high gain in such a regime. Secondly, wetted-foam capsules (where the convergence ratio can be controlled through capsule temperature) mean that the convergence ratio of an implosion can be changed without substantial changes to the capsule, allowing it to be varied almost independently of other parameters.

While the NIF shots discussed so far were indirect-drive, similar links between convergence ratio and 1D-like performance have also been observed in direct-drive shots on OMEGA. Regan *et al.* observed that the hotspot pressure and compressed areal density of capsules using DT ice layers (rather than wetted-foams) reached 90% of the 1D simulated values for convergence ratios below 17 and adiabat greater than 3.5 [71]. When the convergence ratio was increased past 17 a rapid drop-off occurred, which shows good qualitative agreement with the trend in yield observed on the NIF shots. This shows that this relationship applies for both indirect and direct-drive implosions, using both DT ice or DT-wetted-foam layers[2].

4.2.2 Simulation work

This topic has also been explored through simulation work. 2D simulations were previously used to investigate how the yield of wetted-foam capsules changed with varying vapour density (and thus convergence ratio), for different amplitudes of a Legendre P_4 flux asymmetry. As expected, they found that the higher vapour density/lower convergence ratio implosions were significantly more robust to this

[2]It is also worth noting that that a low convergence ratio, high adiabat design using thin shells at high implosion velocities [72] recently achieved the facility record neutron yield at OMEGA. These results were achieved after the work discussed in this chapter was performed.

asymmetry, with a higher ratio of 2D to 1D yield [16]. Simulations performed in preparation for the NIF shots suggested that sensitivity to capsule instability growth and x-ray flux asymmetries increased with convergence ratio, leading to strong agreement between 2D and 1D simulations for CR=15 shots when surface roughness was applied, compared to much worse agreement for high convergence ratio shots. This led the authors of that work to suggest that using low convergence ratios would lead to reduced instability growth and a robust predictive capability [73].

High resolution 2D `xRAGE` simulations performed in 2017 investigated full-scale wetted-foam capsules with a variety of 'asymmetry seeds'. These simulations showed that when surface roughness and drive asymmetries were included, yield increased with convergence ratio through the range of 12<CR<20. However, when the capsule support tent and fill tube were also included the yield was roughly constant as a function of convergence ratio up to CR = 20, at which point it began to decrease [52]. The capsule support tent was found to be the main seed for this instability growth. Despite this, the paper concluded that "as the level of these asymmetries is reduced, the wetted-foam platform will enable layered implosions to be fielded at convergence ratios that optimize the trade-off between enhanced 1D performance and increased implosion instability as the convergence ratio is increased".

There are some differences between the simulations presented in [52] and the designs considered in this chapter, but these are unlikely to lead to significant changes in behaviour. For instance, while the use of a capsule support tent is specific to indirect-drive, direct-drive implosions currently require stalks to hold the capsule in place[3] which will likely also have a significant effect. The effect of the stalk on direct-drive implosions is still being investigated; simulation work on this topic has suggested stalks are responsible for a yield degradation of 10-20% [74], while more recent simulations indicated the impact of the stalk was significant but were unable to quantify it exactly [75]. The other key difference between the low

[3]While stalks are currently required to hold the capsule in position, this is unlikely to be compatible with a fusion reactor, where a higher frequency of shots would be required. In such a reactor, the target would likely be dropped into the chamber and shot while in free-fall, removing the need for a target support structure [41].

convergence xRAGE simulations in [52] and the designs considered in this chapter is their use of a high density carbon (HDC) ablator. Such ablators are typically more susceptible to the effects of the fill tube, but may be more robust to the capsule support tent compared to CH [76, 56, 39, 77].

A key observation from the xRAGE simulations in [52] was the impact of the first shock driven through the capsule. If the first shock is not strong enough to melt the HDC ablator, then then the imprint from the support tent (and the HDC microstructure, which was not considered in [52]) may seed significant instability growth [78, 39]. The xRAGE simulations had a weak first shock, which was likely why the support tent led to such significant yield degradation [52]. Other xRAGE simulations using stronger initial shocks led to the support tent contributing to only a 1% reduction of the 2D yield [39], highlighting that this can indeed be well mitigated using a strong first shock (this was also seen in the xRAGE simulations of N160421, the first liquid-layer DT shot, with a higher adiabat). A number of developments since the 2017 xRAGE simulations were performed have also been implemented to reduce the amplitude of these asymmetry seeds. Alternative designs for supporting the capsule have been proposed [79, 80] and have led to improved capsule performance [80]. In addition, the fill tube diameter has been significantly reduced since the 2017 simulations, leading again to less yield degradation [81, 82]. 2019 xRAGE simulations of a wetted-foam capsule with a plastic shell at CR = 19 showed good agreement with experiment [39], suggesting that these instabilities can indeed be controlled (even at higher convergence ratios) if asymmetry seeds are well mitigated.

4.2.3 Implications of these previous findings

The significance of these previous results are clear. The use of low convergence ratio has been demonstrated in both simulations and experiments to reduce instability growth, and as convergence ratio is increased yield degradation due to instabilities is also observed to increase. Detailed xRAGE simulations investigating the impact of convergence ratio [52] have demonstrated that low convergence alone is not sufficient to mitigate instability growth, but more recent simulations (and significant research

effort to reduce the seeds of these instabilities) indicate that it is possible to control asymmetry seeds and thus minimise the resulting instabilities. The 1D simulations in this chapter do not include the effects of these asymmetry seeds, and thus it is important to note that these are also key to minimising instability growth.

4.3 Defining a 'low instability' regime

The simulation campaign presented in this chapter explores the performance of wetted-foam capsules using one-dimensional radiation-hydrodynamic simulations. 1D simulations are useful for this style of campaign due to their quick run time (a few hours each), which allows them to be performed in large numbers to effectively explore a parameter space. However, they do not include instability growth (which requires higher dimensions to describe), and thus for such simulations to be accurate it is essential to operate in a regime where such growth is expected to be minimal. In this section, a series of 'criteria' are identified which correlate with low instability growth; by identifying implosions which satisfy all of these, instability growth is assumed to be minimised and thus 1D simulations are expected to be reasonably accurate.

The primary way this was achieved was through considering low convergence ratio implosions. For this work, a definition for low convergence of CR < 16 was used. This value was chosen as an expected reasonable compromise between two competing effects; the expected increased 1D fusion yield with convergence ratio, and the decreased agreement between experiment and simulation. A CR of 16 would likely give a much higher 1D yield than the lowest CR of 12 discussed by Olson *et al.* [7], but still fits within the CR range where Olson *et al.* [7] and Zylstra *et al.* [51] observed reasonable agreement with simulation[4].

A number of other implosion parameters are also known to influence instability growth, and so additional criteria based around these were also applied in order to

[4]The `xRAGE` simulations in [52] predicted that such a CR value would fit in the range where yield was flat with performance; this would predict that there was no advantage or disadvantage compared to a lower CR value, but as discussed previously it is likely that the level of instability growth seen in that work could be mitigated. This would result in a more positive trend between CR and performance.

further ensure that 1D simulations would be reasonably accurate. The key parameters considered for this work are discussed in this section, and are described in further detail by Craxton *et al.* [41]. As for convergence ratio, these parameters typically also have a significant impact on the overall fusion performance of the implosion.

In-flight aspect ratio (IFAR), defined in Chapter 3 as the capsule radius divided by the shell thickness during the implosion, is also highly important to the growth of hydrodynamic instabilities. By making some assumptions about the nature of the implosion, it is possible to estimate the growth-rate G of RT instabilities at the capsule outer surface during the acceleration of the shell in terms of IFAR,

$$G_l = \exp\left[\alpha_1 \left(\frac{l}{1+\alpha_2 \cdot l \cdot (\text{IFAR})^{-1}}\right)^{1/2} - \alpha_3 \cdot l \cdot (\text{IFAR})^{-1}\right] , \qquad (4.1)$$

where l is the spherical harmonic mode number, and the three α values are numerical constants [15]. This clearly highlights the significance of a low IFAR value (i.e. a thicker capsule shell for a given capsule radius) for mitigating instability growth. A criteria of IFAR < 30 was thus selected for this work; this value was based on previous works which identified IFAR $= 30$ as a rough threshold above which significant instability growth and risk of capsule break-up was observed [83] (although it should be noted that some works suggest that even lower IFAR values may be preferable to minimise instability growth [57, 84]).

Another key parameter is the peak implosion velocity of the capsule. Implosion velocity does not directly appear in the equations for hydrodynamic instability growth, but does indirectly have an impact through the acceleration [15]. It is also routinely reported for experimental shots, making it a useful comparison to ensure that the simulated shots are feasible; large implosion velocities are desirable for achieving ignition (with Cheng *et al.* reporting a scaling between yield Y, ion temperature T and implosion velocity v of $Y \propto T^5 \propto v^{10}$ [85]), yet these high velocities are challenging to achieve experimentally. A value of 400 $\mathrm{km\,s^{-1}}$ was chosen as an upper limit for implosion velocity for this work, which is comparable to the velocities that have been achieved in previous experiments[5] [41, 87].

[5]The recent N210808 shot was simulated to have an implosion velocity of 391 $\mathrm{km\,s^{-1}}$, which again demonstrates that 400 $\mathrm{km\,s^{-1}}$ is an appropriate choice [86].

The other parameter which is known to significantly influence hydrodynamic instability growth is the adiabat of the fuel. Keeping the cold DT fuel in the shell on a lower adiabat means it is more compressible, which results in higher areal densities and improved efficiency of energy transfer to the hotspot. Cheng *et al.* describe that for a 'high' adiabat implosion where $\alpha = 3$ the internal energy delivered to the hotspot is approximately half of that for a 'low' adiabat $\alpha = 1.5$ implosion, which results in a significantly lower hotspot mass and pressure [85] (two variables that are crucial to achieving ignition [88, 89, 18]). However, the higher ablative velocities achieved at higher adiabat are key for the suppression of hydrodynamic instability growth [90]. Initial implosions on the NIF which targeted very low adiabat achieved poor experimental performance [32], and significant improvement was obtained using higher adiabat 'high-foot' implosions [33]. This effect is coupled with the 'strong first shock' discussed in the `xRAGE` simulations in [52], as altering the strength of the first shock is the most significant way of changing the adiabat [91, 52, 39]. The importance of this parameter has also given rise to adiabat shaping techniques, where the outer fuel layer is placed on a higher adiabat than the inner fuel; doing so minimises the instability growth at the outer edge, while keeping the inner fuel on a lower adiabat to maximise compressibility [84].

No limit on the adiabat was set in this work. Firstly, the adiabat is particularly ambiguous (as discussed in Chapter 3) in terms of how it is calculated and reported, which makes comparison with literature values (and implementation in the analysis code) challenging. Secondly, the main technique used to influence the adiabat is changing the strength of the shocks, which is a parameter which was treated as fixed in this simulation work (as discussed later, in Section 4.4). It is worth noting that many of the parameters discussed here, including the adiabat, are known to be related (such as implosion velocity and convergence ratio, or convergence ratio and the adiabat [88, 92]), and thus it is also possible that a high adiabat may arise as a consequence of the limits on other parameters[6]. The lack of focus on the adiabat

[6]Previous research has also suggested that the combination of multiple parameters is important for 1D like behaviour; for instance, requiring low IFAR and high adiabat [89], or high adiabat and low convergence ratio [84]. For the approach taken here, such limits were not used and each

(or first shock strength) is a potential limitation of this work, and is something that could be addressed in further simulation campaigns, as discussed in Section 4.7.

The significance of many of these parameters to achieving strong fusion performance is highlighted in performance metrics, such as the ignition threshold factor (ITF) [58]. The ITF is designed to vary smoothly around the ignition threshold (unlike yield/gain, which display a cliff-edge behaviour) and to indicate how small changes to key parameters influence the likelihood of ignition. It is given by

$$\mathrm{ITF} = I_0 \left(\frac{M_{\mathrm{DT}}}{M_0}\right) \left(\frac{v}{v_0}\right)^8 \left(\frac{\alpha}{\alpha_0}\right)^{-4} \left(1 - 1.2\frac{\Delta R_{\mathrm{HS}}^{\mathrm{K-wtd}}}{R_{\mathrm{HS}}}\right)^4 \left(\frac{M_{\mathrm{clean}}}{M_{\mathrm{DT}}}\right)^{0.5} (1 - P_{\mathrm{HS}}) \ , \tag{4.2}$$

where I_0 is the ITF of a baseline optimised 1D implosion, M_{DT} is the fuel mass of the capsule, v is the implosion velocity, α is the adiabat of the main fuel, $\Delta R_{\mathrm{HS}}^{\mathrm{K-wtd}}$ is the Kishony-weighted RMS deviation from the the average hotspot radius R_{HS}, M_{clean} is the DT mass not contaminated by mixing, and P_{HS} is a measure of any impurity of the hotspot. ITF is designed so that a value of unity represents break-even, with M_0, v_0 and α_0 the nominal values required to achieve this for the particular design being considered. The strong dependence of ITF on both the adiabat and the implosion velocity clearly indicates the importance of these parameters to achieving ignition (and thus how the criteria applied in this work will necessarily limit the performance). This equation treats implosion velocity and adiabat in an ideal sense, and thus does not capture how instability growth may increase as these two parameters are varied (the effects of instability growth are instead described separately through the perturbation and mixing terms).

Also of relevance to this discussion is the recent statistical mapping model produced using a large database of OMEGA implosions [93], which highlights how a number of parameters contribute to yield degradation. This model was published after the work done in this chapter was performed, and so was not used in identifying the instability criteria. This mapping model included a term related to hydrodynamic instability, which included IFAR and adiabat (linked through the relation identified in [89]), and convergence ratio. It also quantified that the

variable was considered independently against the outlined criteria.

agreement between experiment and simulation scaled with convergence to the power of -0.97. This model therefore clearly shows how the limits imposed in this work (particularly on CR) lead to a greater predictive capability, and thus better expected agreement between simulation and experiment.

Finally, while this discussion has so far related to hydrodynamic instability growth, parametric instabilities arising from laser plasma interactions are also not described in the simulations and must therefore be minimised. Such instabilities lead to decreased laser absorption, and can also result in the generation of hot electrons which preheat the fuel, making it harder to compress and reducing the yield [42, 43]. A fourth criteria was set,

$$I \cdot \lambda^2 < 1 \times 10^{14}\,(\mathrm{W\,cm^{-2}}) \cdot \mu\mathrm{m}^2\ , \tag{4.3}$$

where I is the laser intensity and λ is the wavelength. This is a threshold above which significant levels of parametric instabilities begin to be observed [94].

4.3.1 The four 'low instability' criteria

The four criteria outlined in the previous subsection are used to define a regime where instability growth is expected to be minimised, and thus 1D simulation is expected to be a reasonable estimate of performance. The optimisation campaign was conducted so as to produce implosions with the highest possible gain while satisfying all four of these criteria, in order to ensure this expected level of accuracy. Repeated here for clarity, these four criteria are:

- CR < 16 ;
- IFAR < 30 ;
- $v_{imp} < 400\,\mathrm{km\,s^{-1}}$;
- $I \cdot \lambda^2 < 1 \times 10^{14}\,(\mathrm{W\,cm^{-2}}) \cdot \mu\mathrm{m}^2$.

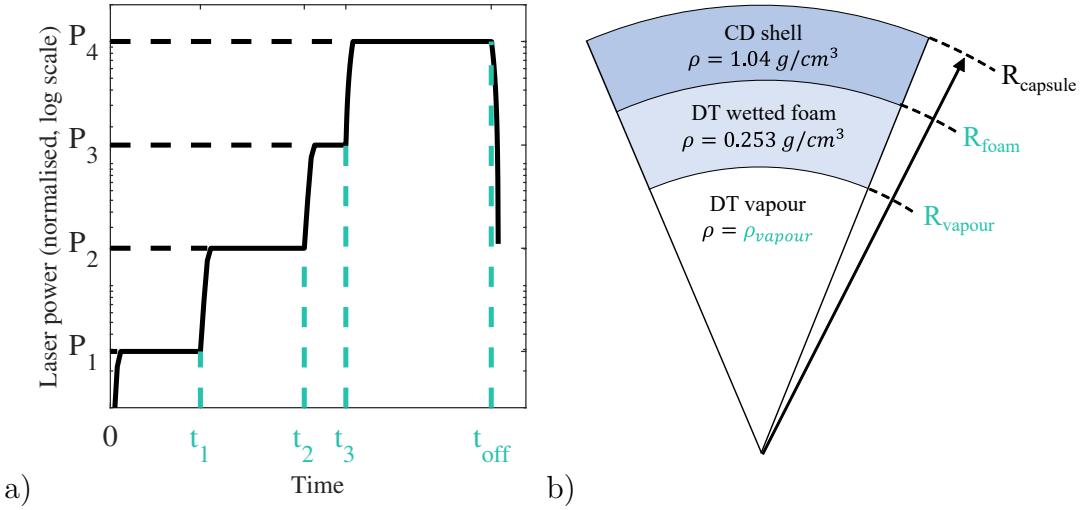


Figure 4.1: Plots of a) the laser profile and b) the capsule design. The laser sequence consists of four pulses of increasing power. The power begins to increase at the times $0, t_1, t_2$ and t_3 (with a linear rise over 200 ps), and the laser is switched off at time t_{off} (with a linear decay over 200 ps). The capsule has three layers, with material interfaces at R_{vapour} and R_{foam} and an outer radius of R_{capsule}. When optimising this design, R_{capsule} is selected and considered fixed. This is used to define the laser pulse powers P_1, P_2, P_3 and P_4, which are also not changed. This leaves seven optimisation parameters, which are indicated in teal: $t_1, t_2, t_3, t_{\text{off}}, R_{\text{vapour}}, R_{\text{foam}}$ and ρ_{vapour} (the DT vapour density).

4.4 Capsule design and laser pulse-profile

In order to perform the implosion optimisations, a general capsule design and laser pulse sequence were developed. These contained a number of free parameters which were varied in the optimisation. This section describes the particular pulse profile and capsule design used.

The capsule, displayed in Figure 4.1, was a simple three layer design consisting of a central DT vapour region surrounded by a DT-wetted-foam layer and a deuterated plastic (CD) ablator. A series of optimisation campaigns were performed, for a number of different capsule outer radii (this parameter was not varied during a given optimisation campaign). The interior dimensions of the capsule (R_{vapour} and R_{foam}) were varied during the optimisation, as was the density in the vapour region (ρ_{vapour}) (which was restricted to the 0.6 mg cm^{-3} $< \rho_{\text{foam}} <$ 4.0 mg cm^{-3} range achievable with wetted-foam capsules). The wetted-foam layer had a fixed density of 0.253 g cm^{-3}, while the ablator had a density of 1.04 g cm^{-3}. This represents

a realistic capsule design similar to those used in previous experiments[7], with minimal complexity.

The laser pulse profile, displayed in Figure 4.1, consisted of a series of 'steps' of constant laser power. A 351 nm laser was used (corresponding to the third-harmonic of Nd:glass, as used at most ICF facilities), with a 0.2 ns linear rise time over which the changes in power occur. Both three- and four-pulse profiles were considered. The timings of each step (t_1, t_2 and t_3 in the figure) were treated as optimisation parameters, along with the time at which the laser was switched off (t_{off}). The relative laser powers were fixed and were not changed during the optimisation; ideally these would have also been varied, but the addition of further optimisation parameters would have made the optimisation infeasible. The pulse powers used were chosen with the aim of minimising the entropy injected into the fuel (but as this was not investigated in detail or optimised, better pulse profiles are likely possible). This laser profile was chosen for ease of optimisation, but again presents a simple realistic pulse profile that is compatible with current facilities such as the NIF.

Optimisation campaigns were performed for a number of different capsule radii. For each campaign, the outer radius determined the peak laser power through the intensity criteria, and the pulse profile was scaled accordingly. The outer radius also significantly influenced the energy required for the implosions, since it determined both the laser power and the implosion time (due to the fixed maximum velocity). The three laser timings, laser switch off time, two internal capsule radii and DT vapour density formed the seven optimisation parameters.

4.5 Performing the simulation campaign

4.5.1 The simulation code `HYADES`

The 1D Lagrangian radiation hydrodynamics code `HYADES` was used to perform the simulations. `HYADES` is developed by Cascade Inc., and was ran on the

[7]The use of a central vapour region surrounded by a dense fuel layer is standard for ignition designs, but the choice of ablator can vary: these can include the use of HDC rather than plastic [78], or can contain layers of dopants (such as silicon) to absorb hot electrons and thus minimise preheating [95].

SCARF supercomputer at the UKRI-STFC Rutherford Appleton Laboratory for the simulations in this thesis. It has previously been used for other large-scale simulation campaigns of ICF implosions [96].

HYADES contains a number of different options for the various physics models. Those used in these simulations were based on the advice and recommendations of Dr Robbie Scott at the CLF, and many of these are discussed in Section 2.5. A multiplier of 0.8 was also applied to the input laser energy to account for losses (accounting for deviations from 1D spherical symmetry, and any losses from parametric processes such as SRS and CBET). A series of benchmarking simulations were performed after the main simulation campaign to confirm that the settings used were sufficient to describe experimental shots; these are described in the Appendix A.

The DT vapour regions were described in the simulations using the in-line HYADES QEOS model, while the thin deuterated-plastic shell used the SESAME EOS table for CH. These are standard ways of describing these materials. No specific equation of state table for DT-wetted-foam was available, and so this layer was instead described as a pure-DT layer at the wetted-foam density, using the DT QEOS model. This neglects the presence of the foam, and is a limitation of this work which is discussed further in later sections.

4.5.2 Capsule meshing

HYADES does not have an in-built meshing routine to discretise the capsule, and thus a meshing script had to be produced to allow different targets to be simulated. This mesh has two major requirements to ensure accuracy; the resolution of the mesh at the capsule outer edge should be sufficiently high, and the relative mass difference between zones should not be much greater than $\sim 2\%$. In addition to this, an increased number of zones/mesh points leads to an increased simulation time, and thus it is desirable to keep the number of zones to a minimum while satisfying these two conditions.

A custom meshing routine was written for this simulation campaign. This routine was automatic and robust, allowing meshes to be created to describe capsules of

a range of different sizes and layer thicknesses relevant to this investigation. The production of such a meshing algorithm was essential to being able to perform the large parameter scans that were the main feature of this work, as it enabled large numbers of capsules with different structures to be simulated with accuracy. The principles of this code (along with the underlying theory) is outlined in Appendix B.

The meshing algorithm worked sufficiently well for this campaign to be conducted. It operated less well as the thicknesses of the layers began to vary too drastically, but this did not prevent the optimisations being performed; if the thicknesses had changed drastically enough that for the meshing code to stop working well, the performance had always decreased to such an extent that it was clearly not worth investigating, or the criteria for implosion velocity/IFAR would have been exceeded. However, if simulations of drastically different designs are required in the future, then the development of a new and more robust meshing algorithm would be necessary.

4.5.3 Conducting the simulation campaign

Separate optimisation campaigns were performed to optimise a number of capsules with different outer radii. For each optimisation, a rough initial simulation was first performed. Large parameter scans were then performed varying either the two layer thicknesses, or two of the three laser timings, with roughly 100-200 simulations being performed in each scan. The results of these parameter scans were considered, and implosions displaying the most promising combinations of the gain and the instability criteria selected. A new round of simulations were then performed scanning over other parameters, and the process repeated. This was highly iterative - optimising the laser timings for example often led to new optimal layer thicknesses, which in turn led to a further change in the optimal laser timings. This process was repeated until no further improvement was obtained.

Throughout the optimisation, the vapour density was also investigated. For instance, if there were a number of capsules with lower yield than the current 'optimal' design, but which also had a much lower convergence ratio than the limit of 16, then the vapour pressure was reduced - this resulted in an increase to the CR

and gain of these capsules, which sometimes led to them outperforming the previous best result. Equally, if a large number of the capsules satisfied the other criteria but the convergence ratio was too high, then the vapour density was reduced. The vapour density was thus also optimised to provide the best performance.

The time at which the laser was turned off was also treated separately from the other optimisation parameters. At some late time in each implosion continuing to apply the laser results in no benefit to the fusion performance, but it continues to cost energy and thus causes the gain to decrease. This effect was often substantial, and trends in the gain were being obscured based on how long after this point the laser was applied for. To resolve this, the laser was applied until the end of the each simulation, and the gain calculation was adjusted so that only the laser energy used up until the bang time was included in the calculation. This meant that the influence of this laser switch-off time was removed. Once the implosion had otherwise been optimised, the laser switch off-time would then be varied and the true gain (using the full applied laser energy) was calculated. This removed the influence of this parameter until the final stage of the optimisation, when the optimal value for it was identified.

Care was also taken to ensure that appropriate time resolution was being used in all simulations, and that the simulation end time was appropriate. If the simulation ended before the capsule had properly converged, the gain was inevitably too low. This was typically spotted by unfeasible data (for instance, a very low convergence ratio). To test for this, a representative number of the simulations were analysed individually to check that the full behaviour was being captured, and that the simulations looked appropriate.

This procedure generated the 'optimal' design for the capsule radius and laser pulse profile being considered. It was repeated multiple times to generate the optimal 3-pulse and 4-pulse sequences for a range of different capsule sizes.

4.5.4 Hydroscaling

When this campaign was first conducted (in 2020) it was done as described above for each capsule, in an attempt to minimise the number of assumptions on the optimal implosion. However, it was clear that all the optimised capsules had relatively comparable thicknesses and timings. In 2022, a further optimisation was performed to fill in a gap in the energies investigated, and this time hydroscaling relations were applied to produce an initial capsule from which to start the optimisation. This significantly reduced the rounds of simulations required to achieve a well-optimised implosion.

Hydroscaling makes use of the fact that much of the physics in an ICF implosion (particularly in 1D) is scale independent, and thus many of the parameters scale in a straightforward manner with capsule size. These relations state that $t \propto R \propto E^{1/3}$ where t and R, are the implosion timings (including all laser timings) and capsule radii (including all internal boundaries) respectively, and E is the laser energy [97]. This meant that, based on the laser pulse energy desired for the implosion, a rough estimate was produced of the necessary capsule size based on the previous results. The interior dimensions of a previous optimal capsule were then scaled accordingly, as were the laser timings. This was then used as a starting estimate (which was already reasonably close to optimal) for the new optimisation campaign.

4.6 Results and analysis

Over 20,000 simulations were performed in total, with over 10,000 satisfying all four instability criteria; these ∼10,000 'valid' simulations are plotted as a function of gain vs laser energy in Figure 4.2. The data forms eight rough peaks, corresponding to the eight different outer radii (and thus energy scales) considered. The four-pulse laser sequences are displayed in black, while the three-pulse sequences are displayed in teal. As expected, the four-pulse sequences outperformed the three-pulse sequences at all energy scales. The best performing three-pulse and four-pulse simulations for each capsule size have been tabulated in Table 4.1.

Number of pulses	3							
Size multiplier	0.25	0.35	0.5	0.65	0.75	0.85	1	1.1
Laser energy (MJ)	-	-	0.80	1.8	-	4.0	6.6	8.9
Gain	-	-	0.11	0.35	-	5.2	26	38
Convergence ratio	-	-	15.3	15.5	-	15.9	16.0	16.0
IFAR	-	-	20.3	18.7	-	15.2	14.3	12.7
Implosion velocity (km/s)	-	-	399.2	395.4	-	373.9	358.4	332.4
Max pulse power (TW)	-	-	173.00	292.38	-	500.12	692.13	837.50
Pulse 2 switch on time (ns)	-	-	3.65	3.80	-	5.00	6.20	7.00
Pulse 3 switch on time (ns)	-	-	4.15	4.90	-	6.30	7.60	8.80
Laser switch off time (ns)	-	-	8.50	10.50	-	13.75	16.50	18.60
Vapour/liquid boundary (mm)	-	-	1.3000	1.6835	-	2.2100	2.6100	2.8710
Liquid/CD boundary (mm)	-	-	1.4000	1.8200	-	2.3800	2.7900	3.0580
Outer radius (mm)	-	-	1.4250	1.8525	-	2.4225	2.8500	3.1350
Vapour density (mg/cm^3)	-	-	1.00	0.85	-	0.65	0.65	0.70
Number of pulses	4							
Size multiplier	0.25	0.35	0.5	0.65	0.75	0.85	1	1.1
Laser energy (MJ)	0.10	0.27	0.77	1.7	2.8	4.2	6.7	8.5
Gain	0.030	0.067	0.19	0.75	6.1	24	42	54
Convergence ratio	15.7	15.7	16.0	15.8	16.0	15.9	16.0	15.7
IFAR	23.4	27.5	29.7	25.1	23.8	21.2	18.4	15.7
Implosion velocity (km/s)	391.4	395.8	399.6	399.6	387.8	347.5	334.1	311.0
Max pulse power (TW)	43.26	84.79	173.00	292.38	389.36	500.12	692.13	837.50
Pulse 2 switch on time (ns)	1.10	2.20	2.60	3.60	4.20	3.70	5.20	5.70
Pulse 3 switch on time (ns)	3.00	4.60	5.60	7.80	8.80	9.40	12.80	15.00
Pulse 4 switch on time (ns)	3.60	5.50	6.80	9.50	10.90	11.40	15.30	18.00
Laser switch off time (ns)	5.80	8.50	11.00	15.00	17.65	19.40	24.50	27.50
Vapour/liquid boundary (mm)	0.6325	0.8950	1.3050	1.6705	1.9275	2.2100	2.6000	2.8270
Liquid/CD boundary (mm)	0.69625	0.976	1.3950	1.8200	2.1000	2.3630	2.7800	3.0580
Outer radius (mm)	0.7125	0.9975	1.4250	1.8525	2.1375	2.4225	2.8500	3.1350
Vapour density (mg/cm^3)	1.35	1.35	1.05	1.00	0.90	0.85	1.00	1.00

Table 4.1: Simulation parameters for the best performing three and four-pulse laser sequences for a range of capsule radii. The first optimised capsule was that with a size multiplier of 1. In order to keep $I\lambda^2$ constant as a function of capsule radius, the pulse powers are multiplied by the square of the size multiplier. A higher level of precision is provided for many quantities than would be expected for 1D simulations, in order to show how the parameters compare to the low instability criteria. While convergence ratio is in places displayed as 16.0, the value to the precision calculated in the analysis is below this upper limit. The original simulation campaign was performed for five capsule sizes; three further sizes (0.25, 0.35 and 0.75) were performed later, but only four pulse sequences were optimized.

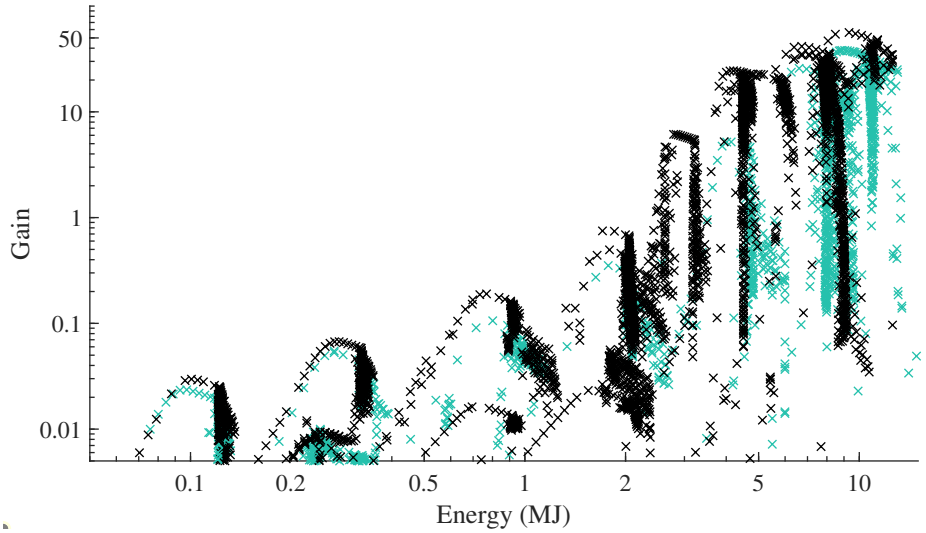


Figure 4.2: A log-log plot showing the simulated gain vs input energy for all simulations satisfying the low instability criteria. Those using a three-pulse sequence are displayed in teal, while those using a four-pulse sequence are displayed in black. The simulations form eight rough peaks, corresponding to the eight different capsule sizes.

These results demonstrate that reactor-level gains can be achieved in this regime, with a gain of 54 being achieved for an energy of 8.5 MJ. However, this is a much higher laser energy than is currently achievable, and is likely too high for an economically viable fusion reactor - suggesting that such a reactor could not be constructed using this particular approach. Nevertheless, the lower energy performance is interesting. The gains of 0.19 and 0.75 at 0.77 MJ and 1.7 MJ respectively are impressive given the conservative nature of this regime. It is also worth noting that when this work was performed in 2020, the record ICF gain was 0.03 - and so this result suggested an improvement of an order of magnitude over the then-current record experimental result. Recent experiments have achieved higher gains, but this work suggests that comparable performance can be achieved in this regime where instability growth is less of a concern (and implosions would hopefully be robust). The recent NIF results (a gain of 0.7 at 1.9 MJ, and a gain of 1.5 at 2.05 MJ) actually fit just below the trend of this data (this is plotted in Figure 5.3 in the following chapter), suggesting that the performance predicted here is comparable to the current record shots on the NIF.

The conditions in the hotspot and shell of the 0.77 MJ capsule are displayed in

4.3. It can be seen that the hotspot conditions are close to that required for ignition. The areal density is greater than the 0.3 g cm^{-2} commonly quoted as required for ignition to occur, but the ion temperature (just below 4 keV) is slightly too low. This suggests that ignition could be within reach, if it was possible to increase the ion temperature just prior to the bang-time. This is investigated further in Chapter 5. It is also interesting to note the relatively low areal density of the shell, which arises from the limited convergence. This is most easily increased in this regime by increasing the size of the capsule.

It is useful to note again here the purpose of this work, which is not to claim that these exact gains are possible - these are the results of 1D simulations only, and thus can not be expected to be fully accurate. However, the expected low instability growth means that such simulations are expected to be a more representative estimate of performance than would otherwise be the case, and thus should be sufficient to give a good first estimate of potential experimental performance. Overall, these results show that it may indeed be possible to achieve interesting fusion performance (comparable to the current best performing shots on the NIF) in this regime.

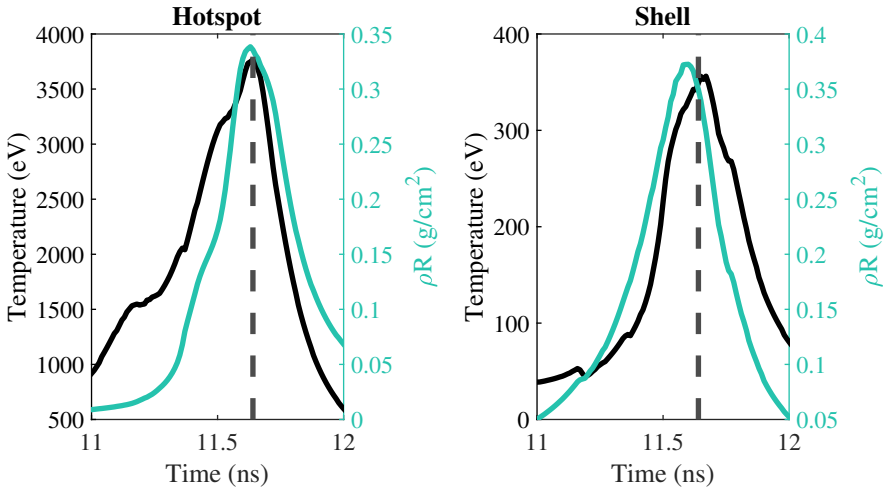


Figure 4.3: Plots displaying the temperature (black) and areal density (teal) in the hotspot and shell for the 0.77 MJ implosion.

4.7 Limitations and suggested future work

There are a number of limitations of this work, some of which will be investigated further in the following chapter. These limitations also suggest some potential avenues for future research.

The fundamental limitation of this work is the use of 1D simulations. While the effort made here to define a low instability regime means that 1D like behaviour is expected, this is still a fundamental and untested assumption upon which this work relies. While 1D simulations are sufficient for an initial investigation of the regime, simulations in higher dimensions are required to estimate the level of instability growth that is likely to be present in a real experiment based on these designs. We have recently secured access to the University of Rochester's RIGEL code, and are collaborating with them to perform confirmatory 2D simulations.

The direct-drive nature of this work coupled with the high laser energies makes experimental verification of the regime challenging. When this work was performed, the use of cryogenic targets and the high energies meant that performing shots based on these designs was not possible, due to a lack of cryogenic capabilities for direct-drive shots at the NIF [31]. Fortunately this has since been rectified, and direct-drive shots of wetted-foam targets are being planned. However, the cryogenic nature of these designs limits the opportunity to test this regime. An alternative solution to this is proposed in the next chapter. In addition, such shots would need to be performed in a PDD configuration. This is contrary to these simulations, where fully symmetric direct-drive was assumed (and the benchmarking simulations in Appendix A suggest that a reduced laser input multiplier would potentially be required to describe PDD shots).

Thirdly, the simulations make a number of approximations that are not particularly accurate, which likely lead to some over-estimation of the performance. The two main ones are as follows: 1) the laser drive is perfectly spherically uniform and focused at the centre of the capsule, which means that it remains perfectly focused as the capsule implodes. This is not the case in real implosions where the laser focus is fixed at the initial capsule radius, and overestimates the laser absorption

(although the 'zoomed focus' technique, discussed in Chapter 6, could be used to reduce losses due to this effect). The benchmarking simulations in Appendix A suggest that the input multiplier used to account for laser losses (including this effect) is reasonable for the case of a spherical direct-drive implosion, although a lower multiplier may be required for a polar direct drive set-up. 2) The description of the wetted foam as simply a DT layer (with the CH foam not accounted for) is a poor one, and is only valid if the foam density is below $\sim$ 10 mg/cm^3 [98]. Other groups have used a less compressible mixed CH + DT equation of state to describe this material with more accuracy [6]; this is also investigated in the following chapter. This is further complicated by the fact that the equation of state for even dry foams at these low densities are not well simulated, which motivated the experiment in the second half of this thesis. Further study is required to better describe this material if wetted-foams are to be simulated to higher accuracy, and ultimately experiments to measure the wetted foam equation of state are required; this is another topic that we are continuing to investigate.

In addition, while these are presented as 'optimisation' campaigns, exploring such large parameter spaces manually is incredibly difficult, and thus there is no guarantee that the most optimal performance has been obtained. This approach would be well suited to a machine learning approach (such as using Bayesian optimisation) to explore the parameter space more efficiently. Machine learning has been applied to ICF for similar problems successfully in the past [96], and this is something that we are currently investigating within our research group. If this could be implemented, it would allow many more simulation campaigns to be performed with relative ease. The impact of adjusting more parameters simultaneously could be explored (for example, the laser pulse powers), and the criteria applied to limit the parameter space also adjusted (i.e. to investigate alternative maximum implosion velocities) or expanded to include other parameters (for instance, applying a minimum adiabat criteria).

Finally, while the regime identified is expected to result in low instability growth, there are other parameters which could be included in future to increase confidence in

this. The most obvious would be the adiabat of the implosion. In addition, no limit was placed on the strength of the first shock, as the impact of particular asymmetry seeds was not considered in this work. This is an important limitation, since the strength of the first shock has been shown to be key to mitigating instabilities from the support tent (in indirect-drive implosions) and ablator micro-structure. Further simulations campaigns should be performed to repeat this campaign, but with laser pulse profile designed to generate a 10 Mbar initial shock (particularly if machine learning can be applied to assist with this).

4.8 Conclusion

A simulation campaign to investigate the performance of wetted-foam capsules in a low instability regime was presented, where the regime is defined by limits to convergence ratio, IFAR, implosion velocity, and laser intensity. While the laser energy necessary for an IFE reactor based on this approach is not feasible with the cost of current-day laser drivers, the performance at lower energies is comparable to the recent record shots on the NIF. This suggests that substantial fusion performance is within reach despite the conservative approach taken here, and that further exploration of this regime is worthwhile. A variety of potential extensions and improvements to this work have been suggested, many of which are explored further in the following chapter.

This chapter has also demonstrated the utility of this style of optimisation campaign - a general platform which can be used to explore fusion performance in a particular parameter space. This approach has again been applied in the subsequent chapter to explore variations on the implosion scheme discussed here.

5

Further simulation campaigns in the low instability regime

Contents

5.1 Introduction 77
5.1.1 An alternative fusion metric: the burning plasma parameter 78
5.2 Estimating the impact of the wetted-foam EOS 79
5.3 Hydrodynamic-equivalent capsules 82
5.4 Alternative laser drivers 84
5.4.1 Background to possible IFE laser drivers 85
5.4.2 ArF lasers 87
5.4.3 Two-colour implosions 88
5.4.4 The burning plasma parameter for alternative laser drivers 92
5.4.5 The zoomed focus technique 93
5.5 Auxiliary heating using relativistic electron beams . . 94
5.5.1 Principles of the auxiliary heating scheme 95
5.5.2 Simulating auxiliary heating in `HYADES` 96
5.5.3 Timing of electron energy deposition 98
5.5.4 Magnitude of deposited electron energy 100
5.5.5 The burning plasma parameter for heated capsules . . . 102
5.5.6 Estimating the gain of heated capsules 102
5.5.7 Comparison with previous simulated and experimental results 105
5.5.8 Auxiliary heating of high gain capsules 107
5.5.9 Optimisation and heating of high areal-density capsules 109
5.5.10 Limitations and suggested future direction 112
5.5.11 Subsequent developments of the auxiliary heating technique 113
5.6 Conclusions 114

5.1 Introduction

The previous chapter presented three key results. Firstly, the new low instability regime was defined, where 1D simulations are expected to provide a reasonable estimate of performance. Secondly, the performance of conventional central hot-spot implosions of wetted-foam capsules using third-harmonic laser light in this regime was explored and well characterised. Finally, the methodology for this style of optimisation campaign (and the necessary code to produce and analyse the simulations) was established.

In this chapter, these previous results are expanded in a variety of ways. In Section 5.2, the impact of the EOS model used to describe the wetted-foam is investigated, in an attempt to address one of the major limitations of the previous campaign. One of the other limitations of this approach was the difficulty of verifying this regime experimentally, and so in Section 5.3 new surrogate targets ('hydrodynamic-equivalent' capsules) are proposed and optimised to enable room temperature experiments.

In the second half of the chapter, two further simulation campaigns are performed. In Section 5.4 the potential performance benefits of using higher frequency lasers are investigated, through additional implosion optimisation campaigns. Finally, in Section 5.5, an auxiliary heating scheme using electron beams is discussed and applied to these implosions. These campaigns further explore the possible performance of wetted-foam capsules within the low instability regime, while also allowing the benefits of these alternative approaches to be assessed (through comparison with the results of the previous simulations).

5.1.1 An alternative fusion metric: the burning plasma parameter

It is useful at this stage to introduce an additional fusion metric, which is used later in this chapter to evaluate fusion performance. So far in this thesis gain has been the primary fusion metric, but calculating gain requires the input energy of the implosion to be known. This is not the case in Section 5.5, due to uncertainty in the exact energy required to drive the auxiliary heating. This dependence on input energy means that the gain in this case is highly affected by 'up-stream' efficiencies (such as how efficiently electron beams can be generated), and thus it is useful to use alternative metrics which do not share this dependence.

A metric known as the 'burning plasma parameter' is therefore introduced for this chapter. The version used here - the 'total capsule burning plasma parameter', as described in [99] - is defined as

$$Q_{\alpha}^{\mathrm{tot}} = \frac{1}{2}\frac{E_{\alpha}}{E_{\mathrm{PdV}}^{\mathrm{tot}}} , \tag{5.1}$$

where E_{α} is the alpha particle energy deposited back into the hotspot over the total duration of the implosion, and $E_{\mathrm{PdV}}^{\mathrm{tot}}$ is the total work done on the capsule hotspot and shell until stagnation (the time of maximum pressure). This metric compares the energy absorbed from alpha self-heating at the bang-time of the capsule, to the external energy that was used to compress it[1].

A value of $Q_{\alpha}^{\mathrm{tot}} > 1$ defines the 'burning plasma' regime, where alpha heating is the primary source of energy in the capsule [63]. The burning plasma parameter is therefore a useful metric for fusion capsules, as it provides direct information about the degree to which fusion reactions are contributing to the energetics of the hotspot. It does not depend on input energy, and thus how efficiently the capsule is driven; this makes it ideal for analysing the auxiliary heating simulations.

Previous works have also attempted to define ignition using such parameters. An early rough definition used an ignition threshold of $Q_{\alpha}^{\mathrm{tot}} = 10$, defining ignition

[1]The burn is assumed to be symmetric around the bang time, such that the alpha self-heating at the bang time is equal to half of the total alpha heating experienced over the full implosion, E_{α}. This is the source of the factor of half in the equation.

simply as when the self-heating was ten times greater than the work done on the hot spot and shell [100]. Recently more rigorous definitions have been created, based instead on an alternative 'hotspot burning plasma parameter', where the total work done on the whole capsule $E_{\mathrm{PdV}}^{\mathrm{tot}}$ is instead replaced by the total work done on the hotspot only. This quantity is typically twice as large as the 'total capsule' version used in this thesis [99]. Analysis of a number of simulations has led to an ignition criteria of $Q_{\alpha}^{\mathrm{hs}} > 1.4$ [63] (where ignition here describes the point at which self-heating results in burn propagation through the cold shell).

5.2 Estimating the impact of the wetted-foam EOS

The `HYADES` simulations used in this thesis are intended to describe wetted-foam targets, but the wetted-foam layer in these simulations is described using the pure-DT equation of state. This approximation does not account for the presence of the foam material, and was highlighted as a limitation of this work in the previous chapter. There is no wetted-foam equation of state that is currently generally available. Other work (at US national labs) has used a mixed equation of state, which does include the presence of the foam [6]; this is reported to be less compressible than pure-DT, resulting in a thicker but less dense shell and a reduced fusion yield [101].

The impact of using a more physical 'mixed' equation of state was investigated using the alternative 1D radiation-hydrodynamic code `HELIOS` [102]. `HELIOS` is compatible with another program, `PropacEOS`, which is used to generate custom EOS tables for materials based on their atomic composition; these can then be used for `HELIOS` simulations involving these materials. A mixed-EOS table was produced for the wetted-foam in `PropacEOS`, assuming that it consisted of a 2% atomic fraction of carbon, 48% tritium, and 50% deuterium (corresponding to a 25 $\mathrm{mg\,cm^{-3}}$ deuterated plastic foam[2], saturated with a 50:50 mix of liquid DT).

[2]The deuterium from the CD foam is responsible for the imbalance between the atomic fractions of deuterium and tritium.

Code	HYADES	HELIOS	HELIOS
EOS	Pure-DT	Pure-DT	Mixed
Gain	0.030	0.040	0.028
Convergence Ratio	15.7	16.6	25.3
IFAR	23.4	24.6	24.8
Implosion velocity (km/s)	391.4	393.8	391.9

Table 5.1: Simulation outputs for the same 0.25 size simulations simulated in HYADES with a pure-DT EOS, and in HELIOS with both a pure-DT EOS and a mixed-EOS.

HELIOS does not model the absorption of alpha-particles, which means that alpha self-heating is not included in any HELIOS simulations. This is a limitation which means that HELIOS cannot be used to accurately simulate the fusion performance of high gain capsules; ignition depends upon alpha particles depositing fusion energy back into the plasma, and neglecting this effect means that initial fusion reactions do not increase the plasma temperature, and thus high gain cannot be achieved. Even the highest gain HYADES simulations presented in this thesis therefore fail to ignite in HELIOS. This means that the impact of this new EOS model on high gain capsules unfortunately cannot be investigated using this code. Due to this limitation, the 0.25 size capsule was instead chosen for the comparison. The gain of this capsule is inherently low and the effect of self-heating is minimal, meaning that HELIOS can be used to describe it with reasonable accuracy. This allows an initial investigation of the impact of this more accurate EOS model on these capsules, but the significance of this change may be different for higher gain capsules.

Two HELIOS simulations using PropacEOS models were performed. The first used a pure-DT EOS to describe the foam layer, while the second used the new mixed-EOS. Both simulations were otherwise unchanged (in terms of laser and capsule properties) from the HYADES simulation. Some differences are expected when changing between codes and EOS models, and so the pure-DT HELIOS simulation was produced to provide a better comparison with the mixed-EOS results. The temperatures and densities achieved at the bang-time in the one HYADES and two HELIOS simulations are displayed in Figure 5.1, while the implosion parameters are displayed in Table 5.1.

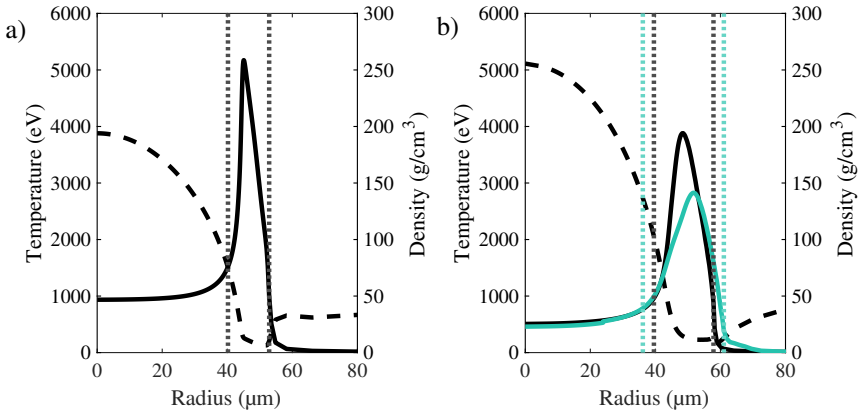


Figure 5.1: Temperature (dashed lines) and density (solid lines) at the bang time for a) the 0.25 size 4-pulse implosion simulated in `HYADES`, and b) the same simulation in `HELIOS`. In b), the teal solid line represents the density profile for the same simulation when a mixed-EOS is used (the temperature profiles for these two simulations display no observable difference). The dotted vertical lines show the hotspot and shell boundaries (where again, the dotted teal lines refer to the mixed-EOS).

Differences are seen between the pure-DT simulations in the two codes; the temperature in `HELIOS` is higher while the densities of both shell and hotspot are lower, and this results in a higher gain and a higher convergence ratio. The hydrodynamic behaviour is otherwise similar. Such differences are expected between different codes and methods for calculating EOS tables. Comparing the two `HELIOS` simulations, it can be seen that the mixed-EOS is less compressible and results in a lower peak density and a thicker shell. The temperature profile is essentially unchanged when changing from pure-DT to the mixed-EOS. These results show good qualitative agreement with the mixed-EOS results presented by Olson *et al.* [101].

The IFAR and implosion velocity using the mixed-EOS are similar to the pure-DT simulation, but the convergence ratio in Table 5.1 is significantly higher. This is a consequence of the particular definition used to identify the hotspot radius[3], rather than representing a physical change in the compression. It is seen in Figure 5.1 that the peak density for the mixed-EOS simulation actually occurs at a greater radius, and the temperature profile (which was used in [101] to define convergence

[3]As the peak density is lower, the $1/e^2$ threshold used to define the hotspot occurs at a lower radius.

ratio) is unchanged between the two simulations. The increase in convergence ratio reported in in Table 5.1 is thus not thought to be significant.

The use of a mixed-EOS table has reduced the gain to roughly 70% of that obtained in the pure-DT simulation. It is difficult to extrapolate these results to larger capsules, as the self-heating instability could mean this results in even lower gains. However, the pure-DT implosion presented by Olson *et al.* [101] had substantially greater fusion performance (with a 100 MJ yield), and a similar 30% reduction was observed (with 71 MJ achieved using the mixed-EOS). This agreement is encouraging, and suggests that as a first estimate a yield reduction of 30% may be expected when switching from a pure-DT to mixed-EOS description of wetted-foam layers.

The comparison performed here is by no means definitive. It would be preferable to estimate the impact that using a mixed-EOS table would have directly in `HYADES`, given that it was this code which was used for the main campaign. The `PropacEOS` tables were converted for use in `HYADES`, but physical results were not obtained when they were used in simulations[4]. This topic should be investigated further in future work. Ultimately, mixed-EOS tables which describe wetted-foams should be produced and made publicly available, and wetted-foam EOS experiments performed to allow the accuracy of such models to be tested and verified.

5.3 Hydrodynamic-equivalent capsules

The implosions presented in the previous chapter are simulated for direct-drive, and feature cryogenic capsules (due to the presence of liquid DT). At the time that this work was performed this presented a major barrier to experimental verification, as there were no facilities that operated at the high (0.8 and 1.7 MJ) energy scales proposed that could conduct direct-drive shots of cryogenic targets[5] [31]. To address this, surrogate 'hydrodynamic-equivalent' or 'hydro-equivalent' versions of

[4]Such simulations showed increased ablation of shell material into the hotspot, reducing the temperature and significantly impacting the yield/gain. This appears to be a problem with the conduction model in `HYADES`, and is currently under investigation.

[5]This has recently been addressed, and it is now possible to perform such shots in a polar direct-drive configuration on the NIF.

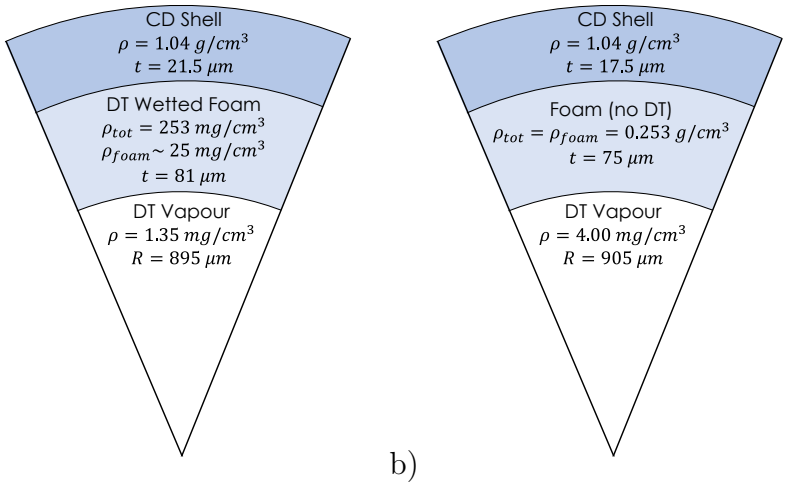


Figure 5.2: a) The original 280 kJ wetted-foam capsule. b) The 280 kJ hydro-equivalent capsule. A higher foam density is used to compensate for the lack of DT, to give an equivalent layer density. The optimisation process is then repeated, resulting in slightly different laser timings and layer thicknesses.

these implosions were developed (following in the footsteps of previous surrogate experiments using symmetry capsules to study hydrodynamic behaviour [103]).

In the hydrodynamic-equivalent capsules, the wetted-foam layer is replaced with a dry foam (without DT-wetting) of equivalent density, removing the need for cryogenic cooling. This change is demonstrated in Figure 5.2. In the original capsule, the wetted-foam layer consisted of a low density foam ($\sim$25 mg/cm^3) saturated with liquid DT, giving an overall layer density of 253 mg/cm^3. In the new capsule this layer is replaced with a higher density 253 mg/cm^3 foam, without any DT, to leave the overall layer density unchanged. Other than the density and presence/absence of DT, the foam used in the two capsules is equivalent.

Hydro-equivalent versions of three different implosions have been developed, and are displayed (and compared to the original designs on which they are based) in Table 5.2. In each case, the new capsule was optimised according to the same optimisation procedure used in the previous chapter, maximising the gain of the implosion while satisfying the four criteria. In each case, this results in an implosion with much reduced yield and gain (due to the significantly reduced quantity of DT). However, the key hydrodynamic variables used to define the low instability

Size multiplier	0.35		0.5		0.65	
Hydro-equivalent?	-	H.E.	-	H.E.	-	H.E.
Laser energy (MJ)	0.27	0.24	0.77	0.78	1.7	1.5
Gain	0.067	0.0094	0.19	0.016	0.75	0.023
Convergence ratio	15.7	16.0	16.0	16.0	15.8	15.8
IFAR	27.5	28.4	29.7	29.4	25.1	29.3
Implosion velocity (km/s)	395.8	396.3	399.6	388.9	399.6	387.5
Max pulse power (TW)	84.79	84.79	173.00	173.00	292.38	292.38
Pulse 2 switch on time (ns)	2.20	2.70	2.60	2.50	3.60	3.70
Pulse 3 switch on time (ns)	4.60	4.90	5.60	5.90	7.80	7.80
Pulse 4 switch on time (ns)	5.50	5.80	6.80	6.90	9.50	9.20
Laser switch off time (ns)	8.50	8.50	11.00	11.20	15.00	14.20
Vapour/liquid boundary (mm)	0.8950	0.9050	1.3050	1.3150	1.6705	1.7225
Liquid/CD boundary (mm)	0.976	0.980	1.3950	1.3950	1.8200	1.8135
Outer radius (mm)	0.9975	0.9975	1.4250	1.4250	1.8525	1.8525
Vapour density (mg/cm^3)	1.35	4.00	1.05	3.90	1.00	3.75

Table 5.2: Hydro-equivalent implosions for three capsule sizes. The wetted-foam implosion for each size has been included for comparison. All implosions use a four-pulse laser sequence. The hydro-equivalent capsules have been optimised, resulting in slightly different values for the optimisation parameters and the laser energy, but the hydrodynamic parameters are broadly similar.

regime - the convergence ratio, IFAR, and implosion velocity - are very similar to the original capsules, and continue to satisfy the low instability criteria.

This provides a series of surrogate capsules, to be used in future room temperature experiments to validate this regime and verify the assumptions of low instability growth. These experiments will allow the instability growth to be measured, and their performance compared to the 1D simulations. If the instability growth is low and the agreement with the simulations good, then this would suggest that the four criteria used to define this regime are indeed sufficient to guarantee good agreement, and thus that the designs presented in the previous chapter are also likely to perform as described in simulations. This therefore provides a route towards experimental verification of these results.

5.4 Alternative laser drivers

A possible way to increase fusion performance is to increase the frequency of the laser driver. Higher frequency lasers have have a higher energy coupling from

inverse bremsstrahlung with the plasma, which results in a greater ablation pressure for equivalent laser intensity. The higher frequency also has a higher critical density and can thus propagate further into the plasma, leading to increased laser absorption (since the light travels through more material) and a higher density blow-off [104]. This results in a higher hydrodynamic efficiency. The lower laser wavelength also means that higher intensities are permitted before the onset of significant parametric instability growth [94] (as highlighted by the $I \cdot \lambda^2$ criteria for the low instability regime).

The simulation platform described in the previous section is well suited to an investigation of different laser drivers. The laser frequency is easily changed in `HYADES`, enabling the same style of optimisation campaign to be performed. The purpose of doing such is two-fold: the possible performance in the low instability regime is further investigated, and the possible impact of increasing the laser driver can be investigated, quantified, and compared to the third-harmonic simulations in the previous section (all in a regime where reasonable agreement is expected between experiment and simulation). Two different laser drivers are considered in this section: a 193 nm laser with a similar pulse profile to the previous simulations, and a novel 'two-colour' laser pulse scheme.

5.4.1 Background to possible IFE laser drivers

Designs for IFE power plants have a number of requirements, including high laser efficiency, broadband laser operation, and high gain. There are two main laser technologies that are seen as potentially satisfying these requirements; diode-pumped solid-state lasers (DPSSL) and electron-beam pumped excimer lasers, such as KrF or ArF [41].

DPSSLs are solid-state lasers that are pumped with high efficiency diodes. Unlike flash lamps, which are used to pump the gain medium at facilities such as the NIF and emit broadband light, diodes emit over a much narrower range of frequencies and thus pump the medium much more efficiently. The wavelength

of a DPSSL depends on the choice of gain medium; the main studies of fusion-relevant DPSSL have considered Yb:S-FAP as the medium[6], which would give a fundamental wavelength of 1.05 µm [105]. This would then be frequency tripled to 0.35 µm for fusion applications - the same frequency used at the NIF, and that the previous simulations were performed for.

Excimer lasers are naturally higher frequency, with fundamental wavelengths of 0.258 µm for KrF lasers and 0.193 µm for ArF lasers. In excimer lasers, the gain medium is a gaseous mixture of a noble gas (Kr or Ar) and a halide (F). An electrical discharge is driven through the medium to excite it, leading to the formation of exciplex molecules - these emit UV radiation when they relax, and it is this relaxation which is used to provide the laser emission. There have been a number of excimer lasers developed (including Sprite/Titania at the Central Laser Facility (CLF) [106], Garpun at Lebedev Physical Institute [107], and Nike at the Naval Research Laboratory (NRL) [108]), but these have so far been limited to energies below 10 kJ.

The 'high average power laser' programme conducted between Los Alamos National Laboratory (LANL) and NRL investigated future laser drivers for IFE power-plants, and led to the development of high repetition rate lasers based on the two technologies [41]; the Electra KrF laser at the NRL (which could produce energies of up to 700 J and repetition rates of up to 5 Hz), and the Mercury DPSSL laser an LANL (which gave shots greater than 50J, with a repetition rate of up to 10 Hz [105]). Both lasers performed well, and this led to the conclusion that either technology could be viable for an IFE power plant and development of both approaches should be pursued.

Both approaches have their advantages. DPSSL lasers (being entirely solid-state) have a higher expected durability, and the lower fundamental frequency is beneficial for preventing damage to laser optics [105]. However, excimer lasers have some strong advantages for direct-drive fusion performance. This includes the higher frequency (and thus greater laser absorption), but also a very smooth beam profile (desirable for direct-drive), good capability for focal zooming, and

[6]Other media, such as the Nd:glass medium currently used at the NIF, are also under consideration.

an inherently high bandwidth (significant for preventing parametric instability growth). These properties mean that simulations suggest that KrF or ArF lasers could be used to achieve higher gains for equivalent energies than DPSSL lasers. There is therefore a continuing interest in these technologies, as demonstrated by recently published simulations suggesting significant gains at low energies using ArF lasers for shock ignition [104].

5.4.2 ArF lasers

Implosion optimisations were performed for a range of capsule radii, using a lower laser wavelength of 193 nm. These results are therefore directly relevant to ArF lasers (for which this is the fundamental frequency), but as the changes depend only on the frequency they also indicate how moving to higher frequencies generally - be that from using a KrF laser (0.258 µm), or using the fourth (0.263 µm) or fifth (0.210 µm) harmonic of Nd:glass - would influence performance. In the low instability regime outlined in the previous chapter, the maximum permitted laser power is set by a limit on $I \cdot \lambda^2$; increasing the frequency thus meant that the laser power could also be increased, while continuing to satisfy this criteria. All the laser powers were thus scaled accordingly. The pulse profile, capsule design, and simulation settings were otherwise unchanged from the previous simulations.

Three different capsule sizes were simulated and optimised by two summer students (Heath Martin and Rusko Ruskov) under my day-to-day supervision. I provided the code required to do the simulation campaign, adjusted this to provide simulations at the new frequency with the new pulse powers, and then trained and supervised the two students as they performed the optimisation. The tabulated results are displayed in Table 5.3, and are plotted in Figure 5.3.

It can be seen in Figure 5.3 that the ArF implosions display a significant improvement in gain for a given energy compared to the previous third-harmonic simulations. A gain of 11 can be achieved for around 1.6 MJ of laser energy, which is an order of magnitude higher than was predicted previously. This clearly displays the potential benefits of using a higher frequency laser driver. It should

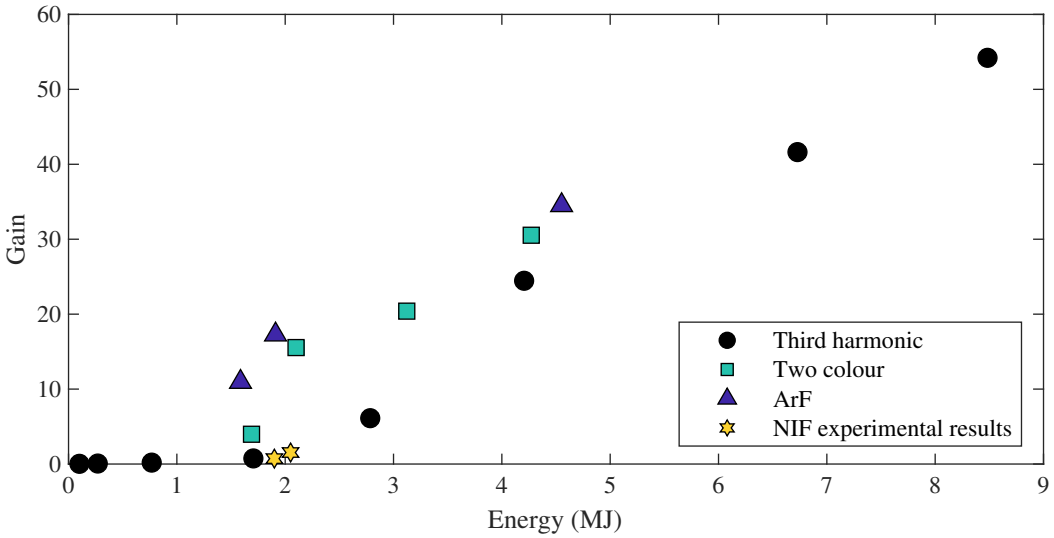


Figure 5.3: A plot of the gain vs the laser energy for simulations using the three different laser drivers. The NIF 210808 experimental shot, and the recent NIF ignition shot, have been included for comparison (although indirect-drive is roughly five-times less efficient in terms of energy deposition into the hotspot than direct-drive [1, 109], and thus the required drive energy for direct-drive equivalents of these implosions would be expected to be significantly lower).

again be noted that this is in the low instability regime, and thus it should be expected that instability growth would be minimal and 1D simulations should give a reasonable estimate of performance.

It is worth commenting on the physical properties of these implosions. The higher laser power but comparable implosion velocity when using ArF lasers means that under this optimisation procedure, for a given capsule radius the ArF laser energy used is much higher than that used for the third-harmonic. Alternatively, this means that for an equivalent energy ArF and third-harmonic laser energy, the ArF implosion will feature a smaller capsule radius. The use of the higher frequency essentially shifts the ignition curve (the relationship between gain and energy) to lower energies, meaning that break-even, ignition, and a particular gain all require lower energies to achieve.

5.4.3 Two-colour implosions

Higher frequency lasers (whether excimer or higher harmonics of Nd:glass) are less technologically developed than third-harmonic Nd:glass lasers, and have not yet been

Laser driver	ArF			Two-colour			
Size multiplier	0.45	0.5	0.65	0.5	0.65	0.75	0.85
Total laser energy (MJ)	1.59	1.91	4.55	1.69	2.10	3.12	4.27
3rd harmonic energy (MJ)	-	-	-	0.92	1.82	2.75	3.96
ArF energy (MJ)	1.59	1.91	4.55	0.77	0.28	0.37	0.32
Gain	10.9	17.3	34.5	4.0	15.5	20.4	30.5
Convergence ratio	16.0	15.9	14.0	16.0	16.0	16.0	16.0
IFAR	11.3	10.5	8.5	16.1	20.4	23.8	23.4
Implosion velocity (km/s)	390.7	398.9	360	391.1	399.9	396.4	391.6
Max 3rd harmonic power (TW)	-	-	-	173	292	389	500
Max ArF laser power (TW)	464	572	967	274	463	617	792
Pulse 2 switch on time (ns)	3.90	3.50	2.50	1.95	3.00	2.70	3.50
Pulse 3 switch on time (ns)	7.10	7.50	8.90	4.35	7.80	7.20	9.10
Pulse 4 switch on time (ns)	8.85	9.25	10.80	8.05	9.40	9.10	11.10
2nd laser switch on time (ns)	-	-	-	10.00	14.70	15.20	18.20
Laser switch off time (ns)	11.95	12.25	15.10	12.80	15.30	15.80	18.60
Vapour/ice boundary (mm)	1.072	1.210	1.435	1.230	1.645	1.933	2.210
Ice/CD boundary (mm)	1.200	1.340	1.780	1.360	1.818	2.100	2.374
Outer radius (mm)	1.2800	1.425	1.8525	1.425	1.8525	2.1375	2.4225
Vapour density (mg/cm^3)	1.00	1.03	0.60	1.02	1.05	1.00	1.01

Table 5.3: Simulation parameters for the two-colour and high frequency implosions.

demonstrated at such high energies. This makes the previous ArF results challenging to achieve in practice. Therefore, an alternative driver scheme is proposed, where the higher frequency laser is used to supplement a third-harmonic driven implosion. This is a compromise solution, which allows some of the performance benefits from using high frequency lasers to be achieved, while reducing the energy demanded from these novel laser technologies. The high frequency energies required are still greater than are currently available, but are far more achievable than those in the previous section. This concept will be referred to as a 'two-colour' approach.

The laser scheme proposed is shown in Figure 5.4. A four-pulse third-harmonic sequence is applied, with the same pulse-profile as used in the previous simulation campaigns. Then, at a late stage in the implosion, a high power short duration pulse is applied at the higher ArF laser frequency. While the four-pulse sequence was already at the maximum allowed power under the $I \cdot \lambda^2$ criteria, the lower wavelength of the final pulse means a higher power can be tolerated. The $I \cdot \lambda^2$ criteria is applied separately for the two lasers, given the large separation in wavelength between them (it has been demonstrated previously that using multiple frequencies

leads to lower instability growth [46]).

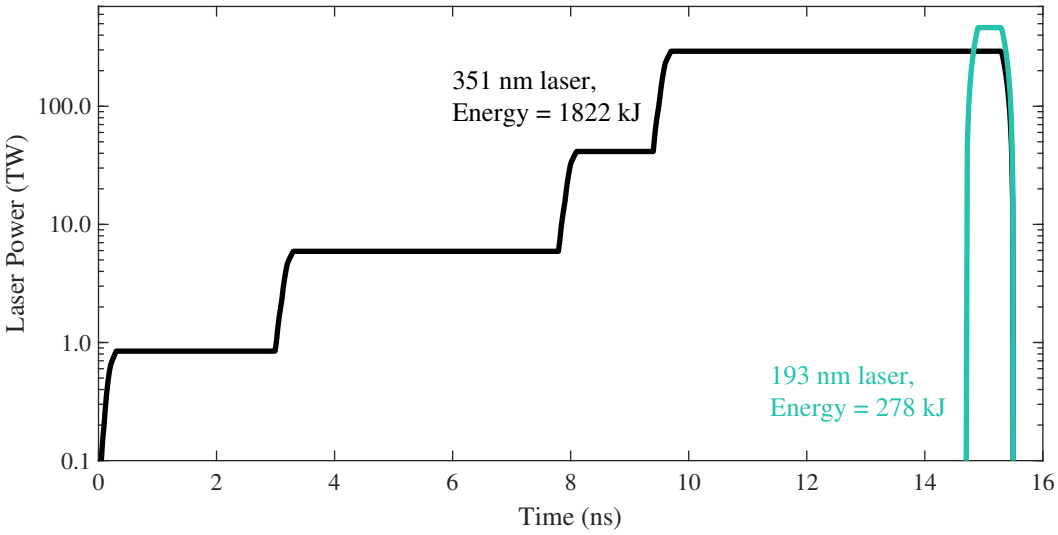


Figure 5.4: A log plot of laser power vs time for the 0.65 size two-colour implosion (which achieves a gain of 15.5). A four-pulse third harmonic laser is applied, which is supplemented by a single late pulse from an ArF laser. The ArF has a higher frequency, and thus a higher peak power can be used than for the third harmonic laser.

As before, the four criteria for the low instability regime continued to be applied, meaning that the peak implosion velocity was still limited to 400 $\mathrm{km\,s^{-1}}$ throughout the implosion. The optimised capsules therefore feature heavier (thicker) shells which are accelerated to lower velocities by the third-harmonic laser, so that they remain under the 400 $\mathrm{km\,s^{-1}}$ limit once the high frequency pulse has been applied. The addition of a late pulse made the optimisation more challenging, as it added an additional optimisation parameter (the switch on time of the final pulse). A lower maximum ArF laser power was used for the late pulse than the maximum permitted under the $I \cdot \lambda^2$ constraint, or than was used for the ArF-frequency implosions; this appeared to be beneficial in preliminary simulations, but the more challenging optimisation procedure meant this parameter was not investigated further. After the code was written to implement these changes and some performed initial simulations were performed, the optimisations were again performed by Heath Martin and Rusko Ruskov under my day-to-day supervision. The results are displayed alongside those of the ArF simulations in Table 5.3 and Figure 5.3.

It can be seen that, as expected, the results for the two-colour implosions consistently fit somewhere between those of the third-harmonic and ArF-frequency results. At 1.7 MJ, the two-colour sequence results in a gain of 4; lower than the gain of 11 achieved for a 1.6 MJ ArF sequence, but a significant improvement compared to the gain of 0.8 achieved using a third-harmonic driver. Perhaps more significant is the gain of 15.5 achieved at 2.1 MJ, requiring only 280 kJ of the high frequency laser. While still a significant increase over current energies for these lasers, this provides a clear demonstration of how this technique leads to significant increases in gain for much more moderate high frequency laser energy; although it should be noted that this scheme would require a facility where the capsule can be irradiated uniformly using two different frequencies at the same time[7].

There are other features of these implosions that it is interesting to comment on. First, note the variation in the proportion of the total energy within the late pulse. This is likely a result of the difficulty in optimising these sequences, due to the additional optimisation parameter. While these simulations are indicative of the performance that may be obtained, it thus appears that further optimisation may be possible. The low IFAR values of these implosions should also be noted, highlighting that this increased performance is achieved by accelerating thicker shells compared to the third-harmonic implosion. Some caution is required here when predicting the Rayleigh-Taylor growth based on these IFAR values, due to the fact that the IFAR quoted corresponds to that measured when the capsule shell is at two thirds of the initial capsule radius. This is before the late pulse is applied, and thus the IFAR value contains no information about the effect of this pulse. Trends between the IFAR value and instability growth are based on experiments and theory which do not include the impact of this late pulse, and thus do not necessarily apply in the same way to the two-colour approach. However, low instability growth is still predicted, due to three reasons: 1) the IFAR values are

[7]This approach would therefore perhaps be more suited to higher harmonics of Nd:glass rather than an ArF laser, as then a single laser system could be used with the higher frequency provided by further frequency conversion on some of the beams - although repurposing existing beamlines for this would degrade the symmetry of the third-harmonic laser drive.

low even compared to the criteria of 30, allowing for some further compression of the shell; 2) the three additional criteria also serve to minimise instability growth; and 3) as the additional pulse is applied late, there is less time for any instabilities that it potentially seeds to grow. When this approach is pursued in future, these predictions should be tested with two-dimensional simulations.

5.4.4 The burning plasma parameter for alternative laser drivers

The burning plasma parameter for all of the implosions shown in Figure 5.3 has been calculated[8], using Equation 5.1. The ArF and two-colour implosions are all well within the burning-plasma regime, with burning plasma parameters of around 10 and above. It can be seen that the 1.7 MJ third-harmonic implosion is the lowest energy third-harmonic result to achieve burning plasma.

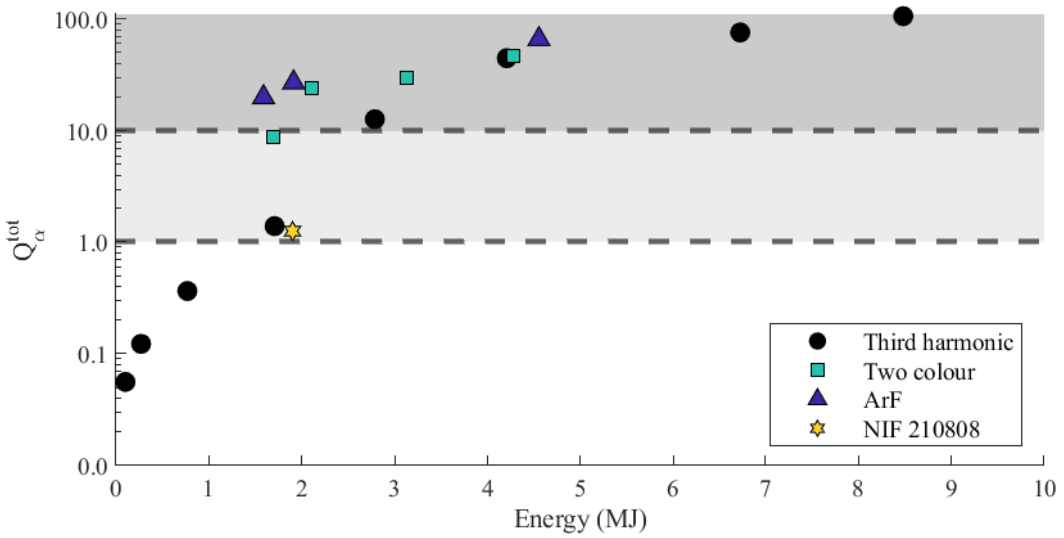


Figure 5.5: A log plot of the total capsule burning plasma parameter Q_α^{tot} vs laser energy for the optimised capsules using the three laser drivers. Values of $Q_\alpha^{\text{tot}} = 1$ and $Q_\alpha^{\text{tot}} = 10$ are indicated by the dashed lines and shaded regions.

[8]The terms were evaluated as follows: 1) the total work done $E_{\text{PdV}}^{\text{tot}}$ was evaluated as the sum of hotspot and shell kinetic and thermal energies, at the point during the deceleration of the shell where these two energies are equal. This corresponds to a period where the total energy of the capsule is constant, just prior to the onset of fusion reactions. 2) the deposited alpha energy over the full capsule was used rather than the hotspot for E_α, assuming that this absorption is dominated by the hotspot [110]. These quantities were evaluated in this way to reduce the dependence of the calculated quantity on the exact position of the hotspot interface, so that it was more robust to issues with hotspot tracking.

NIF shot data for N210808 (the 1.3 MJ yield shot [56]) has also been estimated and plotted on this graph, but this is very approximate and illustrative only[9]. The burning plasma parameter suggests that the performance of this shot was comparable to the simulated performance of the third-harmonic low instability targets (as was likewise observed when comparing the gain of these implosions). Insufficient data was available at the time of writing to plot this for the recent 3.5 MJ yield shot. As with Figure 5.3 the input energy of the NIF shot is high due to the lower efficiency of indirect compared to direct-drive, and a direct-drive version of this shot would therefore likely have a lower drive energy than displayed here.

5.4.5 The zoomed focus technique

Another potential benefit of these approaches are their compatibility with the 'zoomed focus' technique. As the implosion occurs and the capsule is compressed, the capsule radius moves inwards from its initial position. However, the laser focus normally remains fixed at this original location. This smaller capsule radius and reduced surface area means that some fraction of the incoming light will no longer reach the capsule, and instead refracts outwards through the coronal plasma. This results in reduced laser-capsule coupling. In addition, an increase in light travelling through the coronal plasma is expected to contribute to increased cross-beam energy transfer, resulting in further decreased absorption and less uniform compression.

The 'zoomed focus' or 'zooming focus' technique requires changing the focus of the laser during the course of the implosion, allowing it to be better matched to this reduced critical density surface [111, 112]. This is feasible but complex to achieve using conventional Nd:glass lasers [113], but is reasonably straightforward to implement on excimer lasers (due to differences in how the laser is smoothed) and has been successfully demonstrated on the (KrF) Nike laser [111]. This technique leads to reduced laser losses, and thus improvements to the fusion performance. The two-colour scheme is also a natural fit for such a technique;

[9]To do this, the 'hotspot burning plasma parameter' for N210808 was estimated from the figure in [56]. The hotspot compression energy is typically around roughly half that of the total compression energy [99], and so the 'total capsule burning plasma parameter' plotted here is likewise assumed to be half that of the reported hotspot burning plasma parameter.

the high frequency laser can have a smaller focal spot than the third-harmonic laser, providing the 'zoom' late in the implosion.

This effect is not included in the simulations presented here. In the 1D HYADES simulations, the laser provides perfect uniform and spherical illumination, and is focused at the centre of the target[10] - meaning that the aforementioned losses are not included. These simulations represent an ideal case, while the zoomed focus technique is an attempt to solve a non-ideal and practical problem. It is therefore expected that implementing focal zooming (as could be done on ArF or two-colour implosions) reduces a degradation mechanism not seen in the code, and thus give performance closer to that simulated than would be achieved without using this technique.

5.5 Auxiliary heating using relativistic electron beams

In Chapter 4, it was noted that the areal density of the hotspot for the 0.8 MJ capsule exceeded 0.3 $\mathrm{g\,cm^{-2}}$. This is significant, as this is the areal density generally considered to be required for ignition to occur. If the ion temperature could be increased in some way, it may be possible for this capsule to ignite and thus the gain to be significantly increased.

In this section, an auxiliary heating scheme is investigated which could potentially provide this increase in ion temperature. This is based on the scheme proposed and described by Ratan *et al.* [114]. First, a brief description of the scheme is provided. Then, how this was applied to the simulations from Chapter 4 is described, before the results are discussed. This scheme is then applied to some of the higher gain alternative driver simulations, along with new capsules optimised for areal density (rather than yield). Finally, more recent work (to which I contributed, but was performed predominantly by other authors) with implications for how this scheme would be applied in practice are briefly discussed.

[10] It is possible in HYADES to use ray-tracing to simulate the impact of rays of light travelling at non-normal angles. However, it is not possible to include the effect of the finite spot size, rather than an infinitesimally small focus. This is not sufficient to simulate the effect described here.

5.5.1 Principles of the auxiliary heating scheme

The general principles of the auxiliary heating scheme (as first discussed by by Ratan *et al.* [114], and further developed by Lee *et al.* [115]) are as follows. Relativistic electron beams are fired into the hotspot of a central hotspot implosion. These electron beams lead to a bump-on-tail in the electron distribution function, which results in the growth of Langmuir waves. Energy from the electron beam is thus transferred into these waves. The instability eventually saturates due to the trapping of beam electrons. The Langmuir wave modulation instability then leads to the generation of daughter Langmuir waves, which undergo Landau damping. This thus leads to the wave energy transferring to the electrons in the background plasma, raising the electron temperature. Electron-ion collisions mean this also results in an increase in the ion-temperature within the inertial confinement time. This process takes place over the order of a few picoseconds, which is very fast compared to the compression timescales of the capsule.

Ratan *et al.* [114] considered the linear growth rates of the Langmuir wave for two overlapping beams analytically, and found this to be highest where the beams overlapped. Their description of the scheme thus considered two electron beams, overlapped within the hotspot of a central hotspot implosion[11], where it was expected that the energy deposition would occur at the overlap of these two beams. This is the set-up that will be considered in this chapter. The investigation performed by Lee *et al.* [115] (after the work in this chapter was performed) instead demonstrated that in fact the instability quickly saturates in the non-linear regime, meaning that the saturation amplitude (and heating efficiency) is equivalent regardless of whether there is one beam or two. The deposition is thus not preferentially located at the beam overlap, and so cannot be precisely controlled in this way. The consequences of this will be discussed in Section 5.5.11. Ratan *et al.* performed Vlasov simulations of this heating mechanism for fusion relevant conditions, and found that energy

[11]This makes it well suited to low convergence ratio implosions, where the large hotspot makes this easier to achieve.

was transferred from the electron beams to the electrons in the hotspot with an efficiency of 18% - a result which was confirmed by the subsequent work.

How the electron beams are generated is not the focus of this work, but is a topic that has been explored elsewhere due to the application of such beams to fast ignition [21, 116]. Typically, it is expected that an ultra-high intensity picosecond-duration laser pulse would generate the electrons in the plasma through a range of potential interactions. These can include SRS and TPD in the low density region of the plasma, resonant absorption, vacuum heating[12], and $\vec{j} \times \vec{B}$ heating[13][117].

The auxiliary heating scheme appears well suited to the designs presented in the previous chapter. It would be expected to drive a rapid increase in the ion temperature (as is required), leading to an increase in the fusion yield. The low convergence ratio nature of these implosions also provide a large hotspot; this was also an attractive benefit, since it was understood at the time that the scheme required the overlap of two beams within this region.

5.5.2 Simulating auxiliary heating in HYADES

Simulating auxiliary heating of these implosions is challenging, due to the multi-scale nature of such a set-up. Fluid codes such as HYADES can simulate the implosion, but cannot simulate the auxiliary heating scheme as they do not include the necessary physics (it is not possible to generate Langmuir waves or to simulate Landau damping in such a code). Alternatively, particle-in-cell and Vlasov codes do include these effects, but cannot be feasibly ran over the large time and spatial scales required to simulate the implosion. As a result, approximations must be made.

In this work, HYADES is used to estimate the impact of the auxiliary heating, by estimating the effect that such a scheme would have on the electron energy.

[12]Also referred to as 'Brunel heating' or 'Not-so-resonant resonant absorption', this is a case of resonant absorption where electrons within the skin-depth of the critical density plasma are dragged out by the (rapidly decaying) electric field associated with the EM wave, and then accelerated back into the high density region at high velocity a half-cycle later.

[13]This effect is important for high intensities, where the magnetic component of the EM wave cannot be neglected. If the wave travels normal to the critical density surface and the electric field causes electrons to oscillate in the direction of the E field, then the $\vec{v} \times \vec{B}$ component of the Lorentz force will also be normal to the surface, and electrons can be accelerated across the critical density surface.

The Vlasov simulations performed by Ratan *et al.* suggest that, for relevant conditions, the heating leads to an increase in the bulk plasma electron energy in the centre of the capsule. As a consequence, HYADES simulations were performed where electron energy was added to the central zones (for convenience, this was done over those zones which initially contained the DT vapour at the start of the implosion), over a seven picosecond window (corresponding to the timings from [114]) at a point just before the bang time of the implosion. The electron beams, Langmuir waves, and Landau damping are not simulated - but the estimated impact of such a scheme is. The effect this heating has on the fusion performance of the implosion was thereby investigated.

Figure 5.6 shows the kinetic, thermal, and total energy in the hotspot for two HYADES simulations, in a 0.4 ns window around the bang-time. Figure 5.6 (a) considers an implosion without auxiliary heating. As the shell decelerates and the capsule stagnates, the kinetic energy decreases and the thermal energy increases accordingly. The total energy is initially constant, but a small number of fusion reactions occur which leads to a small increase in energy (through alpha self-heating). Figure 5.6 (b) displays the same implosion, but where 20 kJ of electron energy is added in order to simulate the effect of auxiliary heating. This leads to a rapid increase in the thermal energy. This then leads to a significant increase in the number of fusion reactions, as shown by the much greater subsequent increase in total energy. It can also be seen that the increase in energy from the alpha self-heating is larger than that used for the auxiliary heating.

This demonstrates that adding electron energy in this way can indeed be used to improve the fusion performance of the capsule. In the following sections, this approach is applied to and investigated using four of the capsules developed in Chapter 4. The key data for these capsules has been summarised in Table 5.4 for convenience, where they are labelled as capsules A-D. In the simulations discussed in this section, the implosions are unchanged from those reported in Table 4.1 apart from the addition of the electron energy. The capsule and laser parameters are not varied. The laser pulse-sequence used in the simulation - the four-pulse

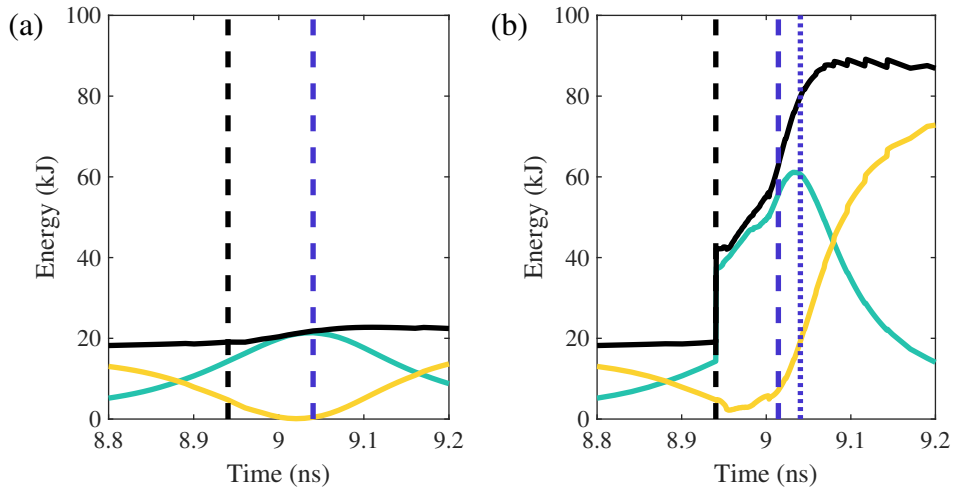


Figure 5.6: Plots showing the total (black), thermal (teal) and kinetic (yellow) energy in the hotspot of the same implosion both with (a) and without (b) 20 kJ of energy added to the electrons. The dashed black line in both plots displays the time at which the energy should be added to achieve the maximum increase in yield. The bang-time for each implosion is displayed by a dashed blue line. For the heated implosion, the bang-time from (a) is indicated by a dotted line; it is clear that applying heating means the bang-time occurs earlier in the implosion.

Capsule label	A	B	C	D
Size multiplier	0.25	0.35	0.5	0.65
'Long-pulse' laser energy (kJ)	101	270	768	1710
Gain	0.030	0.067	0.19	0.75

Table 5.4: Summary of key data for the four implosions (labelled A-D) that the auxiliary heating is applied to. Full details of these implosions can be found in Table 4.1.

third-harmonic profile optimised in the previous chapter - will be referred to as the 'long-pulse' laser energy, to differentiate it from the hypothetical 'short-pulse' laser that is needed to generate the electron beam. The phrase 'heating energy' will be used to describe the added electron energy; this is varied independently of the long-pulse laser, which always has the energy described in Table 4.1.

5.5.3 Timing of electron energy deposition

The timing at which the heating is applied to the implosion has a large impact on the increase in fusion performance. To investigate this, simulations were performed where 4 kJ of electron energy was added at different times around the bang-time

of the unheated implosions. The results are displayed in Figure 5.7. The metric used here to evaluate performance is the relative yield amplitude, which is defined in this paper as the yield of the capsule when the heating is applied, compared to the yield of the capsule when no electron energy is added[14].

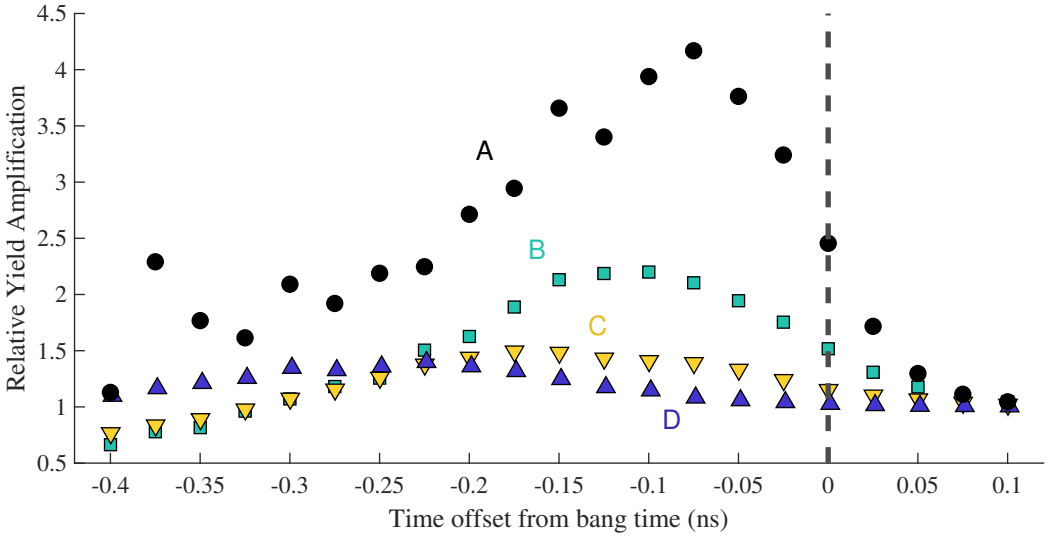


Figure 5.7: Relative yield amplification for the capsules A-D, as a function of the time at which 4 kJ of energy is added to the electrons. The capsule labels correspond to those in Table 5.4. The increased variation in capsule A compared to the others is likely due to the low gain of that capsule, as small changes in performance thus have a larger relative effect.

There is a clear peak in the yield amplification as a function of time for each capsule, indicating the optimal time for the heating to be applied. This optimal time is also that displayed in Figure 5.6, where it occurs just before the minimum in kinetic energy (as the shell stops moving inwards), and just before a significant number of fusion reactions begin. This is unsurprising: adding the energy to the hotspot prior to this point makes it harder to compress and thus reduces the density, but if it is added too late then there is less time over which the fusion rate is increased. The rate of fusion reactions is a strong function of the temperature, and so a later increase in temperature means that the total number of reactions (the fusion rate integrated over the finite burn time) will be lower than for an equivalent

[14]Note that this is different to its conventional usage comparing the yield when alpha self-heating is/is not included.

temperature increase at an earlier time[15]. The same behaviour and optimal timings were found when the simulations were repeated using 20 kJ of injected energy, suggesting these optimal values are independent of the magnitude of energy added.

5.5.4 Magnitude of deposited electron energy

Further simulations then investigated how the fusion performance varied with the magnitude of electron energy added. Between three and sixty kilojoules of electron energy were deposited for each of the four implosions, with the energy injected at the optimal time for each implosion identified from Figure 5.7. The results are displayed in Figure 5.8. It is seen that (as expected) adding more energy has a larger impact on the yield amplification, and even modest quantities of injected energy have a significant effect. The relative yield amplification decreases as the capsule size increases; this is as expected, since 1) the larger capsules are higher energy, and thus the deposited electron energy is a smaller proportion of the capsule energy, 2) the larger capsules have a larger mass, and thus the injected energy per unit mass is lower, and 3) the yield of the unheated capsule is higher, and so a larger increase in yield is required for the same relative amplification.

The relative yield increase is largest for the smallest capsule (A), where 10 kJ of deposited energy gave over ten times yield amplification, and 60 kJ amplified the yield by a factor greater than eighty. Even for the largest capsule (D), 10 kJ of deposited energy is sufficient to double the yield, while 60 kJ amplifies the yield by over twelve times. This clearly demonstrates the potential this scheme has for increasing capsule performance.

The same data is displayed in Figure 5.9, but with the heating energy scaled to the DT mass of the capsule. This shows that relative to the fuel mass, the larger capsules actually show a greater relative improvement in the yield per unit of electron energy added. A similar trend is observed if the deposited electron energy is instead scaled to the 'long-pulse' compressional laser energy of the capsule. This

[15]This is further exacerbated by the fact that the temperature increase from the auxiliary heating will, through the increased fusion rate, lead to increased self-heating and thus result in even higher temperatures and fusion rates later on.

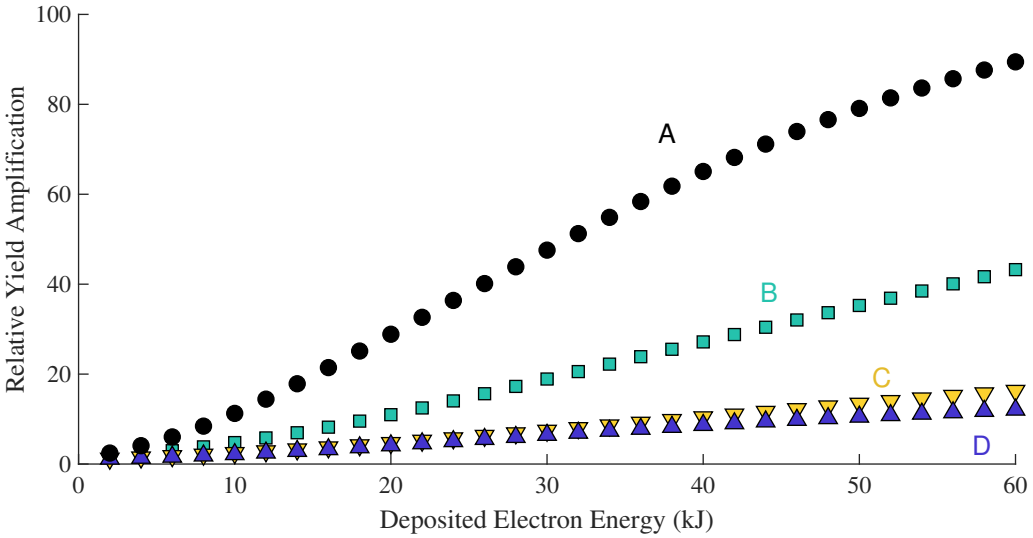


Figure 5.8: Relative yield amplification for capsules A-D as a function of the magnitude of deposited electron energy. The energy is injected at the optimal time for each capsule identified from Figure 5.7.

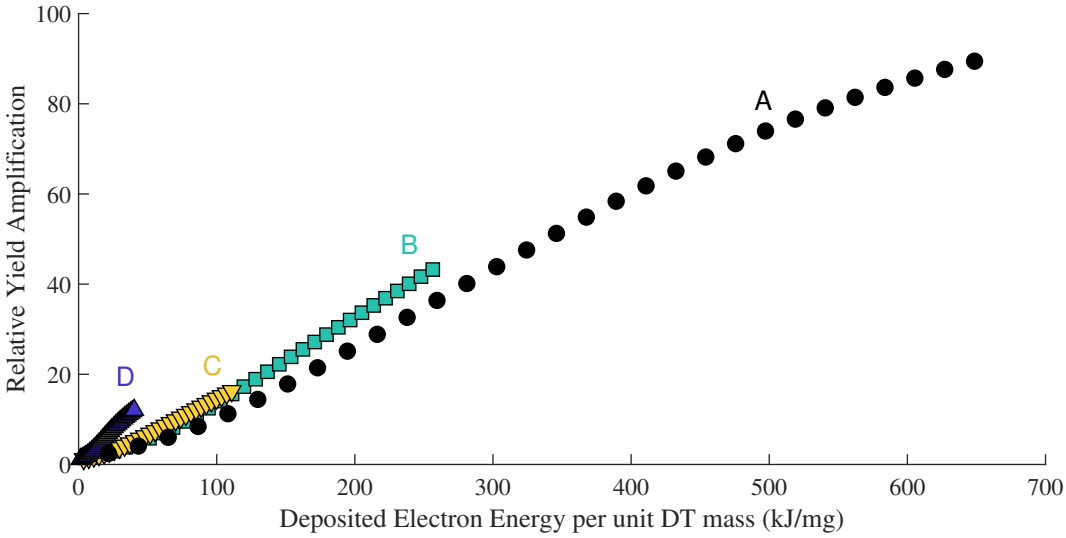


Figure 5.9: Relative yield amplification for capsules A-D, as in Figure 5.8, but as a function of the deposited electron energy scaled to the DT mass of the capsule. A similar trend is observed if the electron energy is scaled to the 'Long-pulse' laser energy.

can be explained by the fact that the larger capsules are much closer to the ignition threshold, around which there is a dramatic increase in yield with input energy. As such, a similar fractional increase in the input energy (in this case, from the increase in deposited electron energy) will have a more substantial impact on the yield than for smaller capsules[16]. For larger capsules that already undergo ignition, the relative

[16]This can also be considered through the DT reactivity. This parameter has an extremely

yield amplification achieved relative to the capsule mass would therefore likely be reduced. However, while it is clear than the larger capsules respond better to auxiliary heating relative to their size, the quantity of energy required to achieve this heating is obviously much larger, and may be challenging to achieve experimentally.

5.5.5 The burning plasma parameter for heated capsules

The (total capsule) burning plasma parameter Q_α^{tot} has been evaluated for the heated capsules A-D, and is displayed in Figure 5.10 as a function of the deposited energy [17]. These results again show the improvement in performance for all four implosions when auxiliary heating is applied. While capsule D is already within the burning plasma regime, auxiliary heating also enables burning plasmas to be achieved using capsules B and C. When the largest quantities of energy are added capsule D sees an increase in the burning plasma parameter of almost a full order of magnitude, approaching $Q_\alpha^{\text{tot}} = 10$. This indicates that self-heating adds almost ten times more energy to the hotspot than is added through external compression, suggesting a very well performing implosion.

5.5.6 Estimating the gain of heated capsules

Despite the challenges in calculating gain for these capsules, it is still the ultimate metric of interest for inertial fusion energy applications. The gain has therefore been estimated using values for the relevant efficiencies from the literature. These efficiencies are not known to high accuracy, and thus the gains presented in this section are indicative only. There are two processes involved for which efficiencies must be estimated: 1) the deposition of energy into the plasma via the electron beams; and 2) the generation of the electron beams via short-pulse laser. The first

strong temperature dependence at low temperatures (i.e. below 8 keV), so causing an equivalent increase in the hotspot temperature for the larger (and hotter) capsule will cause a larger increase in the fusion rate than for the smaller ones.

[17] The terms were evaluated as follows: 1) the total work done $E_{\text{PdV}}^{\text{tot}}$ was evaluated as the sum of hotspot and shell kinetic and thermal energies just after the electron energy was added, such that the deposited energy is included in the calculation 2) the deposited alpha energy over the full capsule was used rather than the hotspot for E_α, assuming that this absorption is dominated by the hotspot [110].

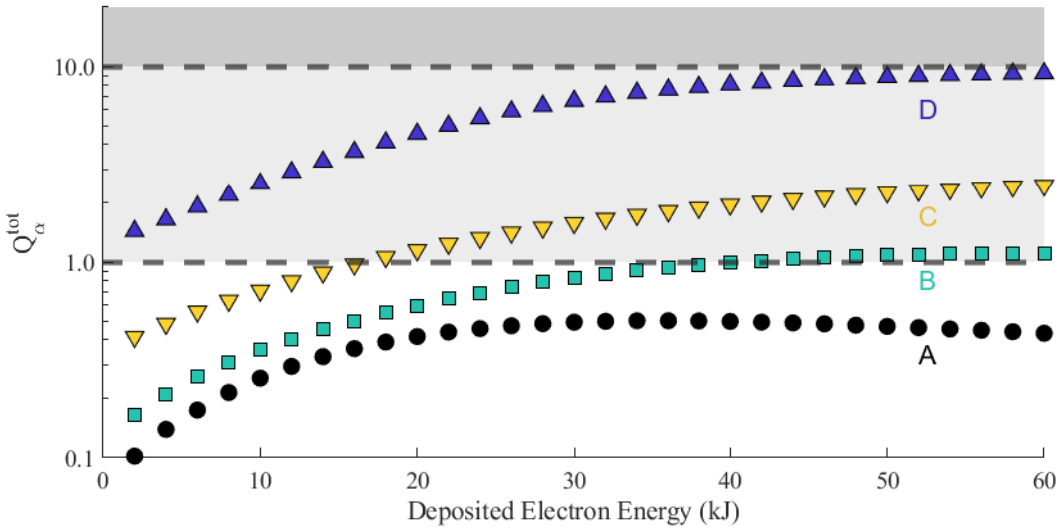


Figure 5.10: A log plot of the total capsule burning plasma parameter Q_α^{tot} vs deposited electron energy for the four capsules. Values of $Q_\alpha^{\text{tot}} = 1$ and $Q_\alpha^{\text{tot}} = 10$ are indicated.

of these conversion efficiencies can be taken from the Vlasov-Maxwell simulations reported by Ratan *et al.* [114], which found that the electron beams deposit electron energy into the hotspot with an efficiency of 18%.

Previous research on fast ignition has investigated electron-beam generation using short-pulse lasers, and have produced a range of estimates for the efficiencies with which such beams can be generated and transported (some examples are [118, 116, 119], and the review article [25]). For the estimates used here, two results are highlighted. The first is that of [120], who performed particle-in-cell simulations for an ignition scale plasma; they estimated an overall laser to electron power conversion efficiency of 52%, which they then coupled to a hydrodynamic code for fast ignition simulations. The second is that of [121], who also performed particle-in-cell simulations but found that as the electron beam passes through the weakly collisional background plasma a significant proportion of the energy is lost. They therefore calculated a lower efficiency of 15% - although they also observed that this increased with both laser power, and the time for which the laser was applied. The maximum simulation time of 2.5 ps used in their work was lower than the 7 ps that the electron energy is deposited over in the `HYADES` simulations, and so if this trend between deposition time and efficiency were to

continue it is possible a higher efficiency could be expected.

The value of 52% from [120] and 15% from [121] have both been used, to provide two separate estimates of the gain. Each of these efficiencies have been combined with the 18% simulated coupling efficiency for electron beam energy deposition into the plasma to calculate an overall efficiency for the auxiliary heating process (capturing generation of the electron beams from short pulse laser, transport of these beams, and deposition into the plasma). The short-pulse laser energies required to cause the deposited electron energies in Figure 5.8 were thus calculated; this was added to the 'long-pulse' laser energy required for each implosion to calculate the total input energy, which was then used to calculate the gain. These two sets of estimates are displayed in Figure 5.11.

The general trends are the same in both figures; auxiliary heating causes an increase in gain, although this begins to saturate for the largest heating energies. For the high efficiency case (a), higher gains can be achieved by using a smaller capsule with some auxiliary heating than through conventional hotspot ignition alone (as seen by the fact that lower energy capsules with a small quantity of heating outperform larger capsules with no heating). This is unsurprising, as the 9.4% heating efficiency is higher than the roughly 6% efficiency for direct-drive laser compression [1, 109]. For the lower efficiency case (b), direct-drive compression is higher efficiency, and so this is no longer observed. The saturation for the highest values of auxiliary heating is also not surprising; this is likely because both high density and high temperatures are required for high fusion yields, and at some point the density starts to become the limiting factor.

The results are obviously far more impressive for the high efficiency case, which predicts break-even for capsule C at around 1.1 MJ of total energy (using $\sim$ 350 kJ of short-pulse laser energy to deposit 32 kJ of electron energy into the hotspot), and gains of up to $\sim$ 3.5 for under 2 MJ of energy using capsule D. Given the observed trends, it appears likely that a capsule with size between B and C would achieve break-even for under 1 MJ of total energy. However, even in the low efficiency case, it is clear that auxiliary heating can still be a useful tool to improve

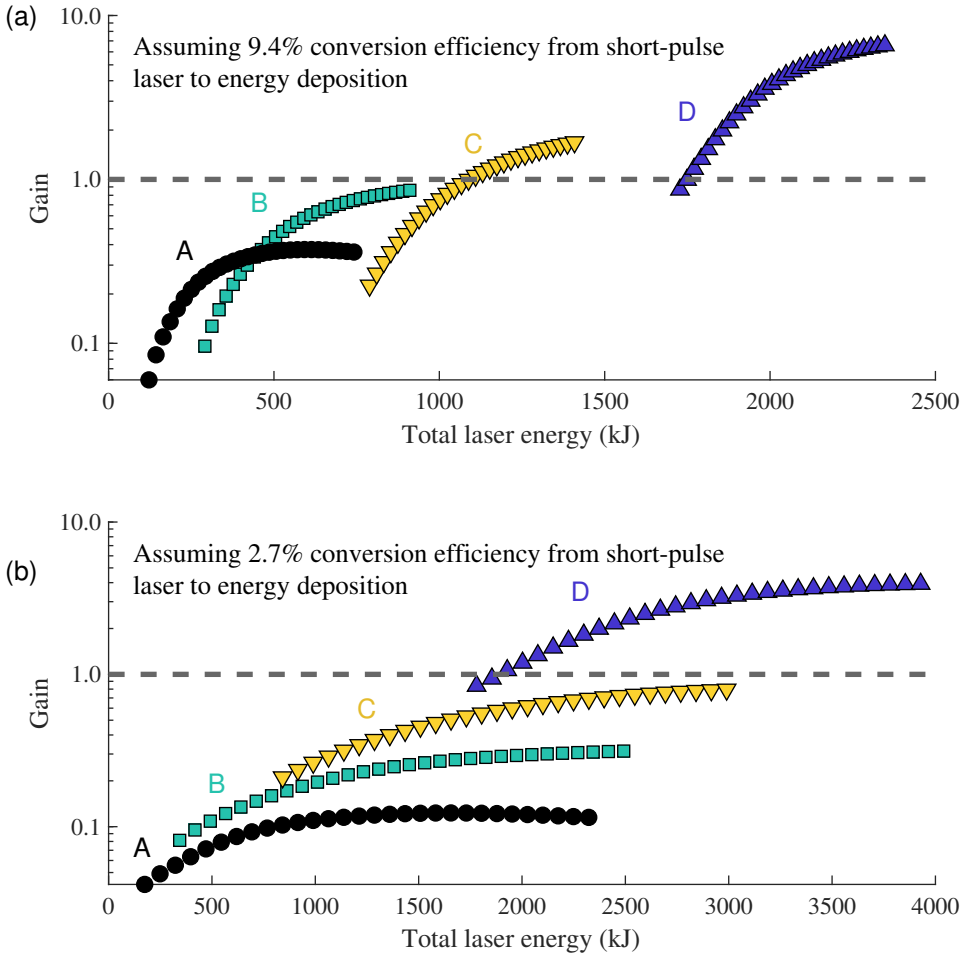


Figure 5.11: Estimated gain vs total input energy, including both the 'long-pulse' laser energy used to drive the implosion as reported in Table 5.4, and the estimated 'short-pulse' laser energy required to generate the simulated auxiliary heating. (a) and (b) use different estimates for the efficiency with which the heating can be generated, as indicated on the plots.

the performance of a capsule, and capsule D can be seen to achieve break-even for under 2 MJ using this technique.

5.5.7 Comparison with previous simulated and experimental results

The results presented for capsules A-D are now compared to a range of previous results. Simulations were previously performed to estimate the impact such a scheme may have on indirectly-driven targets, which were presented in [122]. This

was based on the NIF shot N160421, with a low convergence ratio and a drive energy of $\sim$ 800 kJ. The greater efficiency of direct-drive implosions results in an approximately five-fold increase in the energy deposited in the hotspot [1, 109], which means that such a capsule should be roughly equivalent to an implosion in this work with an energy of around 150 kJ. There are some minor differences between these works; the indirect-drive results demonstrate less dependence on the time at which the heating is applied, and a lower level of amplification than would be expected based on the results presented here (some differences should be expected due to the use of a different code, `xRAGE`, for the indirect drive simulations, and for the differences in the base implosions - with the indirect-drive implosion having significantly lower yield). Despite this the results are broadly similar between the two works, with similar trends observed in both the direct and indirect-drive simulations.

The fusion performance of capsule B is also comparable to that of NIF shots N190918, N191007, and N191110 (the capsules with record inertial fusion yields when this work was first performed [123]), although these shots targeted higher convergence ratios over 25. These NIF shots produced neutron yields ranging from 0.75×10^{16} to 2×10^{16} (compared to 0.8×10^{16} for capsule B without auxiliary heating) and fuel kinetic energies of $\sim$15 kJ (also similar to capsule B), for $\sim$1.9 MJ of indirect drive input energy. These shots led to predictions that laser energies of $\sim$3 MJ [123] may have been required to achieve MJ-level yields (other predictions thought energies of up to $\sim$5 MJ may be required [123, 124]). However, extrapolating the auxiliary heating results presented for capsule B to these implosions suggested that MJ-level yields could be achieved for a total of 2.3 MJ[18]. This highlights how such a heating scheme could potentially be used to improve the performance of real NIF shots. These results could also be extrapolated for the 1.3 MJ yield shot N210808. This has a yield between those of capsules C and D; applying a similar extrapolation suggests that this yield could be doubled for only 10 kJ of deposited energy. These

[18]This was estimated by taking the 56 kJ yield of N191110, and assuming based on implosion B in Figure 5.8 that a yield amplification of twenty times could be achieved for 32 kJ of auxiliary heating, which would require less than 350 kJ of short-pulse laser energy assuming a conversion efficiency of 9.6%.

Capsule label	E	F	G
Laser driver	Two-colour		ArF
Size multiplier	0.5	0.65	0.5
'Long-pulse' laser energy (MJ)	1.69	2.10	1.91
Gain	4.0	15.5	17.3

Table 5.5: Key data for three high gain implosions (labelled E-G) that the auxiliary heating is applied to in this section. Full details of these simulations can be found in Table 5.3.

are illustrative examples only, but serve to demonstrate how this technique could be applicable more widely than just for the low instability regime discussed here.

5.5.8 Auxiliary heating of high gain capsules

Auxiliary heating was also later applied to three of the alternative laser driver implosions. The main purpose of this was to see the effect of auxiliary heating on high gain capsules; auxiliary heating helps the capsule to achieve ignition without adding more fuel, and thus it is likely that for better performing capsules the increase in performance would be more moderate. It also allowed the combination of these technologies to be assessed. These simulations were performed by summer students Tat Li and Eugene Kim, under my day-to-day supervision.

Three capsules were used: the 0.5 and 0.65 size two-colour implosions, and the 0.5 size ArF implosion[19]. These capsules are labelled E-G in Table 5.5. The auxiliary heating was simulated exactly as described in the previous sections, with the optimal timing for the heating identified as outlined previously, before the magnitude of deposited energy was varied. The results are shown in Figure 5.12.

As expected, it was found that for higher gain capsules, the auxiliary heating is much less beneficial. The 0.5 size two-colour implosion has the lowest base gain, and it can be seen that low values of added energy result in moderate increases to gain, but as further heating is applied gain actually begins to decrease. This corresponds

[19]These simulations were performed a year after the capsules were first optimised, and used the updated `HYADES` version 01.12.05 compared to version 01.12.03 for the original simulations. Rerunning the unheated input decks using this new version resulted in a slightly increased yield. This explains why the gain at low heating levels in Figure 5.12 is still higher than that reported in Table 5.5.

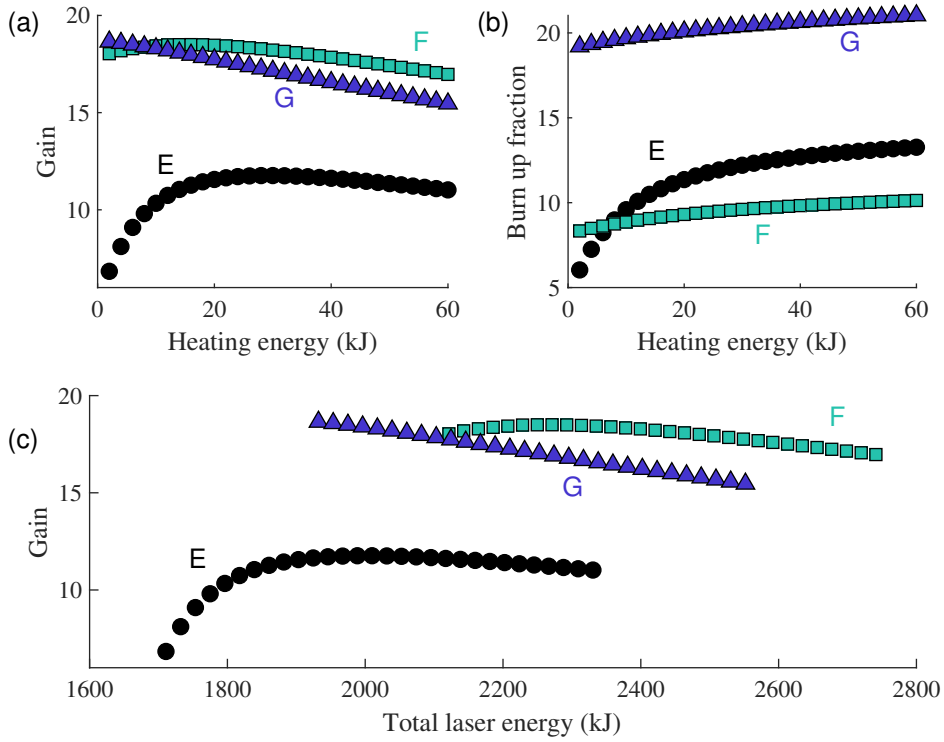


Figure 5.12: Auxiliary heating results for two-colour implosions (E and F) and an ArF implosion (G). In (a), the estimated gain (assuming a 9.4% heating efficiency) is plotted as a function of the heating energy. In (b), the burn-up fraction of the three implosions versus heating energy is displayed. (c) shows the estimated gain, as a function of the total estimated input energy.

to a scenario where the additional fusion reactions generated by the increase in energy no longer compensate for the additional cost. This is not unexpected - this capsule has already achieved the conditions necessary for a large number of fusion reactions to occur, and thus the electron energy deposited (as a fraction of the hotspot energy) is much less significant than for the previous capsules.

As the yield of the unheated implosion further increases, this effect becomes even more dramatic. The 0.65 size two-colour capsule sees an even smaller fractional increase in performance, and less electron energy can be added before the gain begins to decrease. For the ArF capsule, adding electron energy offers no benefit even for low levels of heating, and the gain decreases from the start. This is despite the unheated gain being very similar to that for the 0.65 size two-colour capsule. The reason for this is that the ArF capsule is smaller (the energy is the

same due to the higher laser power), and has a much higher burn up fraction (as shown in Figure 5.12 (b)). The unheated ArF capsule is therefore further along its ignition curve - it is already igniting very well, and is close to the maximum possible yield for this capsule. The two-colour capsule is larger, with more fuel and thus a higher maximum possible gain, and thus the performance can be further increased compared to the ArF capsule.

5.5.9 Optimisation and heating of high areal-density capsules

The implosions A-D (and the high gain capsules E-G) are implosions that were originally optimised for high gain, in the absence of auxiliary heating. If the auxiliary heating was included within the optimisation campaign, it may be that alternative capsule or laser pulse profiles would be found that outperform these designs.

This would add significant complexity to the optimisation, which makes it infeasible to explore without further development. However, it is possible to predict the form such an implosion may take. Auxiliary heating allows the temperature of the capsule to be significantly increased, but the areal density is unchanged. As such, it was suggested that if an implosion could be developed with a higher areal density (at the expense of a lower temperature and yield), then once auxiliary heating is applied to increase the ion temperature the resulting yield/gain may be higher than if a capsule optimised for yield was heated.

To explore this, a further two optimisations were performed. This followed the same procedure as before, but the implosions were optimised to produce a capsule with the maximum possible areal density (i.e. the areal density of DT within both hotspot and shell), rather than yield. I again produced the changes to the code necessary to enable this, and then supervised a summer student, Tat Li, in performing the optimisations. The resulting 'optimal areal density' implosions are included in Table 5.6, along with the corresponding 'high gain' version. Auxiliary heating was then applied, first identifying the optimal timing for the capsules, and

Laser driver	Third harmonic		ArF	
Size multiplier	0.5		0.5	
Optimised for	Gain	ρR	Gain	ρR
Laser energy (MJ)	0.77	0.80	1.91	2.16
Gain (original)	0.19	-	17.3	-
Gain (at time of work)	0.19	0.12	18.7	24.7
Areal Density ($\mathrm{g\,cm^{-2}}$)	0.71	0.76	1.04	1.30
Convergence ratio	16.0	15.9	15.9	15.9
IFAR	29.7	23.2	10.5	9.0
Implosion velocity (km/s)	399.6	358.1	398.9	361.9
Pulse 2 switch on time (ns)	2.60	3.60	3.50	4.20
Pulse 3 switch on time (ns)	5.60	6.80	7.50	9.20
Pulse 4 switch on time (ns)	6.80	8.20	9.25	11.00
Laser switch off time (ns)	11.00	12.60	12.25	14.40
Vapour/liquid boundary (mm)	1.305	1.290	1.210	1.150
Liquid/CD boundary (mm)	1.3950	1.390	1.340	1.340
Outer radius (mm)	1.425	1.425	1.425	1.425
Vapour density ($\mathrm{mg/cm^3}$)	1.050	1.200	1.030	1.405

Table 5.6: Newly optimized 'optimal areal density' implosion simulation parameters, for the 0.5 size third-harmonic and ArF implosions. The data for the original 'optimal gain' implosions have been included for comparison.

then varying the deposited energy. The resulting gains (assuming the optimistic 9.4% conversion efficiency for the heating) are then displayed in Figure 5.13.

The gain for the smaller of the two capsules, seen in Figure 5.13(a), behaved as predicted. It can be seen that for low heating energy the 'optimal gain' capsule outperforms the 'optimal areal density' version. However, as the heating energy is increased further, the 'optimal areal density' capsule begins to perform better, producing a higher gain than the 'optimal gain capsule'. The 'optimal gain' capsule begins to saturate earlier - this makes sense, since it has a lower areal density and a higher temperature, and so the temperature stops being the limiting factor much earlier than for the 'optimal areal density' version.

A number of factors mean that the data in Figure 5.13(b) is more complicated to analyse. Firstly, the optimisation procedure in this case gave an 'optimal areal density' capsule at a higher laser energy than the 'optimal gain' capsule (recall that this is the optimal capsule for a given radius - and not necessarily a given drive energy). The higher laser drive meant that the capsule gain was

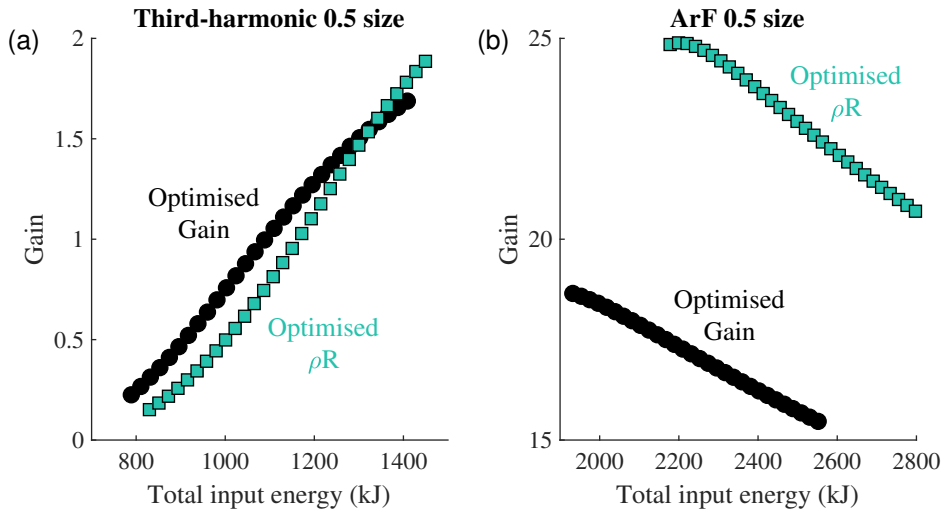


Figure 5.13: Estimated gain vs total input energy for the 0.5 sized capsules optimised for gain (black) and areal density (teal). a) shows the results for the third-harmonic 0.5 size implosions, while b) displays those for the ArF 0.5 size implosions. For each of the four data sets, the left-most (lowest energy) point represents the implosion with no auxiliary heating, and moving to higher energies represents increased auxiliary heating.

also higher, hindering useful comparison between the two capsules[20]. As seen in Figure 5.13 (b), the gain is therefore higher for all values of deposited energy. Secondly, the much higher gain of this capsule means that the auxiliary heating in both cases does not lead to significant increases in gain, as discussed in Section 5.5.8. However, qualitative differences are still visible - although heating adds little benefit in either case, while the 'optimal gain' capsule sees gain decrease for all heating energies, there is a small increase in gain for the 'optimal areal density' capsule for low values of heating. This is confirmed by the relative yield amplification (not plotted), which shows a greater relative increase for the 'optimal areal density' capsule at low energies.

It is clear that while both 'optimised areal density' capsules respond better to auxiliary heating than the original 'optimal gain' implosions, neither of these new capsules shows a particularly significant improvement in performance over the original version. These results also show that the optimised areal density capsules are not too dissimilar from their optimal gain equivalents. A possible explanation

[20]This also indicates, as explained earlier, that the optimisation is not exhaustive - as there is clearly a design for this capsule radius with a higher gain than that initially optimised.

for this is the laser pulse profile used. The relative pulse powers are unchanged (as they were for the previous optimisations), and were chosen to approximate the isentropic compression that is typical for conventional hotspot ignition. It is likely the case that a different laser pulse profile would be better suited to producing a high areal density[21], and that if the optimisation was repeated using such a profile more favourable results would be obtained.

5.5.10 Limitations and suggested future direction

As stated previously, the simulations presented here are an initial investigation of the effects of such a heating scheme. They do not simulate the heating scheme itself, and thus high accuracy should not be expected. However, they indicate a promising level of performance, and justify further investigation. This includes simulations of the full heating process using Vlasov-Maxwell and particle-in-cell codes, along with multi-scale modelling which can integrate the results from such codes into a fluid simulation. These are directions that are under further investigation within our group.

Another limitation of this work is the lack of accuracy to which the efficiencies for the generation and transport of the electron beams are known. These efficiencies are vital for accurate estimates of the implosion gain, and so should be investigated and quantified with greater precision in future short-pulse high intensity experiments.

A final potential limitation of this approach are the high short-pulse laser energies required. These are significantly higher than the short-pulse energies available at current facilities, and are similar to those necessary for fast-ignition schemes [120]. Again, this is an initial investigation of this approach, which is focussed on whether such a scheme may be worthwhile rather than the practical aspects of achieving it. However, recent developments such as plasma beam combiners [126, 127] may offer a route through which short-pulse laser energies could be increased in the future to the levels required [128]. Plasma beam combiners utilise non-linear scattering interactions within plasma-based optical components to produce a single

[21] Producing high areal densities at lower temperature is of interest for fast ignition, where the aim is to produce an isochoric fuel assembly without a central hotspot. Such implosions use very different pulse profiles, such as that discussed in [125].

laser output from multiple input beams, achieving a much higher intensity than is possible with conventional optics. It is expected that continued development of this approach, and of short-pulse laser technology more generally, will enable the energies required for such a scheme in the future.

5.5.11 Subsequent developments of the auxiliary heating technique

Since this work was performed, there have been further developments in understanding this technique which affect how it would be applied; these are described in [115]. While I assisted with this research, the primary contributions to it came from other authors (most notably Jordan Lee, Heath Martin and Rusko Ruskov); it is included in this thesis simply in the interests of completeness, and in ensuring that this section provides the most up-to-date description of the heating process.

In the previous discussions, the heating scheme was assumed to require two overlapping electron beams. This was the set-up proposed in [114]; the highest linear growth rates for the instabilities was seen at the overlap of these beams, and it was expected that this would therefore be where the most significant heating occurred. Having two beams in this way allowed for control over where the energy is deposited, and by overlapping them in the hotspot deposition of electron energy into the centre of the implosion could be achieved. However, these new results instead discovered that there is no benefit to using two orthogonal beams, as the instabilities achieve the same saturation amplitude regardless of whether one or two beams are used. The heating occurs along the beam according to the plasma conditions at each point, and the position of energy deposition cannot be controlled through via beam overlap.

Particle-in-cell simulations performed by Jordan Lee of an electron beam incident on a temperature/density profile representative of a capsule at the bang-time found that the vast majority of the energy is deposited at the edge of the hotspot. The electron beam passes through the cold, dense shell with very little absorption, but is absorbed effectively as soon as it reaches the high temperature material. This is different to the simulations here, where the energy deposition was assumed to be

uniform over the central regions in the simulation. Updated versions of these HYADES simulations (also performed by Jordan Lee) accounted for this, instead depositing the electron energy in a thin shell (of appropriate thickness) around the hotspot. It was found that these actually slightly improved the performance[22]. This assumes that the heating is spherical whereas in fact this would only occur in the direction of the beam, and 2D simulations are thus required to investigate this in further detail.

5.6 Conclusions

This chapter expanded upon the results of the previous simulations in a number of different ways. Firstly, two potential limitations of the previous work were explored. The impact of the EOS used to describe the wetted-foam layer was investigated through simulations in the alternative 1D code HELIOS, where a PropacEOS 'mixed-EOS' table could be used which accounted for the presence of the CH foam. This investigation was limited (due to the need to use an alternative code, and the lack of alpha-heating in HELIOS), but suggested that using a more accurate EOS for the foam layer could potentially decrease performance by around 30%, which is comparable to other results published in the literature. Secondly, surrogate 'hydrodynamic-equivalent' implosions were optimised for a number of capsules, and displayed equivalent hydrodynamic performance to the main implosions. This provides a way that the results of the previous chapter (and the low instability nature of the regime) could be explored in room temperature experiments.

The impact of alternative laser drivers on implosions in this regime were then explored. It was demonstrated that switching from a third-harmonic Nd:glass laser to a higher frequency laser (ArF in this work) could lead to substantial increases in fusion performance for a given energy. A 'two-colour' laser sequence was also proposed, which would allow some of the performance benefits of higher frequency lasers to be obtained, while allowing the bulk of the energy to be provided by the more mature Nd:glass technology.

[22]It is suggested that this is because the material at the edge of the hotspot is likely slightly higher density but colder - and thus the heating has more of an effect.

Finally, the effect of auxiliary heating on a number of these implosions was simulated. It was found that such a technique significantly improves the fusion performance of a range of capsules. The increase in gain was also estimated to be significant, but this is highly dependent on how efficiently electron beams are generated and their energy deposited into the hotspot. Further simulations identified that the impact of such a scheme reduces as the yield of the unheated implosion increases, suggesting that auxiliary heating is most useful for helping sub-ignition or poorer performing implosions to ignite rather than improving the gain of those that are already performing well.

6

Measuring the principal Hugoniot of TMPTA

Contents

6.1 Introduction 117
6.2 Experiment overview and diagnostics 118
6.2.1 The Vulcan laser 119
6.2.2 VISAR 121
6.2.3 Streaked Optical Pyrometry 130
6.2.4 X-ray pinhole camera 133
6.2.5 Fiducial 134
6.2.6 Calorimetry 135
6.2.7 Pulse length 135
6.3 Preparation for the experiment 135
6.3.1 Development of the optical set-up 135
6.3.2 Target design and hydrodynamic simulations 140
6.3.3 Hydrodynamic simulations 144
6.4 Conducting the experiment 149
6.4.1 Changes from planned target 150
6.4.2 Poor VISAR illumination and subsequent changes to optical set-up 151
6.4.3 Later than expected shock breakout times 152
6.4.4 Signal strength issues 154
6.5 Conclusion 155

6.1 Introduction

Despite the recent excitement surrounding wetted-foam ICF capsules and interest in foams more generally, the equation of state of foam materials is poorly characterised. The EOS of wetted-foam capsules was discussed in Section 5.2, where the lack of wetted-foam tables was discussed and the need for further investigation of this material highlighted. This is not a unique problem to wetted-foams however; foam materials generally are described in simulations simply as a low density version of the homogeneous material on which they are based. This does not account for the foam structure or inhomogeneity, and previous experiments have shown such an approach is inaccurate in some cases [129]. This suggests that even mixed-EOS tables for wetted-foams may be a poor description, and also suggests a potential source of inaccuracy for the 'hydrodynamic-equivalent capsules' where dry foam is used.

This chapter presents an experiment designed to help resolve this problem through studying the shock compression of trimethylolpropane triacrylate ('TMPTA', $C_{15}H_{12}O_6$) plastic foam at 260 mg cm^{-3}, to obtain principal Hugoniot data for this foam against which existing and future EOS tables may be compared. This is a comparable material and density to that used in the hydrodynamic-equivalent capsules, making these results of direct relevance to those designs. The experiment was performed using the Vulcan kJ-class Nd:glass laser at the UKRI-STFC Rutherford Appleton Laboratory. A basic overview of the experiment is provided in Section 6.2, along with a discussion of the laser and diagnostics used. The preparation and design work performed before the experiment is discussed and presented in Section 6.3. Finally, the obstacles faced during the experiment and the changes these necessitated are discussed in Section 6.4. The analysis, results and interpretation are summarised in Chapter 7.

It is also worth noting the general current research interest surrounding foams, which have appeal beyond just for their role in wetted-foam capsules. Foams have a variety of other potential fusion applications, including laser beam smoothing and imprint mitigation [130, 131] (where the foam is used to spread energy from an incoming laser beam laterally, reducing the imprint issues associated with

direct-drive), adiabat shaping [41, 132, 133] (where the temporal laser profile is shaped so that the outside of the shell is placed on a higher adiabat - and thus less susceptible to instability growth - while the inner fuel is kept on a lower adiabat and can thus be better compressed), increasing absorption of laser light [28] (as the laser propagates further into the material, due to regions where it is under-dense, compared to a uniform material), and increasing conversion from laser light into x-rays [134]. There is particular interest into how homogenisation (how the inhomogeneous foam transitions over time to a homogeneous plasma) and micro-structure affect foam compression [135, 136, 137, 9], and this experiment is therefore also of potential interest for reasons beyond just its relevance to the fusion capsules considered in this thesis.

6.2 Experiment overview and diagnostics

A multi-layer target, containing a layer of TMPTA plastic foam at an initial density of 260 $\mathrm{mg\,cm^{-3}}$, was shocked via irradiation from the Vulcan laser. The target was a sandwich design consisting of four layers: a $\sim$ 40 µm CH ablator, followed by a $\sim$ 3 µm gold layer, a $\sim$ 40 µm α-quartz reference layer, and then finally the $\sim$ 40 µm TMPTA foam. A schematic of this target can be seen in Figure 6.1. Vulcan was used to irradiate the CH layer, ablating material and driving a strong shock through the target. The gold layer was present to attempt to prevent pre-heating of the quartz/foam, by absorbing x-rays generated in the coronal plasma[1].

The six Vulcan long-pulse beams were overlapped spatially to produce a 'top-hat' beam with a $\sim$ 400 µm diameter, generating a shock that was roughly uniform over this area. The pulse power profile was also a 'top-hat', with a rise and fall time of $\sim$ 200 ps. The target had a step in the rear side, so that the TMPTA foam only covered half of the quartz surface in the $\sim$ 400 µm shocked region. This meant that both quartz and foam were visible to the diagnostics, and is similar to targets used in previous experiments to investigate other materials [138, 139]. The shock

[1]Gold is expected to absorb x-rays with energies below $\sim$ 5 keV effectively, but has significant transmission for higher energies.

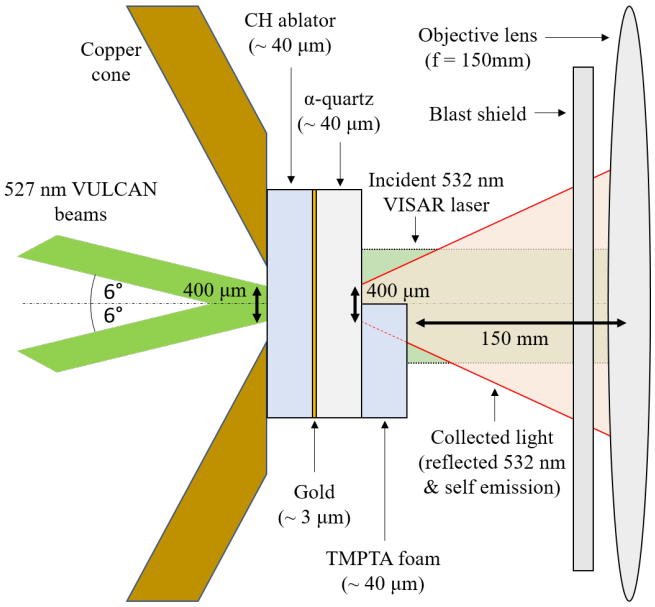


Figure 6.1: Simple schematic showing the target and objective lens (top view, not to scale).

velocity in both layers was measured using VISAR, which was driven using a probe laser reflected off the rear surface of the target. The shock temperature of both layers was determined using streaked optical pyrometry (SOP).

The quartz layer in this experiment was used as a reference material for an impedance matching calculation. The shock velocity in both the quartz and the foam were measured in each shot and used to determine the foam shock state; the theory behind this process is described in Chapter 2 (Sections 2.8.4 and 2.8.5). A single shot is thus sufficient to identify a single data point on the foam Hugoniot. By performing a series of shots at different laser powers, the Hugoniot curve could therefore be populated over a range of pressures.

6.2.1 The Vulcan laser

The Vulcan laser is based at the UKRI-STFC Rutherford Appleton Laboratory's Central Laser Facility (CLF). It is a 2.6 kJ Nd:glass laser with a fundamental frequency of 1053 nm, with 8 beam lines (6 long-pulse, and 2 short-pulse) and two experimental areas providing different operational modes. In Target Area West, the 6 long-pulse beams are available alongside 2-short pulse beams. The long-pulse

beams can each deliver energies ranging from 300 J to 50 J, dependent on pulse lengths ranging from 0.5 ns to 10 ns. These beams can be converted to second harmonic using KDP crystals (with some conversion losses) [140]. The short-pulse beams can either both deliver 1 ps pulses, or offer a mixture of 1 ps and 10 ps pulses, with an intensity up to mid 10^{19} W cm^{-2}. In Target Area Petawatt, a single short-pulse petawatt beam can provide intensities of up to 10^{21} W cm^{-2} at shorter pulse times of 500 fs (or lower powers and intensities for pulse times up to tens of picoseconds), alongside a long-pulse beam line with the aforementioned properties.

Each beam line begins with a 'seed' pulse, which is amplified to produce the main beam [140]. For the long-pulse beams this seed is relatively high energy, and thus the pre-amplification stage consists of a single rod amplifier [141]. The pulse then enters the main amplification stage, where it is amplified by a series of further glass amplifiers. After each round of amplification, the beam diameter is increased; this reduces the intensity, which keeps it below the damage threshold of the laser glass. Rod amplifiers are used up to a diameter of 45 mm, above which more efficient disk amplifiers are used[2] [141]. To produce short-pulse beams, chirped pulse amplification (CPA) [142] is used. Using this technique, a seed pulse is stretched temporally; this is then preamplified, and then amplified using the same conventional glass amplifiers[3] up to the glass damage threshold, before being re-compressed in the compressor chamber to achieve a higher intensity short-pulse beam.

This experiment was conducted in Target Area West, and used only the six long-pulse beams. These were frequency doubled to second harmonic, and focused using custom phase plates into a single 'top-hat' style laser spot with a diameter of 400 µm. The same pulse was requested from each beam, coincident in time, to produce a single flat-topped laser pulse. This pulse time was varied throughout

[2]By angling the amplifier relative to the beam axis disk amplifiers can be pumped uniformly over their entire cross-section, which avoids pump-depletion over the central region. The amplifiers are placed at Brewster's angle to avoid reflections.

[3]The petawatt beam line is pre-amplified using optical parametric CPA [143, 144, 145] rather than glass amplifiers due to the gain narrowing effect, as the short eventual pulse duration provides a limit on the minimum bandwidth ($\sim$ 4 nm for a 500 fs pulse) for the amplified pulse [146].

the experiment between 2 ns and 9 ns; this resulted in a range of pulse energies and intensities, thus driving shocks of a range of different strengths.

6.2.2 VISAR

VISAR (an acronym for 'velocity interferometer system for any reflector') is a velocity interferometry technique which allows the velocity of a moving surface to be determined [147, 148]. A probe laser is reflected from the surface, and the reflected light is passed into the interferometer (or 'the VISAR'). The interferometer delays the light in one leg relative to the other and thus interferes light reflected from the surface at two different times to produce a fringe pattern, from which the velocity of the surface is calculated. In this experiment, VISAR was used in this way to determine the shock velocity of the quartz (while also being used to identify breakout times, allowing the average foam shock velocity to be calculated).

Numerous mathematical descriptions of the VISAR exist. The one presented here is a simple explanation which considers the interference pattern from an interferometry perspective, as presented originally in [149] (and conceptually similar to the explanation given in [147]). This derivation neglects the Doppler shift caused by the moving mirror, and instead simply considers the change in the mirror position in order to derive the correct result. Alternative derivations instead focus on the Doppler effect[4], but then typically ignore the phase difference arising from the movement of the surface. Derivations based on this approach can be found in numerous sources [149, 148, 147, 150]. More complete derivations (such as those by Clifton [151], Goosman [152] and Bennett [153]) reconcile these two approaches, but are much more complicated.

A simple schematic of a VISAR is shown in Figure 6.2. Consider photons A and B, which travel down the VISAR legs A and B respectively. Both arrive

[4]The most simple form of this derivation is that presented in [149]. Using the same variables defined in the main equation, the moving mirror introduces a Doppler shift in the light of $\Delta\lambda(t) = -\frac{2\lambda}{c}u(t)$. The length of the delay leg can be written in terms of multiples of the laser wavelength, $N\lambda = c\tau$. Differentiating this gives $\Delta N(t) = -\frac{c}{\lambda^2}\Delta\lambda(t)$. Substituting in the expression for $\Delta\lambda(t)$ then gives $u(t) = \frac{\lambda}{2\tau}\Delta N(t)$. The interference pattern is based on the difference in path length between the two legs, and thus the fringe shift is proportional to ΔN - leading to the same equation for the VPF as the main derivation.

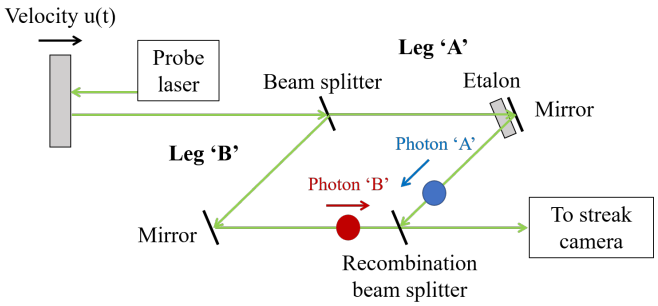


Figure 6.2: Simple schematic of a VISAR. Photon 'A' and photon 'B' travel down different legs to arrive at the recombination beam splitter at the same time and interfere with one another. These photons enter the VISAR (i.e. reached the first beam splitter) at different times; as the delay leg (leg 'A') takes an extra time τ to transit compared to leg 'B', photon A must therefore have entered the VISAR a time τ earlier than photon B.

at the recombination beam-splitter at the same time. Photon A travelled down the delay leg, and thus took a time τ longer than photon B to travel through the VISAR - for them to arrive at the same time, it must therefore have entered the interferometer a time τ earlier.

Before entering the VISAR, both photons travel from the probe laser to the reflective surface, and then from the surface to the VISAR. For the photon 'A', travelling at the earlier time, this distance is denoted as D. If the surface is moving, then photon B travelling at a different time then encounters the surface in a slightly different position. It therefore travels a distance $D + 2d$, where d accounts for this small change (it is multiplied by two, accounting for the fact that the photon will travel this distance both to and from the surface). This distance d is given by

$$d = \int_{t_1}^{t_2} u(t') \cdot \mathrm{d}t' = \int_{t_1}^{t_1+\tau} u(t') \cdot \mathrm{d}t' \ , \tag{6.1}$$

where $u(t')$ is the surface velocity as a function of time t', t_1 is the earlier time at which photon A encounters the surface, and t_2 is the time at which photon B encounters it. The approximation $t_2 = t_1 + \tau$ has been made, which neglects the time taken for the photon to cover the distance d; this is reasonable, since the time $2d/c$ for the photon to travel this distance is typically much less than the time delay τ introduced by the etalon. The time $t_1 = t - T$ is linked explicitly to the time at which the measurement is made at the streak camera t by considering

the time T it takes the photon A to reach the streak camera; in this case, the distance $d = d(t)$ is itself a function of time.

The two photons therefore have a total path difference (when they interfere at the recombination beam-splitter) of $\Delta z = 2d + \epsilon$, where ϵ is the (constant) additional path length introduced by the etalon. This leads to a phase difference of

$$\Delta\phi = k(2d + \epsilon) \, , \tag{6.2}$$

where $k = 2\pi/\lambda$ is the wave vector. Following the standard theory for interference between two waves [154], the resulting intensity is given by

$$I = I_1 + I_2 + 2\sqrt{I_1 I_2}\cos\Delta\phi \, , \tag{6.3}$$

where I_1 and I_2 are the intensity of the two waves in the different interferometer legs. As these are of equal intensity, this simplifies to

$$I = I_0(1 + \cos\Delta\phi) = I_0(1 + \cos k(2d + \epsilon)) \, , \tag{6.4}$$

where $I_0 = 2I_1 = 2I_2$.

In this form, the intensity profile does not depend on spatial position (assuming that velocity is constant across the surface), and thus the interference 'pattern' is constant spatially. In order to produce spatial fringes, the recombination beam-splitter is rotated slightly to introduce a small linear phase ramp. This gives an intensity of

$$I \propto I_0[1 + \cos(k(2d + \epsilon + x\sin(\theta)))] \, , \tag{6.5}$$

where x is the spatial position and θ is the (constant) angle of the beam-splitter, adding spatial dependence and resulting in a repeating pattern of fringes.

At this point, it is useful to consider two illustrative cases. First, consider the situation when the surface is stationary, or moving at a constant velocity u_1. In this case the distance $d(t)$ is given by $d(t) = \tau u_1$, and is a constant. Equation 6.5 then exhibits no time dependence, and the fringe pattern is constant as a function of time, described by

$$I(t_1) \propto I_0[1 + \cos(\phi_1)] = I_0[1 + \cos(k(2\tau u_1 + \epsilon + x\sin(\theta)))] \, . \tag{6.6}$$

Now, let the velocity of the surface change to a new constant velocity u_2. This results in a new fringe pattern described by

$$I(t_2) \propto I_0[1+\cos(\phi_2)] = I_0[1+\cos(k(2\tau u_2+\epsilon+x\sin(\theta)))] \\ = I_0[1+cos(\phi_1+2k\tau\Delta u)] \ , \quad (6.7)$$

representing a small fringe shift from the original pattern based on the difference in the two velocities $\Delta u = u_2 - u_1$. The pattern shifts by one complete fringe when $2k\tau\Delta u = 2m\pi$, where m is any integer. The 'velocity per fringe', or VPF, is the velocity change that accompanies a full fringe shift, and is therefore defined as

$$\text{VPF} = \frac{\lambda}{2\tau} \ . \quad (6.8)$$

A small correction is applied to this equation,

$$\text{VPF} = \frac{\lambda}{2\tau(1+\delta)} \ , \quad (6.9)$$

to account for dispersion within the etalon δ [155]. Further specific corrections can also be required depending on the precise experimental target configuration [156, 157].

This shift in fringes thus allows the velocity of the surface to calculated. The VISAR begins recording when the surface is stationary, and as the surface moves the total fringe shift is recorded. The VPF is used to calculate the surface velocity based on this fringe shift, allowing the velocity profile of the surface as a function of time to be determined. Analysis of VISAR data follows this principle, and a demonstration of how the velocity vs time profile is obtained from the fringe pattern is provided in Appendix B.3.

Etalon selection

The etalon used in the VISAR determines the delay time τ, and thus determines the VPF of the system. Using a thicker etalon results in a larger time delay and a lower velocity per fringe, giving increased velocity resolution. The value of τ also

determines the time resolution of the system; if a velocity change occurs in a time less than τ, it is recorded as a sudden discontinuity in the phase pattern.

Accurate determination of the velocity requires the fringe shift F(t) to be known. A discontinuity in the phase pattern prevents F(t) from being known with certainty, as the fringe pattern repeats if velocity increases by the VPF. A fractional phase shift therefore corresponds to a number of possible velocities, each separated by integer multiples of the VPF. However, the ambiguity in velocity this causes is resolved by introducing a second VISAR, using a different etalon. As long as the VPFs of the two etalons are not integer multiples of each other or a shared root, the different fractional phase shifts seen on the two VISARs are unique to a particular velocity, and thus the velocity is determined without ambiguity.

VISAR design and operation

A labelled diagram of the VISAR used in this experiment is displayed in Figure 6.3. This VISAR was developed and designed by Professor Daniel Eakins. While the main components are the Mach-Zehnder interferometer, the overall device contains a series of optics to aid with alignment.

The VISAR is internally aligned so that the first flip reticle is positioned in the centre of the beam path. This can then easily be used to confirm that the external probe beam entering the VISAR is co-aligned with the system. However, it also serves another important purpose. The etalon and mirror in the delay leg are placed on a piezo-electric motor. As these components are aligned at an angle to the incoming beam, moving them will shift the position of the image.

For VISAR operation it is required that these two images are aligned. To do this, the etalon is first removed from the delay leg. White light is injected into the VISAR, and the motorised stage on which the mirror and etalon mount sit in the delay leg moved until fringes appear. The white light is incoherent, and so fringes are only observed when the two legs have equal path length - this is referred to as the ‘null’ or ‘white light’ position. The system is aligned so that the images are

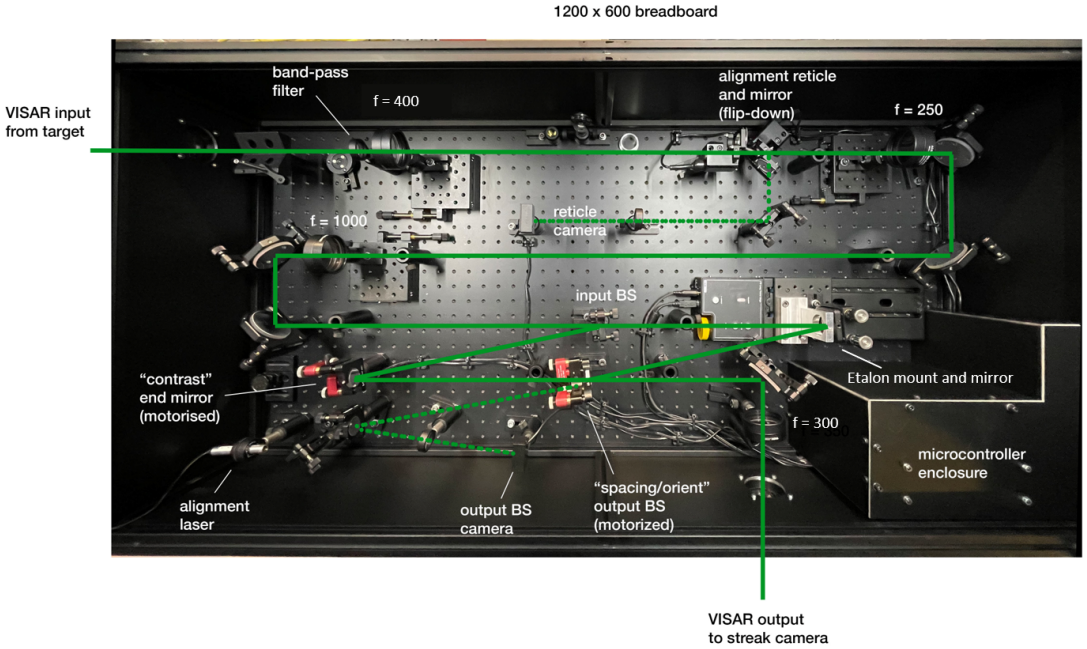


Figure 6.3: The VISAR used in this experiment (developed by Professor Daniel Eakins). The actual interferometer consists of the components between the input/output beam-splitters ('BS'), while the other components are used for alignment and to ensure good image contrast. Image provided by Professor Daniel Eakins, with minor changes to labels to reflect the lens set-up used in the experiment.

coincident for this set-up, which is confirmed by flipping the reticule into the beam and viewing the images at the output beam-splitter camera.

The etalon is then inserted. This introduces some refraction of the beam, and thus shifts the image. The etalon and mirror are thus translated on the motorised stage until the images are collinear again. This offset is calculated[5] using the formula

$$d = t\left(1 - \frac{\cos\theta_h}{n\cos\left(\sin^{-1}\left(\frac{\sin\theta_h}{n}\right)\right)}\right) , \tag{6.10}$$

where t and n are the thickness and refractive index of the etalon, and θ_h is the angle of incidence between the laser and the VISAR components [158]. This results in a difference in path length which would prevent fringes from being observed with white light; however, the probe laser has a long coherence length, and thus this change is path length is not significant. Ensuring the images are coincident again ensures that the image has good spatial resolution, and also serves to increase

[5]In practice, this offset is found experimentally, and this allows the angle θ_h to be calculated. This can then be used to accurately determine τ for the etalon using Equation 6.11.

fringe contrast - making the signal easier to detect. The time delay that the etalon introduces in a VISAR of this design is calculated according to

$$\tau = \frac{2t(n - \frac{1}{n})}{c \cdot \cos\Big(\sin^{-1}(\frac{\sin\theta_h}{n})\Big)} \ , \tag{6.11}$$

where c is the speed of light [158].

The main benefit of a Mach-Zehnder design for the VISAR is that the use of two beam-splitters makes it possible to control the image alignment (and thus contrast) and the linear ramp/fringe spacing independently [157]. The 'contrast mirror' is motorised, and can be adjusted to ensure the best possible overlap between images at the recombination beam-splitter, enabling contrast to be maximised. Independently, the recombination beam-splitter can be rotated to adjust the linear phase ramp, and thus change the number/spacing of the fringes (an important factor in analysing the resulting data).

Making shock velocity measurements with VISAR

The previous descriptions describe how VISAR can be used to dynamically measure the velocity of a moving surface. Measuring shock velocity using VISAR introduces additional complexity, as the surface to be measured is a density interface within the target. To determine what exactly is measured, it is necessary to consider the surface that the probe laser will be reflected from.

The first way to measure shock velocity using VISAR is where the unshocked material is transparent to the probe beam, but the shocked material is reflective. In this case, the probe laser propagates through the unshocked material, and reflects directly off the shock front; this shock front is the measurement surface, and the shock velocity is therefore measured directly [157]. This is the case for α-quartz (the reference used in this experiment), which for pressures above $\sim$ 100GPa becomes a conductive fluid with significant reflectivity at 532 nm [159].

A different approach is required for opaque materials, such as the foam in this experiment. In this case the probe laser is reflected by the stationary front surface of the material; the VISAR thus measures the velocity of the stationary surface,

rather than the shock front. However, the VISAR can still be used to detect changes to this surface. Fringes are produced when the surface is intact, but suddenly disappear when the surface is blown-out and stops reflecting - as occurs at shock breakout. By using steps in the material of known thickness and using the VISAR to identify the time at which the shock transits the steps, it is possible to calculate the average shock velocity through this region.

VISAR application in this experiment

In this experiment the target had a step structure, so that half of the rear surface was exposed quartz, and the other was foam. The probe laser spot was positioned over this step, so that half of the VISAR image corresponds to signal from the quartz and the other half from the foam. To illustrate this, simulated VISAR data is presented in Figure 6.4. Figure 6.4 (a) shows simulated streak images of this VISAR data for two different etalon choices.

The underlying data for these simulated images was chosen to best illustrate the key features/behaviours rather than representing realistic timings or shock profiles, and additional behaviours (such as changes in reflectivity) were neglected for simplicity. A 5 ns streak time/measurement window is used. For the first 1 ns, the shock was in the ablator/gold. The shock was in the quartz for the next 2 ns, and in the foam for a further 1 ns, before breaking out of the rear of the target 1 ns before the end of the data window. The quartz shock velocity was constant at 15 $\mathrm{km\,s^{-1}}$ for 1 ns, and then decayed linearly to 10 $\mathrm{km\,s^{-1}}$ over the next 1 ns. A small quantity of Gaussian noise was included in this data. The streak images were created to match the real experimental measurements, with time increasing with decreasing y position in the image, and the x direction representing different points on the target. In this simple model the velocity was constant with position in the two regions; but as the VISAR has spatial resolution, it would also be possible to measure a spatially varying velocity profile.

The left half of each streak image in Figure 6.4 (a) corresponds to the quartz signal. Initially, the VISAR probe laser passes through the quartz and reflects

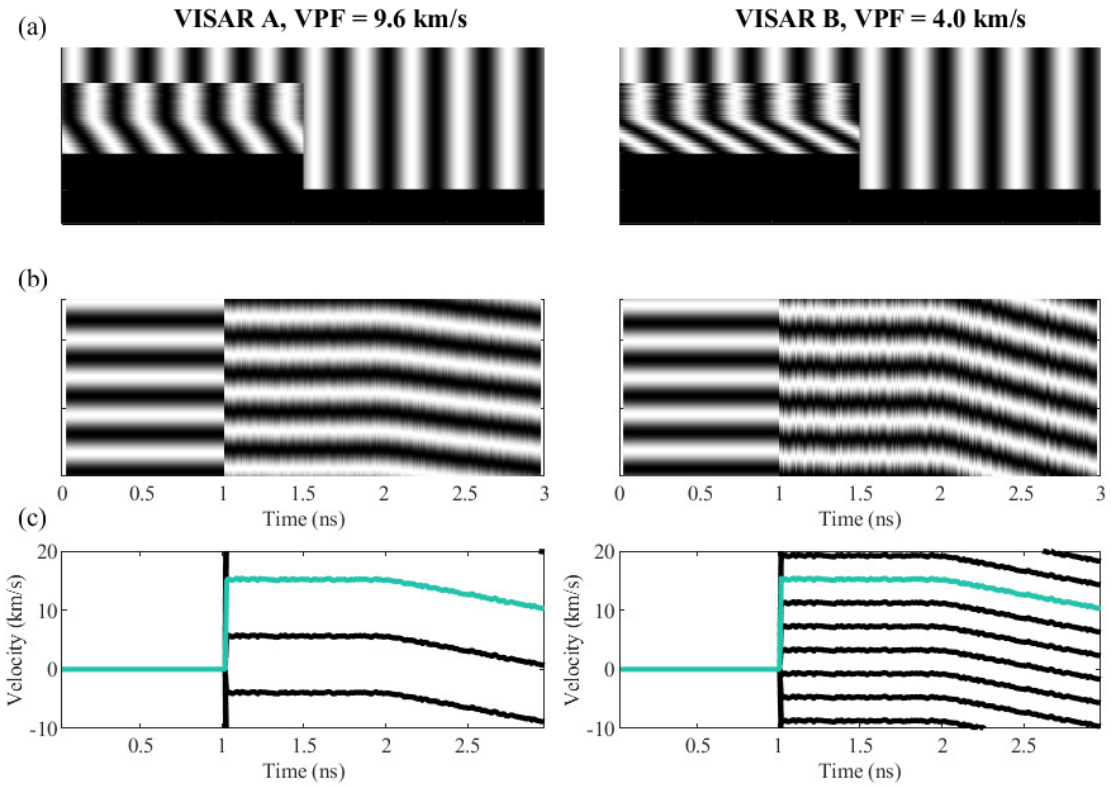


Figure 6.4: Simulated VISAR data for two different etalons. Row (a) shows the simulated streak image from each VISAR (where time increases with decreasing y position and the x dimension represents different positions on the target, as in the real experimental data). In row (b), the portion of the data corresponding to the quartz has been selected, rotated and flipped, and a timescale added based on the simulated streak window. This has then undergone the filtering and phase unwrapping procedure described in Appendix B.3. In row (c), the resulting velocity traces for each VISAR are displayed. The discontinuity could correspond to any integer number of VPFs. However, using two VISARs with different etalons allows for unambiguous determination of the real velocity, since only one of these velocity traces (displayed in teal) is common to both etalons.

from the stationary gold layer. When the shock transits through from the gold to the quartz, the laser begins to reflect from the shock front instead - leading to a sudden jump from a measured velocity of 0 to a measured velocity of U_s^{quartz} (15 $\text{km}\,\text{s}^{-1}$ in this case). The sudden change in velocity results in a sudden fringe shift. These fringes are static in the new position while the shock velocity is constant, but as the shock velocity begins to decay the fringes are seen to shift smoothly as a function of time. The right side of the image shows the foam signal. The foam is opaque to the probe laser, which continues to reflect off the foam surface - this surface is stationary, and thus the fringe pattern is static.

When the shock breaks out from the quartz, the laser on the quartz side has no surface to reflect off and thus the fringes disappear. The laser on the fringe side continues to reflect off the stationary foam surface until the shock breaks out from this surface too, at which point fringe extinction again occurs. The y position in the streak images relates to a time scale, based on the streak time, and thus the shock transit times through the quartz and foam layers are identified from these images, enabling calculation of the average shock velocity.

The quartz velocity is also obtained as a function of time through VISAR analysis of these images. In Figure 6.4 (b), the quartz section of the image has been selected and rotated, and the y-axis converted to a time-scale. The VISAR analysis procedure[6] described in Appendix B.3 is then applied to calculate the shock velocity profiles seen in Figure 6.4 (c). While the shock decay is larger than the 4 $\mathrm{km\,s^{-1}}$ VPF of the second VISAR, the extended period this change occurs over means this is easily resolved. However, the sudden change in shock velocity as the shock enters the quartz cannot be resolved and results in a fringe discontinuity. This corresponds to an infinite number of possible velocity jumps each separated by an integer multiple of the VPF, as seen for each VISAR in Figure 6.4 (c). However, only the true velocity profile is common to both etalon choices[7]. This is the teal profile displayed in Figure 6.4 (c). The shock velocity profile is thereby unambiguously determined.

6.2.3 Streaked Optical Pyrometry

Streaked Optical Pyrometry enables the temperature of a body to be determined, based on its thermal emission [48]. Emitted light from the target is captured using an objective lens. This light is broadband, but is reduced down to a narrow wavelength range using appropriate filters. This filtered emission is then recorded on a streak camera, which records the intensity vs time vs a spatial dimension.

[6]The code I used to do this was based on a VISAR analysis script provided by Professor Daniel Eakins.

[7]Mathematically, this leads to two simultaneous equations: $u = (m_1 + a) \cdot \mathrm{VPF}_1$ and $u = (m_2 + b) \cdot \mathrm{VPF}_2$, where u is the true velocity, m_1, m_2 are the appropriate integer number of fringes missed in the discontinuity, and a, b are the fractional fringe shifts measured in the two streak images. As both signals correspond to the same true velocity, the integers m can be calculated and thus the true velocity determined.

Following the derivation given in [160], the intensity recorded on a streak camera at a given pixel is given by

$$I = G \int_{\Delta t} \mathrm{d}t \int \Phi_S(\lambda) T_x(\lambda) \mathrm{SR}(\lambda) \mathrm{d}\lambda \,, \tag{6.12}$$

where G is a factor proportional to the streak camera gain, Δt is the dwell time (the time over which signal is acquired for a given pixel), $T_x(\lambda)$ is the transmission between the object and the camera, and $\mathrm{SR}(\lambda)$ is the SOP response function. $\Phi_S(\lambda)$ is the radiant power from the source. The power from the source that maps to a given pixel is calculated as

$$\Phi_S(\lambda) = \int_{A_{pixel}} \mathrm{d}A \int_{\Omega_{\mathrm{lens}}} L_S(\lambda) \mathrm{d}\Omega \,, \tag{6.13}$$

where A_{pixel} is the source area represented per pixel, and Ω_{lens} is the solid angle captured by the imaging system. $L_S(\lambda)$ is the spectral radiance of the source. For a given optical set-up of the streak camera (i.e. where the target is placed at the same location, the imaging system is not being changed, and the streak camera settings are constant), then these equations simplify to

$$I = \frac{B \Delta x W_s \Omega_{\mathrm{lens}} G}{\eta M^2} \int L_S(\lambda) T_x(\lambda) \mathrm{SR}(\lambda) \mathrm{d}\lambda \,, \tag{6.14}$$

where the variables outside the integral all relate to the specifics of the optical set-up and streak camera settings (B is the CCD binning, Δx is the dimension of one pixel on the CCD, W_s is the slit width, η is the sweep rate, and M is the magnification).

The source is then assumed to emit as a black body, such that the spectral radiance is calculated using Planck's law,

$$L_S(\lambda) = \frac{2hc^2}{\lambda^5} \frac{1}{[\exp\left(\frac{hc}{\lambda T}\right) - 1]} \,, \tag{6.15}$$

where T is the black-body temperature. The SOP signal is therefore given by

$$I = \frac{B \Delta x W_s \Omega_{\mathrm{lens}} G}{\eta M^2} \int \frac{2hc^2}{\lambda^5} \frac{T_x(\lambda) \mathrm{SR}(\lambda)}{[\exp\left(\frac{hc}{\lambda T}\right) - 1]} \mathrm{d}\lambda \,. \tag{6.16}$$

The incoming light is filtered down to a narrow wavelength band. The Fourier spectrum of the incoming light is therefore approximated as a delta function, at the central wavelength of this band λ_0. This simplifies the SOP signal to

$$I = \frac{B\Delta x W_s \Omega_{\text{lens}} G}{\eta M^2} \frac{2hc^2}{\lambda_0^5} \frac{< T_x \text{SR} >}{[\exp\left(\frac{hc}{\lambda T}\right) - 1]} , \tag{6.17}$$

where $< T_x \text{SR} >$ is the averaged product of the transmission and response functions over this wavelength range. This equation is then rearranged to give an expression for temperature in terms of the measured intensity,

$$T = \frac{T_0}{\ln\left(1 + \frac{A}{I}\right)} , \tag{6.18}$$

where $T_0 = \frac{hc}{\lambda_0}$ is a constant, and A is a calibration constant,

$$A = \frac{B\Delta x W_s \Omega_{\text{lens}} 2hc^2 < T_x \text{SR} > G}{\eta M^2 \lambda_0^5} , \tag{6.19}$$

which is constant for a constant experimental set-up and streak camera settings.

Real samples do not emit as perfect black-bodies. This can be accounted for by factoring in the emissivity of the body [161], which can be calculated as $\eta = 1 - R$, where R is the reflectivity (from Kirchhoff's second law, stating that absorption and emissivity are equal for an object in thermal equilibrium [48]). To account for this in the above derivation, the SOP streak intensity I is replaced in every instance by $I/(1-R)$. This results in an equation for the grey-body temperature of the sample of

$$T = \frac{T_0}{\ln\left(1 + \frac{(1-R)A}{I}\right)} . \tag{6.20}$$

The SOP diagnostic can be absolutely calibrated for a fixed set-up, where the components are not regularly changed. This process is outlined for the OMEGA-60 set-up in [161]; by making measurements over a series of wavelengths using an absolutely calibrated light source, the different wavelength dependent functions (such as spectral response and transmission) of the camera are determined. The value of A is then calculated with confidence for different wavelengths and different camera settings.

SOP in this experiment

In this experiment, SOP was used to measure the signal from both the quartz and foam. However, as the SOP was built from scratch during the experimental run there was not sufficient time nor equipment for an absolute calibration to be performed. Instead, the quartz was used as a known temperature reference, and calibration was performed on-shot.

To do this, the expected quartz shock temperature is required. This was estimated from the VISAR-measured shock velocity using a power law fit,

$$T(K) = 1421.9 + 4.3185 \cdot U_S^{2.9768} \,, \tag{6.21}$$

published in the supplementary material for [162] based on experimental data in [159]. The corresponding SOP intensity signal was thereby used to calculate the on-shot calibration constant A in equation 6.20. With A for the shot determined, the foam SOP intensity was therefore used to calculate the foam temperature. This on-shot calibration necessarily leads to greater uncertainty, as the error in the shock velocity measurement and the above model both affect the calculated foam temperature.

6.2.4 X-ray pinhole camera

As the Vulcan beam irradiated the target, high energy x-rays were emitted in all directions; an x-ray pinhole camera with a 50 µm pinhole was used to detect these and thus image the laser pulse focal spot. A basic schematic of a pinhole camera is shown in Figure 6.5. In such cameras, an x-ray detector is placed a distance L behind a small aperture (or 'pinhole'), which in turn is placed a distance D from the object being imaged. Only x-rays passing through the aperture reach the detector. The magnification M is equal to the ratio $M = L/D$, where L is the length of the camera and D is the distance from the object to the pinhole. In this experiment the magnification $M \sim 2.5$ (with lengths $L \sim 400$ mm and $D \sim 160$ mm). An image of the camera used is shown in Figure C.5 in Appendix C.

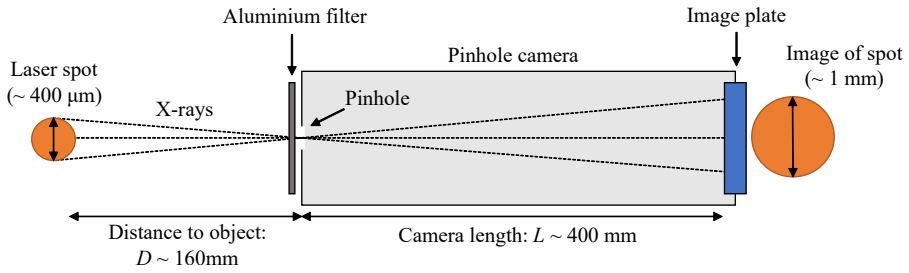


Figure 6.5: A schematic of a pinhole camera. The lengths D and L set the field of view and the magnification. Appropriate values have been included for the camera used in this experiment.

A Fuji BAS image plate [163] was used to detect the x-rays, and was scanned after each shot to obtain the two-dimensional x-ray image[8]. It is possible to quantify the x-ray signal received by the image plate, but such an analysis was not necessary for this experiment. Instead, the pinhole camera was used to indicate the spatial characteristics of the Vulcan spot. This was useful for confirming where the laser spot hit on the target, that there was a good overlap between the six beams, and to estimate the focal spot size.

6.2.5 Fiducial

A fiducial was set-up on one of the VISAR streak cameras. A fibre-optic cable was set-up behind one of the mirrors in the Vulcan laser chain. These mirrors were slightly lossy, and the fibre 'picked-off' some of this parasitic signal. The other end of the fibre was positioned so that it produced a small dot on the side of one of the VISAR streak cameras. The fibre length was such that this parasitic Vulcan signal reached the streak camera in the same streak window that the main signal from the target was detected. This set-up can be seen in Figure C.8 in Appendix C.

Two-minute shots (low power Vulcan shots, so-named as they can be fired every two minutes) were performed without a target present, so that the Vulcan pulse could be measured directly on the streak camera. By doing so, the time delay between the Vulcan shot reaching the target and the fiducial signal showing on the streak could be determined. This meant on subsequent full power shots, where the Vulcan

[8]A description of the operation of such image plates is provided in [164].

pulse could not be seen directly due to the presence of the target, it was possible to calculate the time at which the pulse was applied based on the fiducial signal.

6.2.6 Calorimetry

Calorimeters, which returned a voltage signal based on the fluence of light received, were in place behind a leaky mirror in each of the six beams. These were calibrated to provide an estimate of the energy in each beam for each pulse (one of these was partly blocked by the fiducial, and so returned a low signal - as the six beams were delivering similar energies, an average of the other beam energies was used for this beam when estimating the total pulse energy).

The Vulcan beam is generated at IR, and had to pass through a frequency-doubling crystal to be converted to 527 nm. The facility staff performed a series of calibration shots. On these shots, larger calorimeters were set up in the main beam path to measure the full energy in each beam, while the orientation of the crystal was optimised. This also allowed for calibration of the smaller, on-shot calorimeters; the voltage the parasitic calorimeters measured as a function of the energy measured by the full calorimeters were recorded, and used to find the relationship between voltage and energy, so that the energy of the beam on shot could be calculated.

6.2.7 Pulse length

Photodiodes were placed behind the final set of mirrors in the Vulcan chain (a different set of mirrors to the calorimeters). The mirrors can be seen in Figure C.9 in Appendix C. These photodiodes measured the pulse shape of each beam on every shot, and saved it to an oscilloscope - allowing the pulse length and temporal shape of each beam to be observed.

6.3 Preparation for the experiment

6.3.1 Development of the optical set-up

An optical set-up was required to collect and transmit the self-emitted light to the SOP. Further optics were also required to transmit the probe laser to the rear of

the target, and to collect the reflected beam and deliver it to the VISAR.

Both VISAR and SOP were operated normal to the rear target surface. While it is possible to operate VISAR off-axis, this has not previously been done in experiments at these intensities (and would also require a correction to the data). Operating both diagnostics in this configuration meant that a shared optical system would be required to transmit the light from the target to the diagnostics, with a dichroic mirror used to separate the reflected 532 nm VISAR probe laser from the self-emission. Similar set-ups are used on OMEGA-60 [160, 161].

This optical system had to minimise transmission losses to ensure a strong signal at the diagnostics. There was a risk that the self-emission could be relatively weak, and thus it was desirable to avoid placing any beam-splitters in the shared optical relay. This meant that the probe laser injection needed to occur prior to the dichroic mirror, resulting in the basic experiment design displayed in Figure 6.6.

A number of factors were also considered in the design of this system:

- **Optical relay spacing**: The spatial configuration of the target chamber and area set limits on the optical relay design. Injecting the probe laser through the VISAR beam-splitter required the laser to be located relatively close to the diagnostics (to avoid a long beam path around the target area). This set constraints on the positions of the optical benches in the target area, and meant that it was necessary for the optical relay to transmit the light through the south chamber window (which was furthest from the target position). This window was 1.5 m away from the target position, with a further 1.3 m of window/free space propagation before the optical table would be reached.

- **Magnification**: The SOP streak camera had a slit width of 4.46 mm, while the VISAR streaks both had a width of 17.4 mm. The imaged area was 400 µm in diameter, and needed to be magnified to fill as much of the camera slit as possible. The VISAR already had a number of built in lenses (of which only the first could be changed), and so it was necessary to design a system

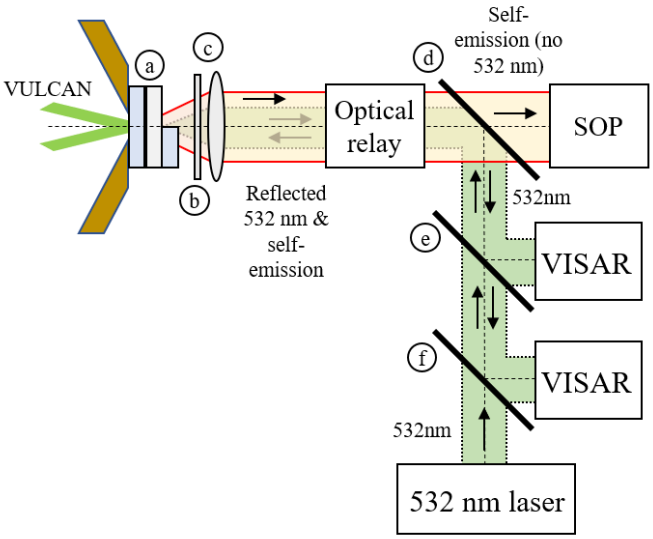


Figure 6.6: A simple schematic of the initially proposed experimental layout. Reflected 532 nm light from a probe laser, and the self-emission, are captured from the target (a) by an objective lens (c) protected by a blast shield (b). This captured light passes through an optical relay (series of lenses), which magnifies it and transports it to the diagnostics. A dichroic mirror (d) reflects the 532 nm light out from the broadband light. This reflected 532 nm light is split and sent to the two VISAR diagnostics via two beam-splitters (e and f). The remaining broadband light passes through the dichroic mirror to the streaked optical pyrometer. The 532 nm laser is injected through the back of the second VISAR beam-splitter, and travels through the full optical relay to illuminate the target. This set-up avoids the need for the self-emission to pass through a beam-splitter on its route from the target to the SOP.

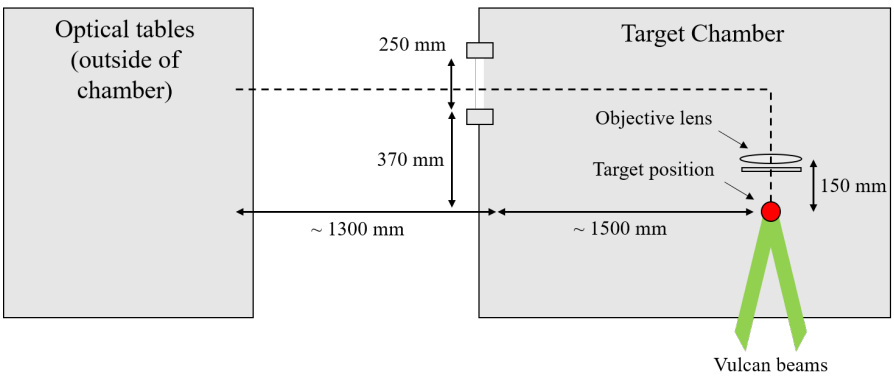


Figure 6.7: A simple schematic (not to scale) showing the spacial considerations for the optical relay. There was a useable distance of $\sim$ 1500 mm of optical table within the chamber before the window. There was then a further $\sim$ 1300 mm gap (in which no lenses could be placed) before the optical tables outside of the chamber, which housed the diagnostics.

which would provide appropriate magnification for both diagnostics as well as transmitting the beams with minimum losses over the relevant distances.

- **Components**: The optical components required for the relay and for the diagnostic systems also needed to be identified and purchased. This included: a dichroic mirror, low-pass and high-pass filters for the SOP, and notch filters (to prevent any 527 nm Vulcan noise affecting the VISAR signal). Many of these components, as well as the optical components within the VISAR, were only available in 1-inch versions - which set limits on the beam size.

- **Diagnostic positioning**: Finally, it was necessary to consider how the diagnostics and beam paths would actually fit on the optical table. The VISARs in particular were large pre-assembled systems with fixed entry points for the beam path. The 1-inch optical components meant it was necessary to minimise the distances between the VISARs and the final lens in the optical relay to avoid transmission losses, but this was challenging to achieve practically.

To assist with this design work, a `MATLAB` script was produced to perform ray-tracing (using simple geometrical optics) of light through a lens system. A number of set-ups were explored using this code to find a system which minimised losses while being compatible with the above considerations. The final design achieved this, with transmission (according to the simple approximations used within the code) above 99% from the first lens to the target. Details of the operation of this code are provided in Appendix B.2.

Figure 6.8 shows a full schematic of the experiment. The optical relay consisted of 5 lenses (L1 to L5). The dichroic mirror then took the collimated beam, and sent the signal to the SOP/VISARs. The SOP consisted of a 400 nm long pass filter and 500 nm short pass filter, leaving only 400 - 500 nm wavelengths, before a final lens S1 focussed the beam on to the streak camera. Beam-splitters sent the VISAR beam into the VISAR system (the 'VISAR' here is used to refer to the full system shown in Figure 6.3, rather than just the interferometer). The beam passed through a notch

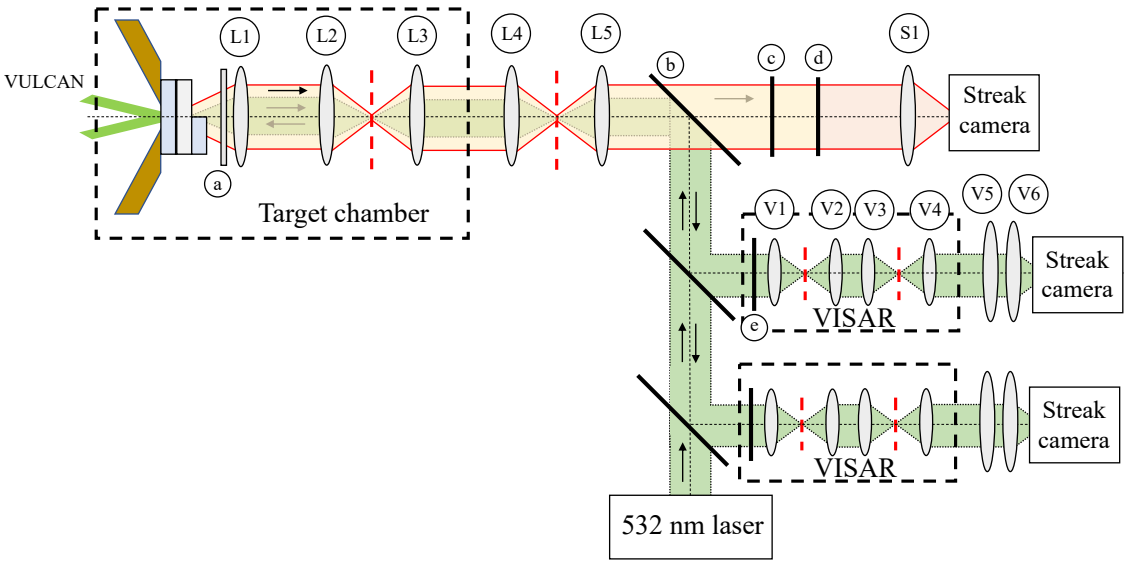


Figure 6.8: The full proposed schematic. Lenses labelled 'L' form the optical relay, while those labelled 'S' are part of the SOP and 'V' are part of the VISAR. The two VISARs are identical in terms of components. The dashed red lines represent the positions of image planes. The interferometer in the VISAR is not pictured (it can be viewed in Figure 6.3), but is positioned so that the second beam-splitter is located at the second image plane. The focal lengths are as follows: L1 = 150mm, L2 = 700mm, L3 = 500mm, L4 = 700mm, L5 = 250mm, S1 = 300mm, V1 = 400mm, V2 = 250mm, V3 = 1000mm, V4 = 300mm, V5 = 300mm, V6 = 200mm (cylindrical). Also indicated are the other components in the beam line: a blast shield to protect the objective lens (a), the dichroic mirror (b), the short-pass (c) and long-pass filters (d) for the SOP, and notch filters (e) for the VISARs.

filter, followed by the 4 lenses within the VISARs (and the interferometer, which is not shown), before being focussed onto the beam splitter by lens V5 (and further focused in the vertical direction by the cylindrical lens V6). The lens spacings were significant for preventing losses, and are shown in the technical drawings in Appendix C. The spacing between focusing pairs were fixed by the focal distances of the lenses. There was some flexibility in the spacing between L1 and L2, as well as between L3 and L4; however, the distance between L5 and V1 for each VISAR had to be kept below 1 m to avoid losses (otherwise the beam would be too large for the 1 inch components inside the VISAR), and the distance from V4 to the VISAR streak camera also needed to be under 1 m. The short pass filter (c) was also a 1 inch component, and so it was important to place this close to the dichroic mirror.

Figure 6.9 is a CAD model of this experimental set-up (produced by Michael

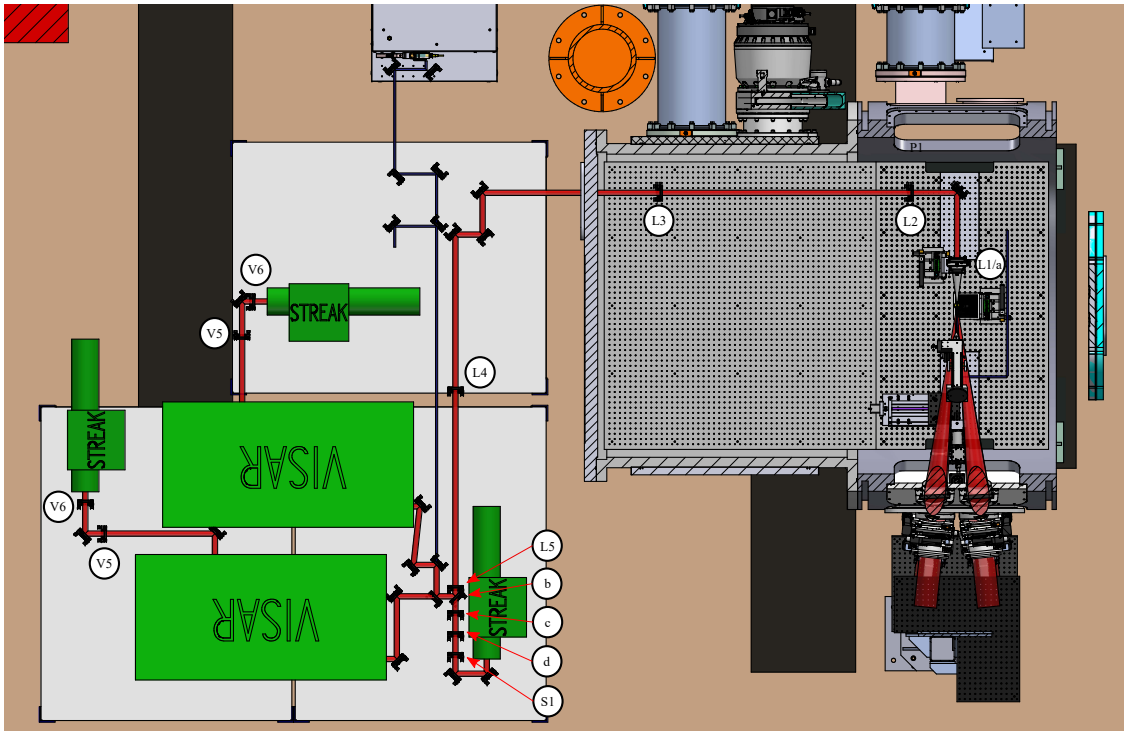


Figure 6.9: Scale CAD model of the experimental set-up. The key components are labelled as described in Figure 6.8; mirrors and beam-splitters are unlabelled. The beam line in blue is the probe laser, which enters the main beam line (red) through a beam-splitter. There are two 'branches' from the beam lines that do not terminate in a component. These represent where diode lasers were used for alignment purposes, and flip mirrors were used to insert these beams into the main beam paths.

Woodward of UKRI-STFC), which shows how this set-up was arranged spatially in order to satisfy all the requirements discussed previously. Figure 6.10 then shows the output of the ray tracing script for the target to SOP and target to the second (furthest) VISAR, for the optical path lengths corresponding to this experimental arrangement.

6.3.2 Target design and hydrodynamic simulations

Layer choice and thickness

The target was based on multi-layer step target designs successfully used in previous experiments [138, 139]. The initial proposal used an aluminium layer as the reference, but this was replaced with quartz so that the velocity of the shock could

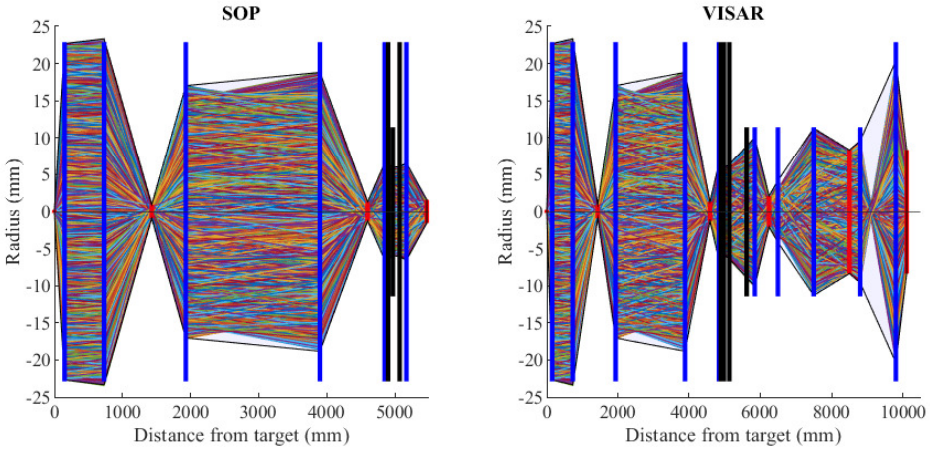


Figure 6.10: Ray tracing simulations from the target through to the SOP and VISAR. The components are placed with the real experimental spacing. Lenses are represented by blue lines, while other components (dichroic mirror, beam-splitters, filters) are represented by black lines. 10,000 rays are generated, each at a random position on the target and with a random angle. These rays are propagated through the system according to simple geometric optics through thin lenses. At each component position, the radial position of each beam is compared to the component size; if it is bigger, the ray is considered 'lost'. This approach suggests transmission through the system greater than 99% for both the VISARs and SOP (with the small losses occurring between L1 and L2).

be measured directly using the VISAR[9]. The four layers and step structure led to a relatively complicated target.

The thicknesses of the different layers were influenced by a number of factors. These are listed below:

- **Vulcan pulse length/energy capabilities**: In order to ensure a sustained shock, it was intended that the laser pulse should last until the shock broke out from the rear of the target. There was an upper limit on how long Vulcan could sustain a pulse for (particularly at high intensity, where high energy is required), and this therefore prevented very thick targets with long shock propagation times.

[9]The original proposal had suggested that the VISAR could be used to measure the free surface velocity of the aluminium and foam. Under this approach, the VISAR reflects from the rear of an opaque target. The target is shocked, and the shock wave reaches the rear surface, which undergoes free surface expansion. The free surface velocity is measured by the VISAR; this is generally approximated as being twice the particle velocity within the shocked medium. This is a common approach at lower pressures, but it is not possible in this regime - as the shocked material becomes a plasma, the rear surface is destroyed and thus the free surface velocity cannot be measured. As such, the target design was changed to replace the aluminium with quartz, and the foam shock velocity was instead determined from the transit time.

- **Secondary shocks**: Hydrodynamic simulations showed that when the shock crossed from the gold layer into the quartz, a release wave would propagate back through the target. Upon reaching the ablation front, this then generated a second shock which travelled through the target. Depending on intensity/target thickness, this secondary shock could catch up and merge with the primary shock. This affected the shock strength, preventing accurate use of the impedance matching calculation. This effect was heavily influenced by the thickness of the different layers, and needed to be avoided in the real experiment. An example of this behaviour can be seen in Figure 7.26, in Chapter 7.

- **Fabrication constraints**: Target fabrication indicated that the 40 µm was the minimum thickness of quartz/foam that they were able to produce.

- **Preheating**: The gold layer needed to be sufficiently thick to prevent preheating of the quartz/foam.

- **Avoiding direct laser heating of gold**: The ablator needed to be sufficiently thick so that the ablation front remained in this layer throughout the laser pulse, and thus the gold did not undergo direct laser heating.

`HYADES` simulations were performed to find the optimal target design. The minimum 40 µm thickness for the quartz and foam was chosen to mitigate the pulse length concerns. This also was required for the second shock; the thicker the quartz/foam layers were, the more distance there was for the second shock to catch the primary shock in. Even for 40 µm quartz/foam thickness the second shock was already a significant problem.

Increasing the gold and ablator thicknesses were both found to delay the second shock, improving this issue (the gold/quartz interface was moved further from the ablation front, which meant the second shock (which has to travel this distance twice once it is generated at this interface) had a longer distance to travel relative to the first). Simulations were performed investigating the potential gold/ablator

thicknesses that could be used, and the impact this would have on the required laser pulse time/energy at different intensities.

The selected design used a 40 µm ablator, and a 3 µm gold layer. In `HYADES`, this led to the required pulse times/pulse energies as a function of intensity seen in Figure 6.11 - remaining under the 600 J maximum energy of Vulcan for the full planned intensity range[10]. The target did not experience second shock merger over this intensity range. The simulations also suggested that the ablation front remained in the ablator for all of these intensities, and that the 3 µm gold layer would prevent preheating of the foam by less than 10 K. Hydrodynamic simulations cannot be expected to describe preheating with high accuracy as they do not include many of the necessary laser plasma effects; however, previous experiments have used similar gold thicknesses and did not report substantial levels of preheat, suggesting that this thickness was sufficient.

Additional considerations

There were also a number of practical concerns with the target assembly. It was important to avoid glue layers in the targets; a glue layer between the quartz and foam would prevent impedance matching from being performed accurately, and glue wicking into the foam could change the foam properties. In addition, a glue layer between the gold and the quartz could prevent clear identification of when the shock entered the quartz. As such, the ablator and gold coatings were grown directly on to the quartz, removing the need for glue between these layers. The foam was tacked onto the quartz using glue at the corners only; as the foam and quartz was much larger than the shocked/imaged region, this meant that no glue was present in the area of interest. An anti-reflection coating was also applied to the quartz, to prevent spurious reflections and 'ghost fringes' (where reflections from other surfaces lead to interference and persistent fringes in the data). The

[10]In the spherical hydrodynamic simulations in Chapter 4 a multiplier is applied to the incoming laser energy to account for losses. These are likely reduced for this set-up, but some losses would still be expected. The simulations presented in this section do not include such a multiplier, and so it's likely that the shots would require a larger energy/pulse time for a given intensity. However, as the highest intensity in Figure 6.11 requires only 500 J as opposed to the 600 J Vulcan can deliver, there is some additional energy available to account for this.

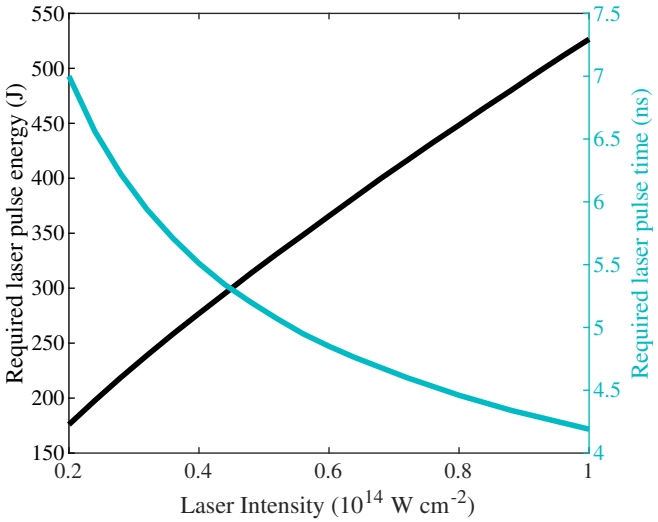


Figure 6.11: The required Vulcan laser pulse time (teal) and pulse energy (black), as a function of laser intensity. The Vulcan pulse was planned to be maintained for the full duration of the shock propagation through the target; the required pulse time thus decreases with intensity, since higher intensities lead to stronger shocks which breakout from the rear of the target earlier. However, the increase in power that the higher intensities require leads to an overall increase in pulse energy as a function of intensity, despite this shorter pulse length.

manufacture of the targets was handled by CLF target fabrication (particularly Chris Spindloe, Sam Irving, David Haddock, and Donna Wyatt).

6.3.3 Hydrodynamic simulations

The hydrodynamic simulations mentioned in the previous section also enabled the expected conditions in the quartz/foam to be investigated. A selection of the key shock variables can be seen in Figure 6.12. These simulations were performed in `HYADES` and were used to inform etalon selection, as well as confirming that the quartz pressure was sufficient for a reflective shock front. Experimental collaborators also ran simulations of this in `MULTI`, `HELIOS`, and `FLASH`, and good agreement was observed (a comparison between these codes can be seen for post-experiment simulations in Figure 7.24 in Chapter 7).

Two-dimensional simulations of this set-up were also performed to ensure that the shock was planar and that similar shock behaviour was observed in higher dimensions. The simulations presented in this section were performed in `h2D` (the

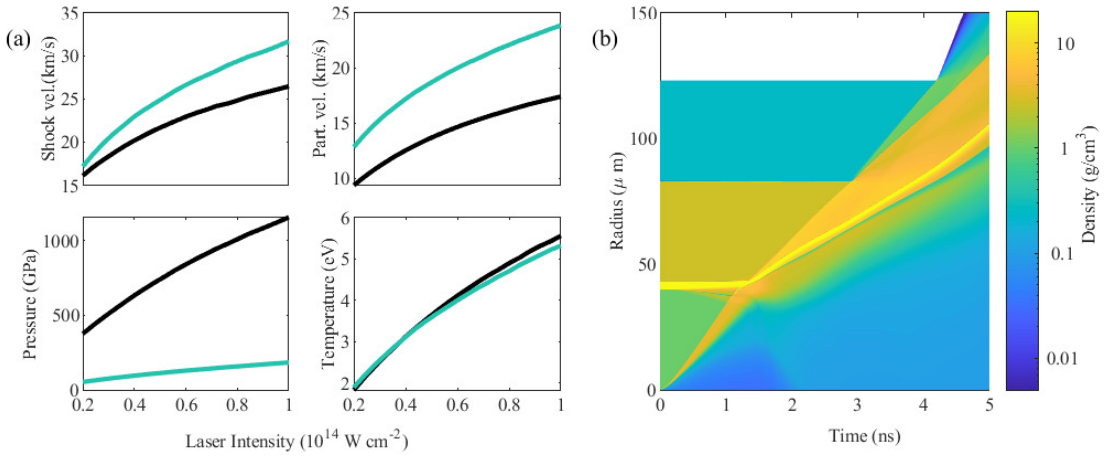


Figure 6.12: Results from HYADES simulations showing a) shock conditions in the quartz (black) and foam (teal) as a function of intensity, and b) shock propagation through the target for an intensity of 1×10^{14} W cm^{-2}. The quartz shock velocity informed decisions on the required etalons. The shock pressure in the quartz is also confirmed to be above the threshold for a reflective shock for the full range of simulated intensities.

two-dimensional version of HYADES). This is a Lagrangian code and as such suffers from mesh tangling, where the movement of the simulation zones can cause the mesh to become distorted and prevent the simulation from continuing. This requires the user to manually rezone the problem, and frequently resulted in the simulation code crashing. After a number of iterations a sufficiently detailed simulation was ran to completion. Other 2D simulations were performed in MULTI and FLASH by the other collaborators and were used to confirm these results; as these codes were ran with more success they were used instead of h2D for the post-shot simulations, and will be discussed in more detail in that section.

In h2D, a simple four layer target was simulated without a step in the foam. This was sufficient to measure shock propagation and to estimate the shock variables, as well as to indicate the shock planarity. 2D FLASH simulations were used to confirm that the step did not have a significant impact on the shock behaviour (a post-shot 2D FLASH simulation with this structure is shown in Figure 7.23 in Chapter 7). h2D is an axisymmetric code, and so the target is simulated as a cylinder where the shock propagates along the long axis. R=0 is the axis of symmetry, and the 2D plane can be rotated around this axis to produce the full target (this also means that any off-axis lasers represent rings of laser light when extruded to 3D). The

lasers were simulated as 3 beams at -25°, 0°, and 25° to the target normal. Given the real experimental configuration (three pairs of two beams) a simulated set-up consisting of two beams at 6° and -6° would also have been valid, but the three beam set-up was chosen as it included larger deviations from target normal, and was thus more likely to exhibit changes from 1D behaviour.

Figure 6.13 shows a number of schematics representing the simulated set-up, and how it compares to the real experimental configuration. Figure 6.13 (a) shows the real laser set-up (not to scale), with the true 3D configuration of the six beams. This cannot be fully described in 2D. Figure 6.13 (b) shows the 2D simulation. Three laser beams are simulated by specifying a number of 'rays' with an initial angle and position. The rays were created using a 'random-to-uniform' lens-to-target mapping. 10,000 uniformly spaced points were generated over the laser spot at the target, and for each point an incident ray was generated. Each ray was given an origin position at random, corresponding to a position on the lens for one of the three laser beams (considering the real lens size and lens-to-target distances). This method ensures that uniform illumination of the laser spot is achieved with the finite number of rays, while also using a random mapping to ensure behaviours such as defocussing past the laser focus. The energy of the ray was scaled according to the square of its radial position at the target; this accounted for the axisymmetric symmetry, and the fact that a ray incident on a higher radial position is spread over a much larger area if the 2D slice is converted to a 3D cylinder. While only positive radii are simulated (as the simulation assumes rotational symmetry around the axis), as the rays are described at a single point (on the target face), it is possible to include rays which would have originated on the other side of the axis by specifying an appropriate angle. This is shown in Figure 6.13 (c) where the simulation is reflected around the axis, and it is shown that the rays shown in (b) do in fact originate from all 3 beams. Finally, Figure 6.13 (d) shows the rotation of the simulated set-up in (b) around the axis. Here, it can be seen that this rotation means that the off-axis beams in fact represent a hollow cone of incoming laser light.

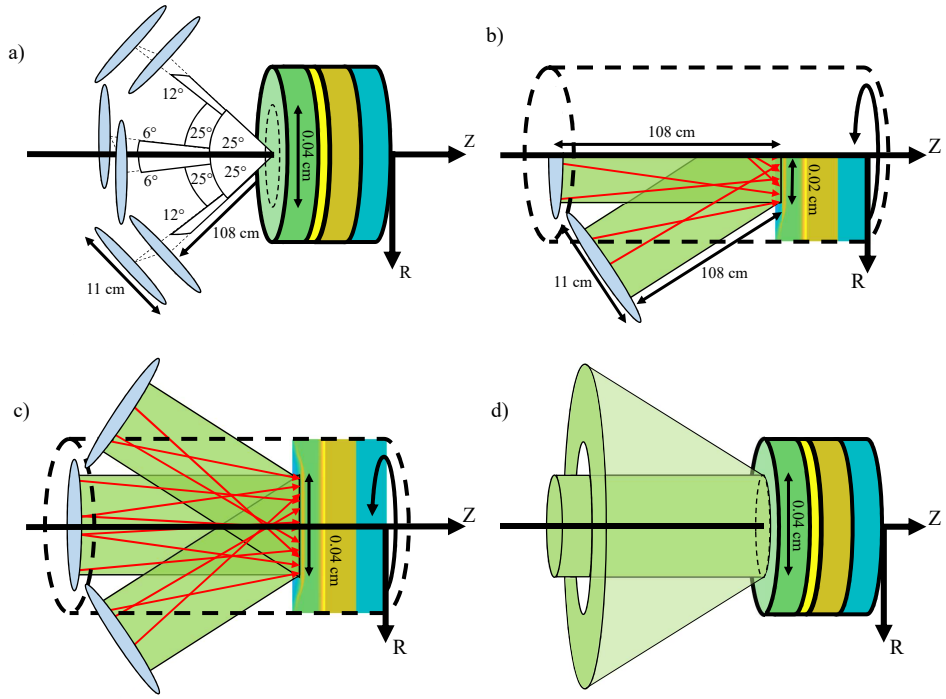


Figure 6.13: A schematic (not to scale) showing how the set-up is simulated in `h2D`. a) The real experimental layout of the 6 beams. The beams are arranged in three planes at -25°, 0°, and 25° to the horizontal, with each plane containing two beams at 6° and -6° to the vertical. Each beam is 11 cm in diameter at the lens, which is 108 cm from the target, and is focused to a 400 µm meter spot. b) The `h2D` simulation. This is done in 2D with cylindrical symmetry. The laser is described using rays, which are mapped from a random position on one of three lenses (at -25°, 0°, and 25° to the horizontal, at the correct distance) to a uniform grid of points in the laser spot. c) The `h2D` simulation has been reflected around the x-axis, and the rays traced back to show their sources. As the rays are generated in the laser spot, the configuration in (b) allows for the presence of rays that would have originated from all three lenses (even if they have to cross the axis to reach the target). d) The 3D extrapolation of the 2D simulation. As the simulation has cylindrical symmetry, the two beams at -25° and 25° actually form a laser 'ring', with a hollow cone of incoming laser light.

Figure 6.14 (a) shows four snapshots of the 2D shock structure. It can be seen that the shock front is relatively planar over a nearly 200 µm radius (and thus the 400 µm diameter of the imaged region). Figure 6.14 (b) shows the shock propagation through the target (at the position of the dashed black lines in (a)) as a function of time, which shows good qualitative agreement with the 1D simulation seen in Figure 6.12.

Figure 6.15 shows how the pressure just behind the shock front in this simulation

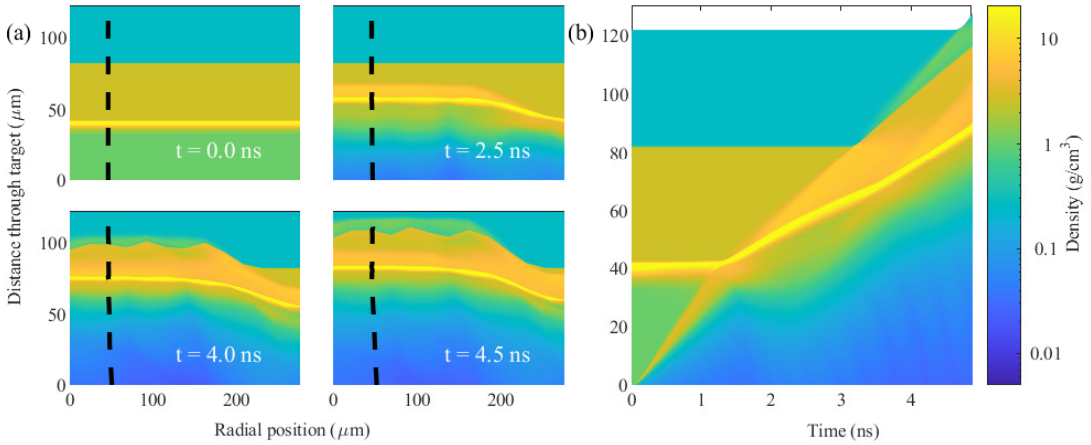


Figure 6.14: A two-dimensional `h2D` simulation of the set-up using a simple four-layer target (ablator, gold, quartz and foam without the step structure), for a laser intensity of 1×10^{14} W cm^{-2}. (a) shows four snapshots of the 2D simulation at the indicated times, while (b) shows the shock propagation through the target as a function of time for a given radial position, represented by the black dashed lines in (a). It can be seen that for the later two times, the black dashed line is noticeably curved. This is because the Lagrangian mesh moves as the simulation occurs, and thus the mesh points (where the simulation parameters are returned) do not have fixed radii. This does not significantly affect the results - particularly as it is only seen in the ablated material far behind the shock.

varies across the radius of the target. It is clear from the figure that, while the pressure reduces towards the edge of the target, over the central region the shock pressure is relatively planar in both the quartz and the foam. Meanwhile, Figure 6.16 displays the shock breakout times from the gold, quartz and foam layers. This again indicates that the shock is planar over the central region of the laser spot.

Both Figure 6.15 (a) and Figure 6.16 also compare these results to the 1D `HYADES` simulation shown in Figure 6.12. It can be seen that the timing agreement is reasonably close, but there is a small increase in the shock transit time in both the quartz and the foam. This increase in transit time is due to a decreased shock velocity, which in turn indicates a lower shock pressure as seen in Figure 6.15. While the change in shock transit time is relatively small, the corresponding impact on the pressure is significant. Interestingly, as can be seen in the post-experiment simulations, such an effect was not observed to the same extent in the simulations performed in other 2D codes. This effect (particularly the pressure) was found to vary with the resolution of the `h2D` simulations; unfortunately it was not possible

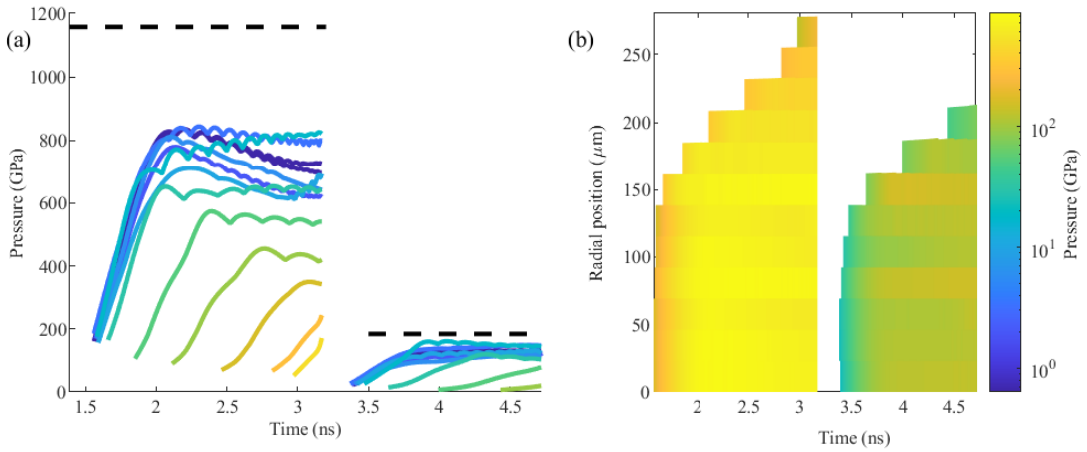


Figure 6.15: Plots showing the shock pressure in the simulation displayed in Figure 6.14 as a function of both time and radial position. In (a), the different lines represent the different radial positions, with the colour transitioning from dark blue through to yellow (i.e. up the colour bar in Figure 6.14) as the radius moves further away from the centre of the target. It can be seen here that the pressure is relatively uniform over the central regions, but drops off towards the edge. The dashed line represents the average shock pressure obtained in both materials in the 1D simulation. It is clear that this is significantly higher, but the difference appears to depend on simulation parameters (such as resolution) and was not observed in the other 2D codes. (b) represents the same data, but here the pressure is represented by colour on a 2D plot of position vs time. On both plots, the shock front pressure is only recorded when the shock front can be automatically identified in the simulation, and thus there are blank regions when the shock crosses material interfaces and the shock front cannot be clearly identified.

(due to persistent issues with mesh tangling and crashing) to increase the resolution of the simulation to a point where the results converged. As such, while the `h2D` simulation was useful for confirming the shock planarity and expected behaviour in 2D, it was decided that the results should not be relied on for the precise shock timings/pressures, and the other codes (such as `MULTI` and `FLASH`) were used for the post-experiment 2D simulations.

6.4 Conducting the experiment

The experiment was performed over a five and a half week period. Roughly two and a half weeks of this were set-up, with a further three weeks of experimental shots. A series of images showing the experimental configuration can be seen in Appendix C.

Throughout the run a number of changes were required from the original set-up and experimental plan. These are outlined and explained in this section.

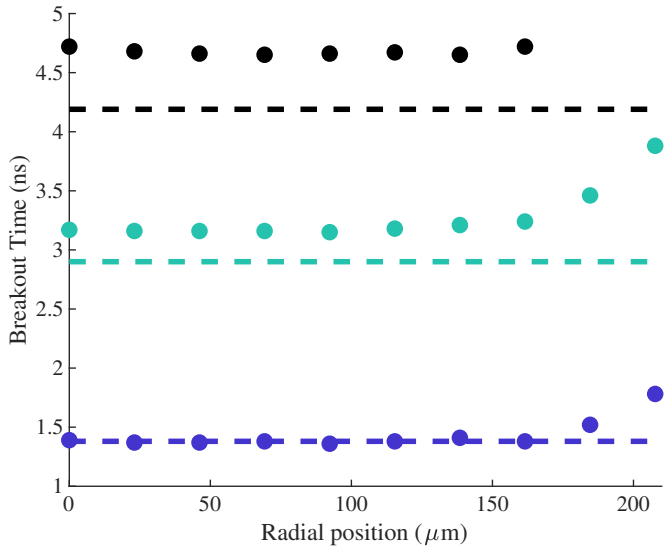


Figure 6.16: Shock breakout times from the gold (blue points), quartz (teal points) and foam (black points) as a function of radial position in the target. It can be seen that the breakout time is roughly steady over a radius of 150 µm, and then increases towards the edge of the laser pulse. The dashed lines represent the breakout times obtained from the 1D `HYADES` simulation. While there is good agreement for the gold breakout, the 2D simulation returns a later breakout from the quartz and the foam, suggesting a slower and weaker shock. This is not observed in the other 2D simulations.

6.4.1 Changes from planned target

The actual targets that were delivered for the experiment differed slightly from the planned targets. This was due to challenges in fabricating the targets, as well as one of the suppliers failing to deliver a shipment of quartz (which also limited the number of available targets).

- **Different target dimensions**: the delivered quartz was thicker than expected, which led to the quartz layer in each target being $\sim$ 50 µm thick (rather than the planned 40 µm). The foam thickness also varied, and was frequently significantly larger than the planned 40 µm.

- **Fewer targets with an AR coating**: the reduced quantity of quartz meant that some targets did not have the planned anti-reflection coating. Fortunately, this did not seem to have a significant effect.

- **Delamination in some targets**: in a small number of early targets, the ablator/gold delaminated from the quartz, and had to be glued back on. The

glue layer prevented an accurate timing measurement in the quartz and meant that the target could not be used for impedance matching; however, they were shot as test/set-up targets. These later became important in understanding some features of the results (see Section 7.8).

6.4.2 Poor VISAR illumination and subsequent changes to optical set-up

Once set-up was complete and shots started, it became apparent that the probe laser was only illuminating a small section of the imaged region. This prevented sufficient VISAR data from being collected. This was due to the fact that the probe laser travelled through the optical relay before reaching the target, which demagnified it and brought the laser to focus at the target position, as seen in Figure 6.8.

An initial solution was to place a 700 mm focal length lens between the laser and the VISAR beam-splitter; this meant that the laser was not collimated when entering the relay, and so was not perfectly focussed by the objective lens. This did not sufficiently improve the illumination. In response, a modified set-up was designed which would solve the problem while requiring minimal changes to the optical relay (which could not be realigned due to time constraints). In the new set-up, the first mirror after the objective lens was replaced with a beam-splitter, and the probe laser was injected through this beam-splitter rather than after the VISARs. This meant the probe laser bypassed the relay, and only the objective lens was shared. A lens was placed in the probe laser beam path prior to this beam-splitter to form a lens pair with the objective lens; this meant that the probe laser was collimated by the objective lens, and was thus collimated when it reached the target[11]. A simple schematic showing this new set-up can be seen in Figure 6.17, with a CAD model[12] showing the new spatial configuration shown in Figure

[11]While this was initially the case, extra mirrors were later added to the set-up to give additional control over the probe laser position. Spatial constraints meant these mirrors had to be added between the lenses, which increased the P1-L1 distance from the intended 550 mm to 875 mm. This meant that the beam was no longer perfectly collimated. However, burn paper tests showed that the spot size at the target position was still sufficient, and the illumination remained uniform. Simulation of this set-up with the ray tracing code suggested the spot size was around 1 mm in diameter.

[12]This CAD model was again produced by Michael Woodward.

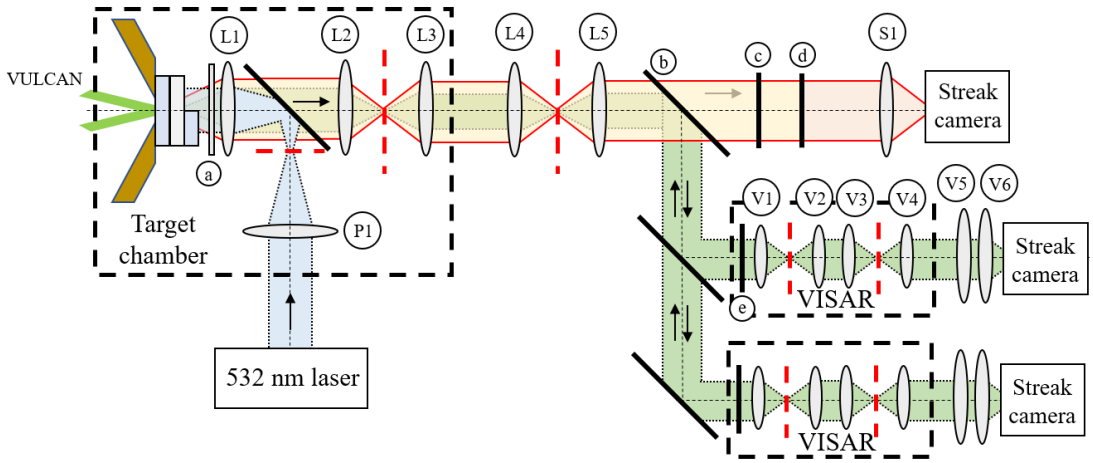


Figure 6.17: The full schematic, with the new probe laser injection. The probe laser is displayed in blue to differentiate the beam path from the reflected probe light shown in green, but both have the same 532 nm frequency. The labels are unchanged from Figure 6.8, except for a new 400 mm lens P1. This forms a lens pair with the objective lens L1; the beam passes through P1 and is focussed at the focal point of L1, so that L1 collimates the beam before it reaches the target. This results in a larger spot diameter at the target position, and thus illumination of a larger area of the target. The imaging system continues to image the target as before. The only other changes from the original set-up are the addition of a beam-splitter between L1 and L2 (replacing a mirror in the previous set-up), and the replacement of the beam-splitter prior to the bottom VISAR with a mirror.

6.18. Technical drawings of this set-up are provided in Appendix C.

This set-up significantly improved the illumination, as shown in Figure 6.19. The downside of this approach was that the beam-splitter in the relay would reduce the intensity of the captured emission (both the VISAR signal and the self-emission) by half (although the second VISAR beam-splitter could be replaced with a mirror, which compensated for the signal loss on one of the VISARs). This was a necessary sacrifice though, as without this change the illumination was not sufficient for the VISAR to be used.

6.4.3 Later than expected shock breakout times

It was noted mid-experimental run that the shock breakout times being measured for a given intensity were significantly longer than those predicted by simulations. This meant that the Vulcan drive beam was frequently ending prior to shock breakout - contrary to the experimental plan, where the drive beam would be kept on until

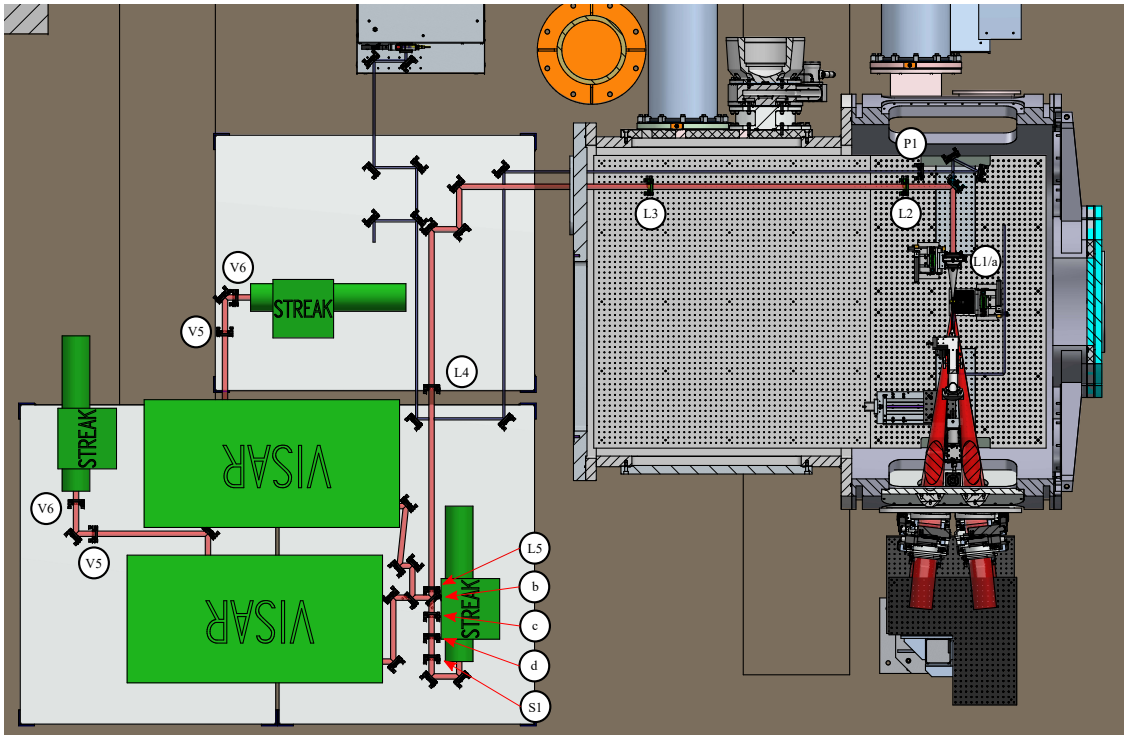


Figure 6.18: Scale CAD model of the updated set-up. The labels are as in previous figures. The main changes are the new beam path route for the probe laser (including the new lens P1), which leads to a beam-splitter between L1 and L2 (previously a mirror), and the removal of the beam-splitter before the second VISAR (replaced with a mirror).

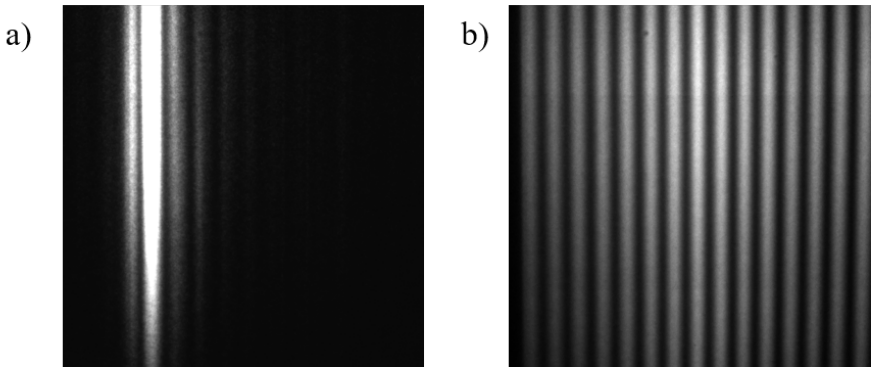


Figure 6.19: Images taken in focus mode on the VISAR streak camera from a piece of quartz a) using the original set-up and b) using the updated set-up. The quartz sample is stationary, and the camera is used in focus mode so that the signal at a single time across two spatial dimensions is recorded; uniform fringes are expected across the whole image. It is clear that in (a) only part of the sample is well illuminated, whereas the new set-up well illuminates the full imaged region. The signal strength was also better using the new set-up; (a) is only brighter due to different filtering.

shock breakout to ensure a sustained shock. Additional simulations were therefore performed mid-run to determine the expected impact of the drive laser pulse ending before shock breakout time. It was found that this in fact did not lead to substantial differences in shock behaviour, and thus was not a problem.

There was also concern that this could be due to a lower than expected laser intensity. However, the pulse time and pulse energy were directly measured for each shot using the diagnostics, and the laser spot diameter could easily be estimated from the pinhole image (the complete post-experiment analysis of this is given in Section 7.4). This suggested that the laser intensity on target was as expected. While the intensity would likely be reduced due to transmission losses between the calorimeters and the target, and due to not all the laser energy being contained within the focal spot, the magnitude of these losses are expected to be too low to account for the discrepancy in shock strength.

6.4.4 Signal strength issues

Throughout the experiment, both diagnostics suffered from a lack of signal strength. Some of this could be explained by the streak cameras. Three were used in this experiment; two high dynamic range Hamamatsu C7700 cameras were used for the VISARs, while a Hamamatsu C5680 was used for the streaked optical pyrometer. The older C5680 was expected to be less sensitive, and this was used for the SOP accordingly (as the VISAR was the more important diagnostic). This camera struggled for signal throughout. One of the VISARs also consistently recorded low signals compared to the other; it was later realised that this camera had an S1 streak tube as opposed to the S20 used in the other. The S1 tube (which was also present in the SOP streak) is around two orders of magnitude less sensitive to 532 nm light.

The signal strength for both VISARs was also heavily influence by the probe laser. This laser was relatively high power, and significant ND filters were needed to prevent it damaging the optics in the optical chain. However, the long optical relay likely resulted in some decrease in signal strength. There were also significant operational issues with the probe laser. It frequently did not seed, resulting in the pulse being

significantly weaker, and being emitted at a different time. This would occur on shots, where the low strength and change in time (relative to the Vulcan pulse) would prevent useful data from being collected. After attempting a number of fixes, the facility staff eventually replaced the seed laser within the unit with one from a different laser. This improved the situation and allowed the experiment to continue, but the timing and pulse intensity continued to vary drastically from shot to shot.

6.5 Conclusion

An overview of the experiment has been presented, along with an explanation of the diagnostics that were used to make the measurements. The preparation for the experiment was discussed, including how the optical set-up and experimental configuration was designed, and the simulations performed to assist with design of the targets. Finally, challenges that arose during the actual experiment were considered, along with the changes that were made to allow data to be collected. The analysis of this data and the results will be discussed in the next chapter.

7

Measuring the principal Hugoniot of TMPTA: analysis and interpretation

Contents

7.1 Introduction 157
7.2 Analysis of VISAR data 158
7.2.1 Identifying shock velocities from raw VISAR data . . . 158
7.2.2 Impedance matching calculation 160
7.2.3 Uncertainty quantification - Monte Carlo analysis 163
7.2.4 Achieved quartz shock states 167
7.2.5 Calculation of velocity from VISAR fringe curvature . . 168
7.3 Analysis of SOP data 171
7.3.1 Calculating grey-body temperature from SOP data . . . 171
7.3.2 Comparing SOP timings to VISAR timings 174
7.4 Analysis from additional diagnostics 175
7.4.1 Using the fiducial to estimate laser turn-on time 176
7.4.2 Using the pinhole camera to estimate spot size 177
7.4.3 Measuring laser pulse length using the fiducial and photodiode traces 180
7.4.4 Estimating pulse intensity and comparison with shock strength 181
7.5 Shot selection for final data set 183
7.5.1 Possible evidence of second shocks 183
7.5.2 Confidence ranking of shots 184
7.6 Results and discussion 185
7.7 Comparison with simulation 189
7.7.1 Shock propagation time 191
7.7.2 Shock stability 192
7.7.3 Shock merger 194

7.8 Investigation of the shock strength discrepancy 196
7.8.1 Analysis of non-data targets 196
7.8.2 Proposed explanation for the discrepancy 197
7.8.3 Simulation of a gap between layers 198
7.9 Suggested improvements 199
7.9.1 Target improvements 199
7.9.2 Equipment 200
7.9.3 Experimental layout 201
7.10 Conclusion 201

7.1 Introduction

This chapter outlines the steps taken to analyse the data from Chapter 6. The results are then discussed and compared with simulation and previous experimental results, and some unexpected features in the data investigated.

In Section 7.2, the analysis of the VISAR data is discussed, showing how the streak images are used to calculate the shock state achieved in the foam on each shot. Section 7.3 focuses on the SOP data, describing how this was used to calculate the shock temperature, and how SOP timing data was compared to that obtained from the VISAR. Section 7.4 then looks at the information that was obtained using the fiducial, pinhole camera, and photodiode traces of the laser pulse. The data analysis then concludes with Section 7.5, which broadly describes how decisions were made regarding the final data set.

The main results are then discussed in Section 7.6, where they are compared to both theoretical models and data from previous experiments on similar materials. Comparison with post-experiment simulations is then presented in Section 7.7. Section 7.8 goes on to describe how a certain feature of the data - a weaker than expected shock in the quartz - was investigated, and explores a potential explanation for this. Finally, Section 7.9 discusses suggested improvements that could be made were such an experiment to be performed again.

The structure used here has been chosen to aid explanation of the results, and roughly follows how it was actually performed. However, the analysis was in reality an iterative process; early stages were refined and revised later on, and early

results from the preliminary analysis meant that this was performed with some knowledge of the final data set. A basic confidence ranking had to be performed early on to obtain preliminary results and spot trends in the data, and this was then updated as understanding and interpretation of the data set improved. In recognition of this, the results from each stage of the described analysis are included in that section, although these are not discussed in detail until Section 7.6. The data set used in these results is the data set selected in Section 7.5 (highlighting the iterative nature of this analysis).

Throughout this chapter, experimental data is displayed to aid understanding of the analysis. For consistency, data from the same shot has been used where possible. Shot 38 is used for the general VISAR and SOP analysis, as this shot contained the clearest data. However, it did not have a fiducial present, which meant that the ablator/gold transit times could not be determined - the simulations therefore focus instead on an alternative shot, Shot 47. Data from other shots are also sometimes used where necessary to best show a particular feature or behaviour.

7.2 Analysis of VISAR data

7.2.1 Identifying shock velocities from raw VISAR data

In the VISAR streak images, the x-axis position is a spatial dimension, and the image contains two distinct halves - one corresponding to the quartz, and the other to the foam. The y-direction represents time, with later times at the bottom of the image. A labelled example can be seen in Figure 7.1. Three 'transitions' are identified in the image. Firstly, shock entry into the quartz is identified by a sudden change in signal intensity on the quartz half of the image. Secondly, shock breakout from the quartz rear-surface is identified by fringe extinction on this half. This also corresponds to the shock entering the foam (although as the foam is opaque to the probe laser, no change is seen in the foam signal at this point). Finally, shock breakout in the foam is identified by fringe extinction on the foam half of the image.

The two VISAR images for each shot were considered together. In each image, where possible, the time at which these three changes in behaviour occurred were

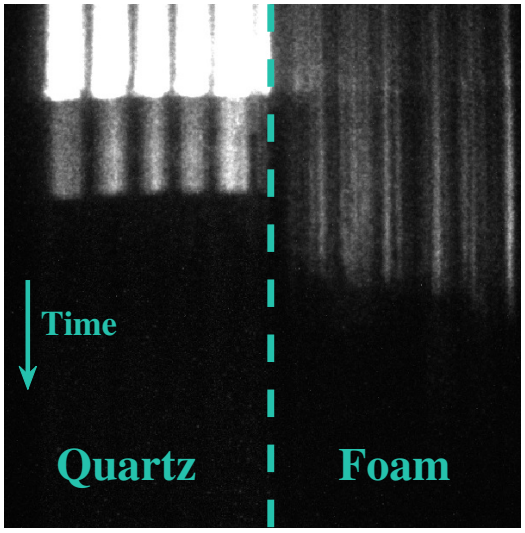


Figure 7.1: A raw VISAR streak image for shot 38, labelled to indicate the quartz and foam signals. Three clear changes can be observed; a change in fringe intensity in the quartz (corresponding to the shock entering the quartz from the gold), fringe extinction in the quartz (shock breakout from the quartz), and fringe extinction in the foam (shock breakout from the foam).

identified. These changes typically occurred over a range of pixels (rather than at a single clear point), and so a 'region of interest' was selected for each to cover the full range over which the transition could possibly have occurred. Example selections are displayed in Figure 7.2.

The streak time was then used to convert the pixels (in the y-direction) to time. The ROI's were thus converted from a pixel range into a time (the mean position of the window) with associated uncertainty (the width of the window). The finite slit width of the camera also introduced a minimum possible time resolution (which was used for the uncertainty if the ROI width was less than this value). The differences between these three timings were then used to calculate the shock transit time through the quartz and the foam.

For many shots, the timings could only be confidently determined from a single VISAR. In these cases, only the transit times from that image were used for the analysis. When both VISARs returned good data, these two independent measurements were used to reduce the uncertainty; a single transit time (with

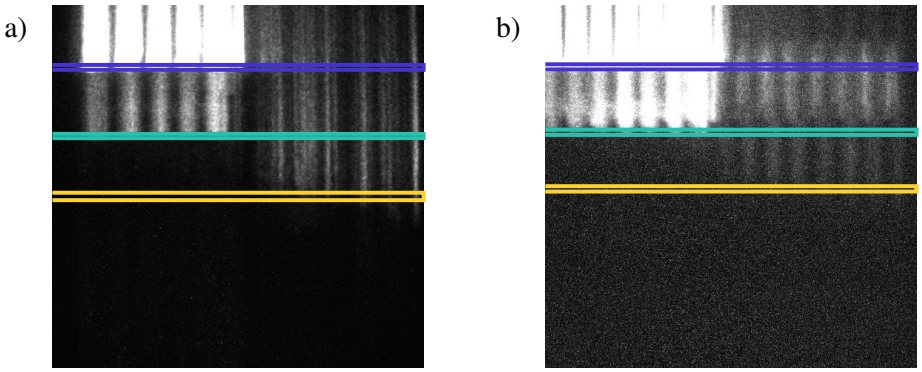


Figure 7.2: VISAR images from the two VISARs for shot 38. The selected regions where the shock transitions are believed to have occurred are shown. The blue region corresponds to shock entry into the quartz, while the teal region corresponds to shock breakout from the quartz and the yellow corresponds to shock breakout from the foam. The edge of image (a) appears to show some curvature; this is expected based on the simulations due to the size of the Vulcan laser spot. The region was chosen to capture the transition at the centre, where the shock is expected to be more planar.

smaller uncertainty) was produced, corresponding to the range of values compatible with both diagnostics.

Finally, the thickness of the quartz and foam layers (with estimated uncertainties) were then used to calculate the average shock velocity in each material. The shock velocity calculated for a single shot using the independent time measurements from each VISAR, along with that calculated from the combined time measurements, are shown in Figure 7.3.

7.2.2 Impedance matching calculation

The impedance matching calculation was then performed to calculate the foam shock state from the two shock velocity measurements. The principles behind this calculation is described in Sections 2.8.4 and 2.8.5, where in this case the quartz is the reference material (for which the Hugoniot and release isentrope are known) and the foam is the material of study.

An initial faster but less accurate impedance matching calculation was performed to assess the data (and to provide a sanity check on the more accurate, but more complicated, subsequent calculation). This used a common assumption that the

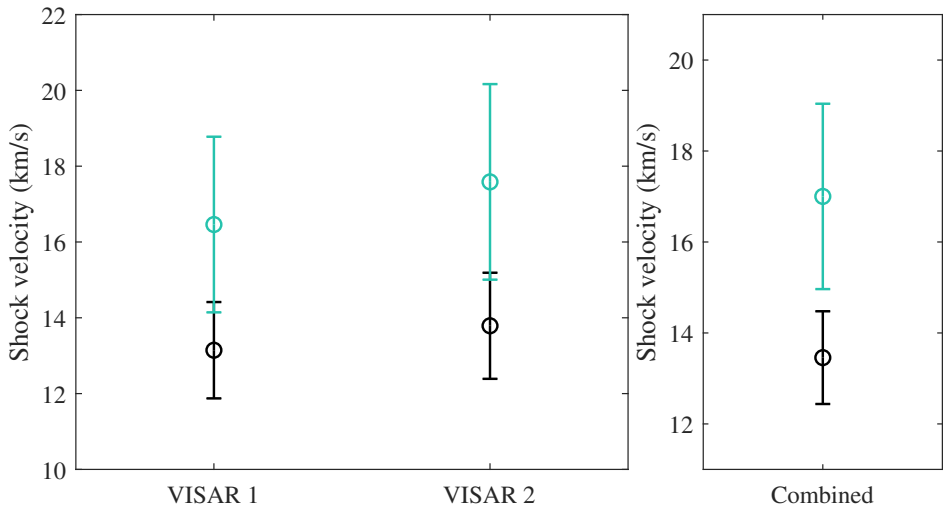


Figure 7.3: Calculated shock velocities in the quartz (black) and the foam (teal) for shot 38, using the shock timings obtained from the VISAR images in Figure 7.2. In this shot, all three timings can be identified from both VISARs. The shock velocities corresponding to the timing data from each of the two VISARs independently is displayed, along with the shock velocity calculated from the 'combined' timings. The combined measurement uses both sets of VISAR measurements to reduce the error, giving a 'combined' set of possible timings that is compatible with the uncertainty ranges measured by each VISAR.

quartz release isentrope can be approximated by reflecting the quartz Hugoniot around the quartz shock state. This is a reasonably accurate approximation at low pressures, but introduces errors as the pressure is increased [49]. In this calculation, the quartz Hugoniot described in [165] was used, with a functional form of

$$U_s = a + bu_p - cu_p \exp\{-du_p\} \,, \tag{7.1}$$

with coefficients of $a = 6.26 \pm 0.35$ km s^{-1}, $b = 1.20 \mp 0.02$, $c = 2.56 \mp 0.15$, and $d = 0.37 \pm 0.02$ (km s^{-1})$^{-1}$. Then, as described in Section 2.8.5, the measured shock velocities were used to produce the quartz and foam Rayleigh lines. The intercept of the quartz Rayleigh line with quartz Hugoniot was used to determine the quartz shock state. The quartz Hugoniot was then reflected about this state, and the intercept of this approximate release isentrope with the foam Rayleigh line allowed the foam shock state to be calculated. This calculation was performed for each shot, and the preliminary results indicated that there was indeed a trend in the data and thus that the more complicated analysis should be conducted.

Table 7.1: Coefficients for the quartz Hugoniot in equation 7.2, reproduced from [166].

a_0 (km/s)	a_1	a_2 $(\mathrm{km/s})^{-1}$	a_3 $(\mathrm{km/s})^{-2}$
1.754	1.862	-3.364×10^{-2}	5.666×10^{-4}

The more accurate calculation used a full and complete calculation of the release isentrope. For this version of the impedance matching calculation a slightly different functional form for the Hugoniot was used [166] - a cubic fit to the same data as used previously [165]. The new Hugoniot was described by

$$U_s = \sum_{n=0}^{3} a_n u_p^n \,, \tag{7.2}$$

with coefficients in Table 7.1. This newer form was used to facilitate the use of Monte Carlo error analysis (described in Section 7.2.3). The release isentrope was then accurately calculated according to the relevant quartz shock state. This was done largely according to the method outlined by Knudson and Desjarlais [166], but with a few small differences.

In their paper, Knudson and Desjarlais model the isentrope using a new 'linear-reference' Mie-Grüneisen model. The key departure of this model from other approaches is the use of a variable Grüneisen parameter. They compare the new model to the previous standard - a conventional Mie-Grüneisen model with a fixed Grüneisen parameter of $\Gamma = 0.64$ - and found that the new model led to reduced errors. However, the range of pressures over which this model is valid does not extend across those achieved in this experiment, and the model returned implausible results for some of these values. As a result, the conventional Mie-Grüneisen model, with a fixed Grüneisen parameter of $\Gamma = 0.64$, was used. This value of the Grüneisen parameter was calculated from a range of EOS models [167], and is in good agreement with that derived [167] from experiments [168, 169]. Other than this change, the method for calculating the release isentrope for a given quartz shock state was as Knudson and Desjarlais described [166].

The impedance matching calculation was therefore conducted as in the previous method, with this more accurate model for the quartz isentrope used, and the foam shock $(u_p^{\mathrm{foam}}, P^{\mathrm{foam}})$ state was determined from the intercept of this curve with

the foam Rayleigh line. A graphical representation of this for the experimental data from Shot 38 is shown in Figure 7.4. This calculation was performed for all shots. A comparison with the previous method, using the reflected Hugoniot as the isentrope, found the difference to be essentially negligible; this suggests the isentrope is comparable to the reflected Hugoniot in this regime, and provided a sanity check that the release isentrope had been calculated correctly.

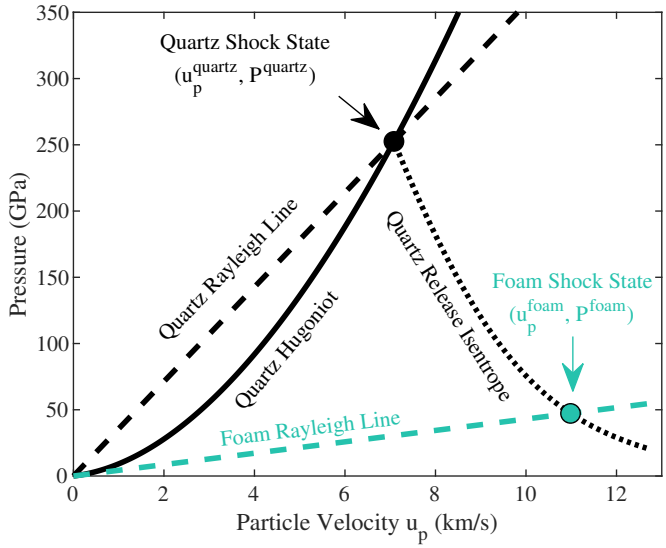


Figure 7.4: The impedance matching calculation for the shock velocities obtained from shot 38. The intercept of the quartz Rayleigh line and the quartz Hugoniot defines the quartz shock state. The release isentrope for this shock state is then calculated, and the intercept of this curve with the foam Rayleigh line describes the foam shock state.

The resulting foam shock states for the final data set are displayed in Figure 7.5, where they are compared to theoretical models. These results are discussed in Section 7.6. The error bars are calculated as explained in Section 7.2.3.

7.2.3 Uncertainty quantification - Monte Carlo analysis

Uncertainty quantification for the shock variables obtained using the impedance matching calculation was performed using a Monte Carlo approach [170]. This technique allowed the experimental uncertainty to be propagated through the complex IM calculation, while also including uncertainty contributions from the quartz Hugoniot fit and foam initial density. This section is accompanied by a flow chart in Figure 7.6, which shows the different steps of this calculation for shot 38

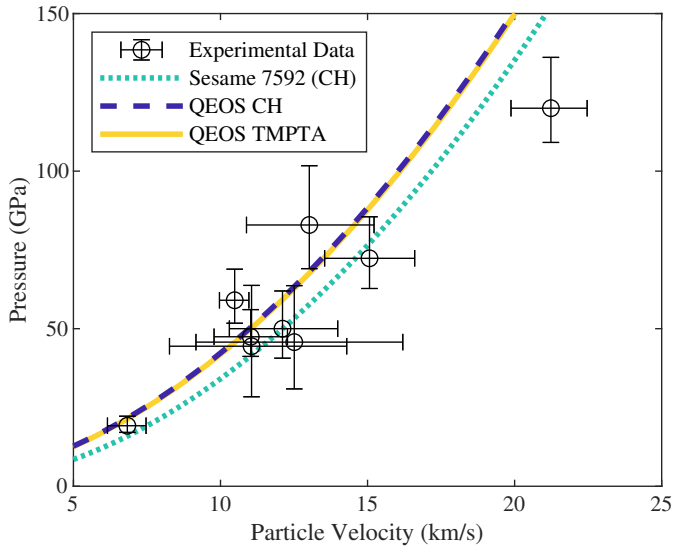


Figure 7.5: The foam Hugoniot data calculated from the impedance matching calculation, compared to Hugoniots generated from theoretical equation of state models.

with example distributions and samples. Uncertainty quantification in this way has previously been performed on other similar experiments [166, 170, 171, 172].

The calculation begins with the measured shock velocities for the quartz and foam for a particular shot. These are of form $U_{\text{shock}} \pm \epsilon$; a quoted value U_{shock}, with an experimentally measured error ϵ that is symmetric around this value. A normal distribution is then produced for each of the two shock velocities, with a mean of U_{shock} and a standard deviation of ϵ.

These distributions were then used to perform 10,000 Monte Carlo iterations of the impedance matching calculation for each shot. For each iteration, a sample is randomly generated for each shock velocity and the foam density according to the relevant distribution. Samples are also generated for each of the four Hugoniot fit parameters, according to the covariance matrix[1] provided in [166],

$$C = \begin{bmatrix} 2.097 \times 10^{-2} & -6.159 \times 10^{-3} & 5.566 \times 10^{-4} & -1.572 \times 10^{-5} \\ -6.159 \times 10^{-3} & 1.877 \times 10^{-3} & -1.742 \times 10^{-4} & 5.017 \times 10^{-6} \\ 5.566 \times 10^{-4} & -1.742 \times 10^{-4} & 1.650 \times 10^{-5} & -4.834 \times 10^{-7} \\ -1.572 \times 10^{-5} & 5.017 \times 10^{-6} & -4.834 \times 10^{-7} & 1.438 \times 10^{-8} \end{bmatrix}.$$

[1]The printed matrix is that provided by [166], but actually attempting to use this to generate the parameters returned an error due to the matrix not being positive definite. To resolve this, a small correction was added, $C + 1 \times 10^{-10} \cdot I$ was required, where I is the identity matrix [173].

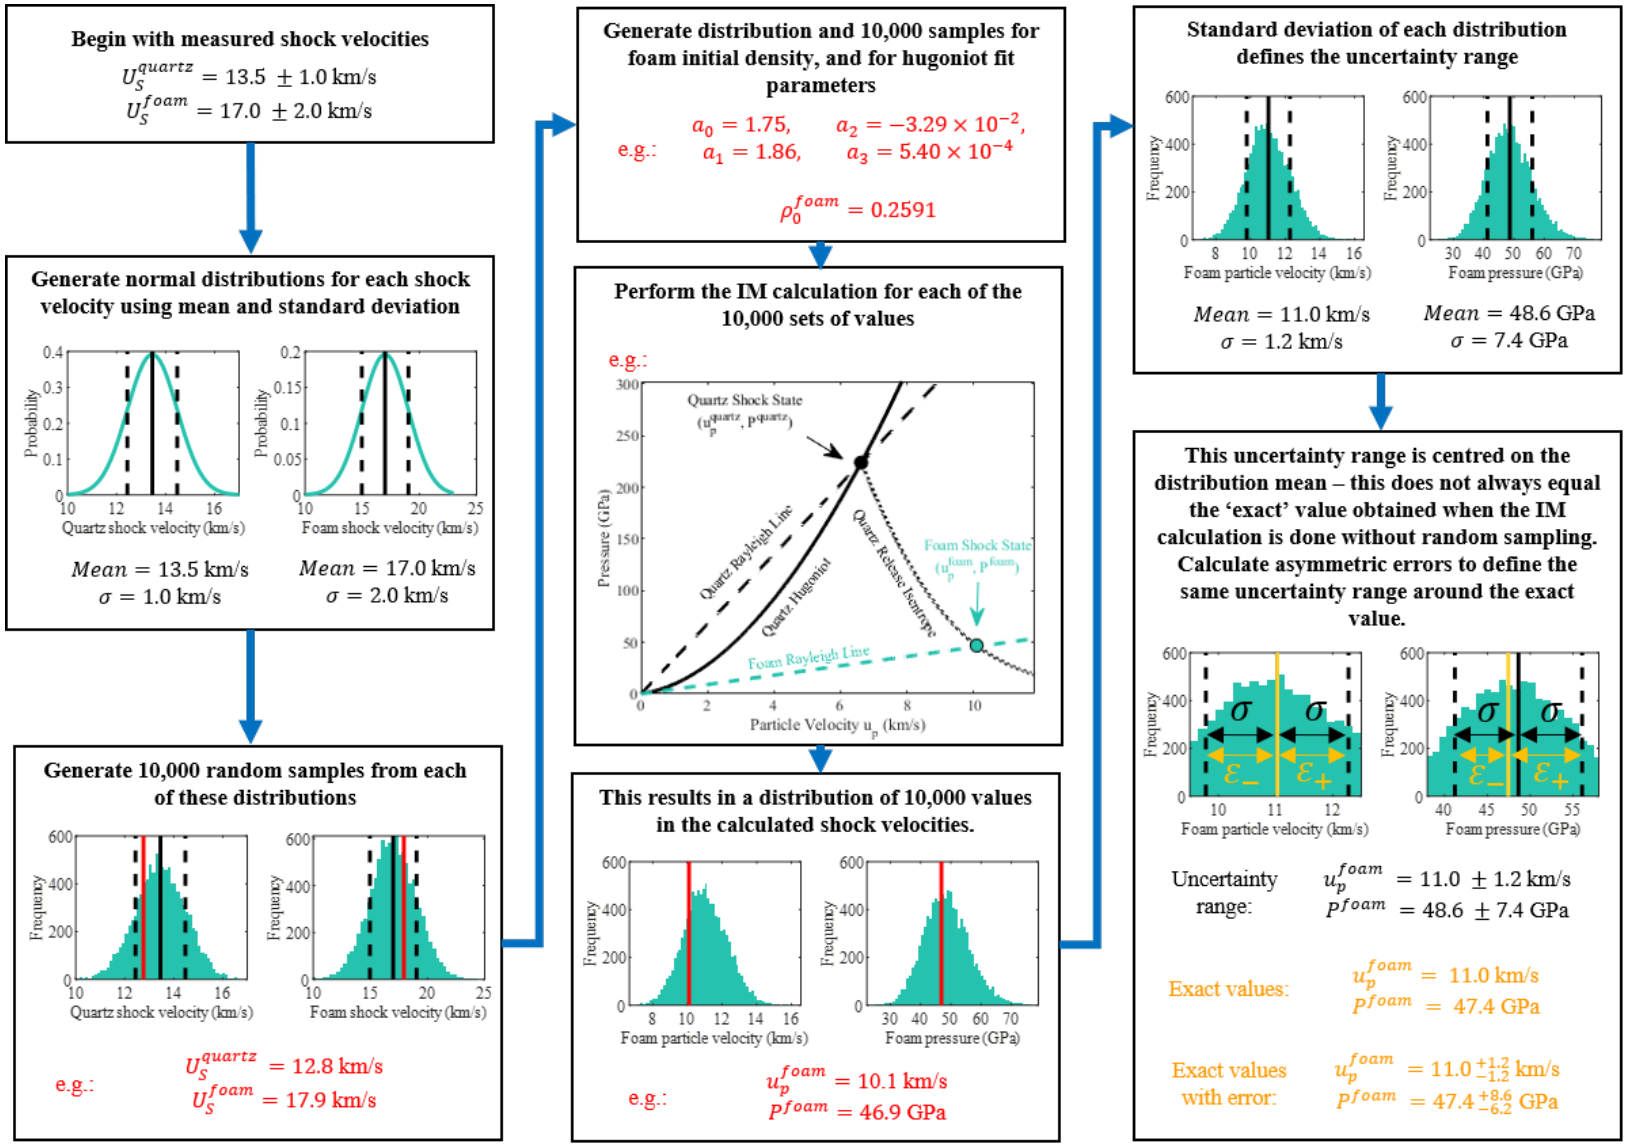


Figure 7.6: Flow chart showing the MC methodology for quantifying the uncertainty in the calculated variables. The MC analysis is shown for shot 38. The values/lines in red show the IM calculation performed for an example random sample from the set of the 10,000 random samples.

The impedance matching calculation was then performed, as described in Section 7.2.2, using these randomly sampled values. The 10,000 iterations produced 10,000 different estimates of the quartz and foam shock states, providing distributions for each of the relevant shock variables. The standard deviations σ of these distributions were then calculated, and used to define the uncertainty range of these values[2].

The mean of the MC distribution of a particular value does not always equal the 'exact' value obtained if the IM calculation is performed without considering uncertainty (i.e. if no MC sampling is performed, and thus for all shock velocities $U_{\text{shock}} \pm \epsilon$ the values of U_{shock} are used, along with the exact Hugoniot fit parameters quoted in Table 7.1). This was most noticeable for the shock pressure, where there is a small skew in the distribution[3] (the particle velocity distributions are effectively normal). This effect arises from a non-linearity in the calculation. Consider a shock velocity U_{shock} and associated pressure P. A small positive increase in U_{shock} results in a larger change in P than a small decrease in U_{shock}, leading to a distribution mean that is skewed towards higher values.

The most obvious way to present the results of this calculation would be to quote each shock variable as the MC mean plus/minus the standard deviation. This could lead to confusion; for instance, the quartz shock state would then be quoted as $\bar{u_p}^{\text{quartz}} \pm \sigma_{u_p^{\text{quartz}}}$ and $\bar{P}^{\text{quartz}} \pm \sigma_{P^{\text{quartz}}}$ (where a variable with a bar is used to indicate the MC distribution mean), but the state $(\bar{u_p}^{\text{quartz}}, \bar{P}^{\text{quartz}})$ would not actually be a valid state on the quartz Hugoniot described by Equation 7.2. To avoid this, the shock variables are instead quoted as the 'exact' values - those obtained if the IM calculation is done without random sampling of any of the variables. These values are quoted with asymmetric error to describe the same uncertainty range as the standard deviation (which was symmetric about the mean). This is shown in the last

[2]It should be noted that there is some dependence between the quartz and foam shock velocities; if the quartz breakout time is perceived to occur earlier, then the quartz shock velocity will be higher while the foam shock velocity will be lower. Each of these velocities can also vary independently of the other, however, through the time of shock entry into the quartz and the time of shock breakout from the foam respectively. They are therefore treated as independent for this analysis.

[3]Skew in the MC distributions of such an experiment was also previously noted in [171].

box of Figure 7.6, where one standard deviation σ from the mean (solid black line) is described by asymmetric errors ϵ_+ and ϵ_- around the 'exact' quoted value (orange)[4].

7.2.4 Achieved quartz shock states

The achieved shock states in the quartz are displayed in Figure 7.7. This plot includes the eight shots from the main dataset, plus an additional two shots[5]. As the quartz particle velocity and pressure were not measured directly, but were inferred using the Hugoniot based on the measured shock velocity, these data points all lie perfectly on the Hugoniot curve. It can be seen that the pressures achieved in these shots are significantly lower than were expected for the pre-experiment simulations, presented in Figure 6.12 (this is also the case for the foam pressures). While it is not surprising that shock pressure is higher in 1D simulations, the size of this discrepancy is unexpected. This will be explored further in Section 7.8.

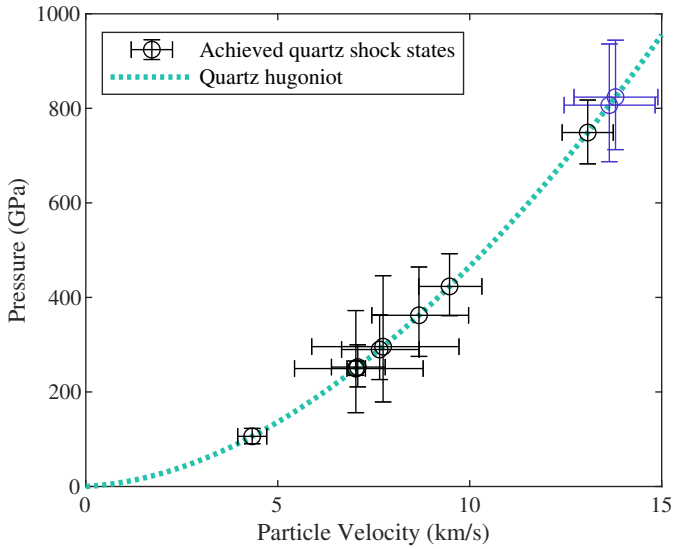


Figure 7.7: The quartz shock states achieved in the experimental shots, for the main data set (black) and two additional shots where the quartz signal could be determined with confidence (blue). The shock pressures and particle velocities are calculated based on the quartz Hugoniot, and thus all lie exactly on this curve.

[4]It is also worth noting that the difference between the mean and the exact value is relatively small, and this uncertainty range actually represents 34% (i.e. 1σ) of the generated MC samples in either direction from the 'exact' value to within one or two percent).

[5]These two shots were not included in the main dataset because of the lack of foam signal, but the quartz velocity could be calculated with confidence. In one, the target did not contain a foam layer. The other did have a foam layer, but the signal was too weak for foam breakout to be detected.

This weaker than expected quartz shock strength potentially explains two of the behaviours noted during the experiment. Firstly, the weaker shock will have a lower shock velocity, which explains the longer than predicted shock transit times. Secondly, many of the shots in Figure 7.7 are at pressures around 200-300 GPa, which (particularly when the error range is considered) are comparable to the threshold pressure at which the quartz is expected to become reflective. This provides a potential explanation for the lack of fringe curvature exhibited on the VISAR in most shots; it could well be that the shock pressure simply was not high enough for the shock front to exhibit sufficiently high reflectivity.

A small number of shots in this experiment did demonstrate some degree of fringe curvature, although the weak signal strength meant this could not be analysed to a sufficient level for velocities to be determined. It is useful to compare these shots with Figure 6.12 to check the pressures for which this occurred. Figure 7.7 displays 10 shots in total; of these, 4 displayed some degree of curvature, at pressures of 290 GPa, 360 GPa, 810 GPa and 820 GPa. There were only 2 shots with pressures above 290 GPa that did not display any curvature (at 450 GPa and 720 GPa). No shots displayed curvature for pressures below 290 GPa. This threshold behaviour suggests that it may indeed have been necessary to achieve a pressure above 290 GPa in order to obtain sufficient reflectivity for the VISAR, and thus that the lack of fringe curvature observed on the other shots was likely just a result of the shock pressure being too low. This correlation between the pressure and the shots displaying fringe curvature gives confidence that the shock velocity measurements and calculation of shock pressure are likely accurate.

7.2.5 Calculation of velocity from VISAR fringe curvature

The intention for this experiment was that the VISAR signal would display fringe motion in the quartz, and the quartz shock velocity as a function of time could thus be obtained following the analysis procedure described in Appendix B.3. Unfortunately this was not possible; very few shots displayed any fringe motion, and those that did had low signal strength or were too noisy for meaningful analysis to

be conducted. However, attempts were made to produce such an analysis for shots where some curvature was observed. This was performed based on code written by Dan Eakins, which was adapted for use in this experiment[6].

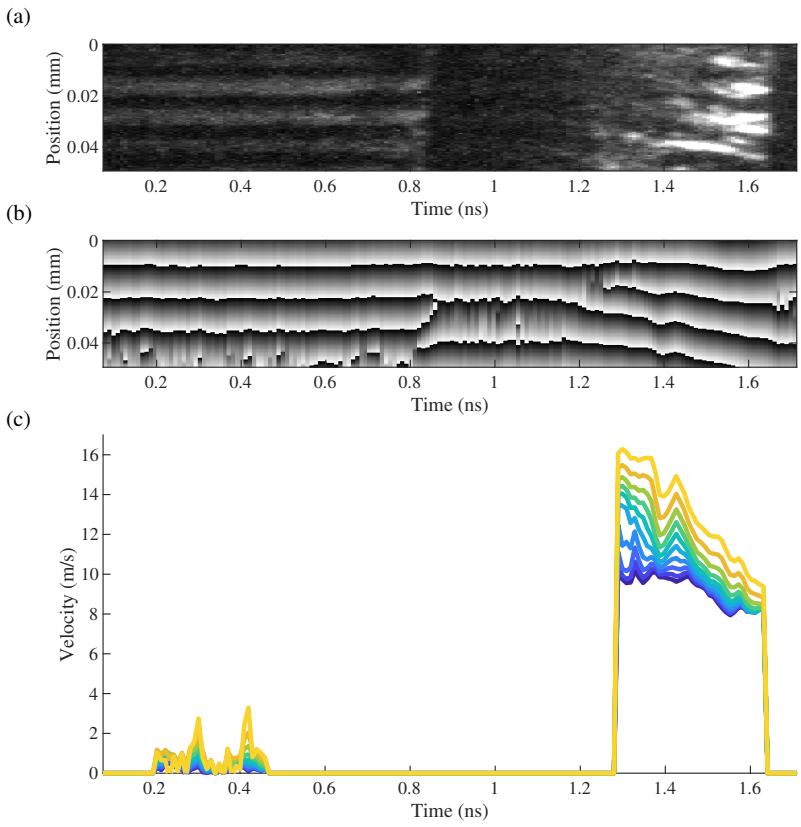


Figure 7.8: Analysis of the curved VISAR fringes in shot 39. In a), the fringes in the quartz in the VISAR image (cropped and rotated) are shown. In b), the phase has been extracted from this image. The line-outs in c) then show the velocity extracted from these fringes at different vertical positions in the image. The velocity calculation only included the time period where the fringes are strongest, as including earlier times led to erroneous results.

Figure 7.8 shows this analysis for data from shot 39; the only shot in the final data set to show significant curvature. Figure 7.8 (a) shows the raw VISAR data up until the time of quartz breakout, where it can be seen that the fringes have clearly curved and shifted from their original position. Figure 7.8 (b) shows this data after

[6]In particular, the code was adapted so that it could be used in cases where there were fringe discontinuities which could increase the velocity by an integer number of VPF. It was also adapted so that it could be used on data where there were only small regions where the fringe signal strength was sufficient for analysis.

filtering has been applied. Here the fringe movement can be seen more clearly; but it is also clear that the low signal leads to discontinuities. In particular, there are regions in the original data where the signal strength is low, and the filtered image is obviously also very poor in these positions. The signal between roughly 1.3 ns and 1.6 ns is stronger and produces clearer fringes; this range has been selected, and the analysis applied to allow velocity to be determined from these fringes (relative to the original 'zero velocity' fringes before the shock enters the quartz.

Figure 7.8 (c) then shows velocity traces for different positions within the image. As only one VISAR returned a curved signal, it was not possible to unambiguously determine the velocity. As such, the average velocity measurement was used to determine the appropriate integer multiples of the VPF to add to the velocity signal (for this shot, the VPF was 8.85 $\mathrm{m\,s^{-1}}$).

The measured average quartz velocity for this shot was $15.8 \pm 1.8\ \mathrm{m\,s^{-1}}$, and it can be seen that the velocity measured on the VISAR is roughly compatible with this. The trend in velocity seems to vary significantly along the target; the darker traces (corresponding to higher positions in the image) suggest a shock which is reasonably steady, whereas the lighter colours suggest significant levels of shock decay (this can also be seen by the different levels of curvature in the fringes in the filtered image). It is unlikely that such a discrepancy in shock behaviour was actually being recorded in the target, particularly as shock breakout (which occurs on the right of the image) is seen to occur at the same time at all positions. This discrepancy is likely due to the poor signal strength - the data is clearly very noisy, and the fringes are much less prominent than would be desired. This analysis therefore shows that the VISAR could in principle be used for such measurements, and the agreement with the average velocity lends a small amount of confidence to this approach; but it is also clear that even for one of the best examples of curved fringes achieved in the experiment the signal was too poor to return useful results, and it is thus difficult to make any conclusions about the shock velocity or stability from this data.

7.3 Analysis of SOP data

7.3.1 Calculating grey-body temperature from SOP data

As for the VISAR streak images, the x-dimension in the SOP streak image was spatial, while the y-dimension was temporal. The streak time for the camera on each shot was used to convert pixels in the y-direction into time values. An example SOP streak camera image (for the same shot considered in the previous sections) is displayed in Figure 7.9.

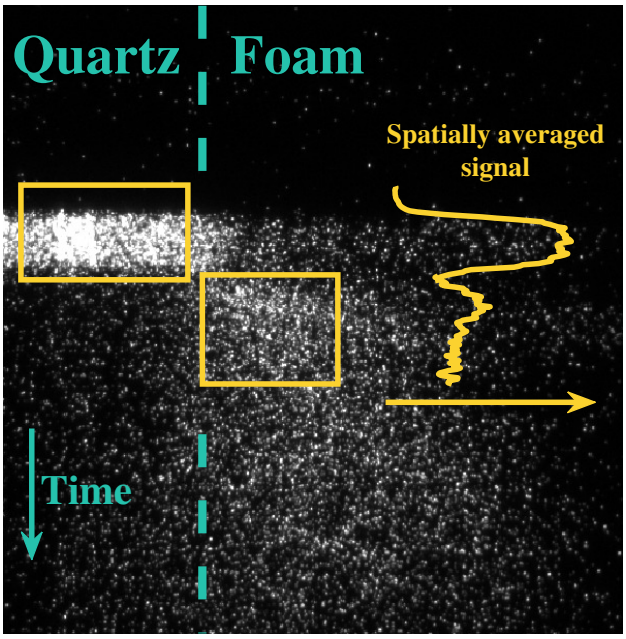


Figure 7.9: SOP streak data for shot 38. As with the VISAR data, the quartz is positioned on the left of the image, with the foam on the right. Time increased in the negative y-direction. The yellow rectangles represent the chosen regions covering the SOP emission for both materials, while the trace shows the spatially averaged signal across these regions.

In each image, a region was selected covering the emission from the two materials. This region was used for spatial averaging of the signal, to obtain an average intensity vs time for each material; this was chosen to cover the section of the image where the emission was uniform (i.e. to avoid including weaker signal in the averaging from the edge of the target), while being as wide as possible to improve the averaging. The maximum in the averaged signal for the two materials were then identified,

and used in the temperature calculation. The selected regions for each material in Figure 7.9 are indicated by the orange rectangles. As no temporal averaging occurs and only the maximum intensity is used, the temporal range of these selections is not significant; the only concern is ensuring that both peaks are captured, and that the foam selection does not capture any quartz signal (as this would be higher intensity than the foam signal).

In order to calculate the grey-body temperature using Equation 6.20, the reflectivity of the two materials was required. This was not measured in the experiment, and thus need to be estimated from other sources. The quartz is a well characterised reference, and thus relationships between shock velocity and reflectivity have been previously published. The equation

$$R_{532} = 4.614 \times 10^{-3} + \frac{(0.3073 - 4.614 \times 10^{-3}) \cdot U_s^{9.730}}{U_s^{9.730} + 16.185^{9.730}} \,, \tag{7.3}$$

published in the supplementary material for [162], was used to estimate the quartz reflectivity using the shock velocity values determined from the VISAR.

Less literature data exists for the foam. The only relevant data available was for polystyrene [174], but this was for solid density plastic rather than a foam. To improve this, collaborators at the University of Rochester performed density functional theory simulations to calculate the foam reflectivity for the conditions achieved in two separate shots. These were performed using the same method used for the polystyrene data (which they were also responsible for) [174, 175]. They found that for a density of 1.0906 $\mathrm{g\,cm^{-3}}$ and a temperature of 4.8454 keV, the reflectivity was 0.227, and for a density of 1.0568 $\mathrm{g\,cm^{-3}}$ and a temperature of 1.9743 keV, the reflectivity was 0.174.

There were a few options on how reflectivity should be estimated for the other data points. The most basic method is to assume a linear relationship between these two points. However, it was clear from the polystyrene data that reflectivity saturated at higher pressures; assuming a similar relationship for the TMPTA foam would mean a linear relationship would therefore over-estimate reflectivity at these pressures. Instead, the polystyrene fit was scaled using the two new

data points. First, the polystyrene data was extrapolated at low pressures to ensure it covered the full intensity range. It was then scaled and shifted so that it passed through both of the new simulation data points. This gave a linear relationship over most of the pressure range of interest, but ensured that the reflectivity saturated at high pressures. This reflectivity curve (with the two points indicated) can be seen in Figure 7.10.

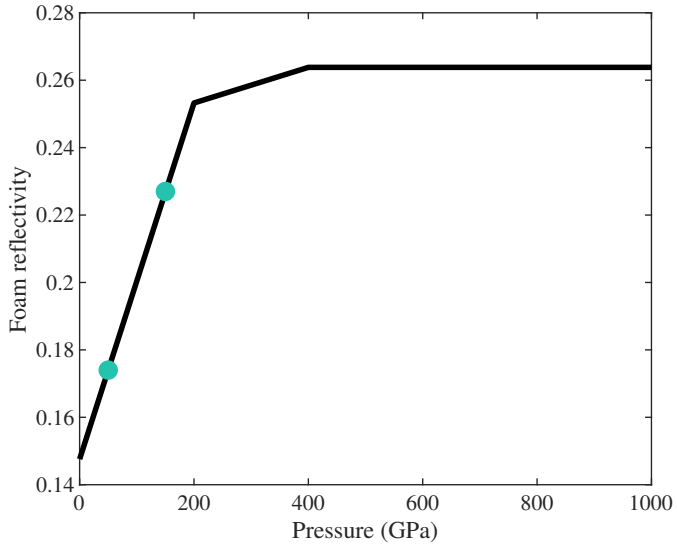


Figure 7.10: The foam reflectivity model used in this work. The polystyrene pressure reflectivity model from [174] has been scaled so that it passes through the two teal points; these were the results of two DFT simulations for the TMPTA material, performed for this experiment by collaborators at the University of Rochester. It can be seen that for the pressures of interest to this experiment, this is effectively just a linear relationship between the two points. However, this fit ensures that for any MC ransom samples at significantly higher pressures (statistically unlikely), the reflectivity will saturate rather than increasing indefinitely.

As discussed in Section 6.2.3, the SOP set-up was not absolutely calibrated and thus the calibration constant A in Equation 6.20 could not be determined in advance. The quartz was therefore used as an on-shot temperature reference. Equation 6.21 was used to estimate the expected quartz temperature, based on the the VISAR-determined shock velocity. This temperature and the maximum quartz SOP intensity from the SOP data were substituted into Equation 6.20 and used to determine A. This same calibration constant was then used for the foam data, and the maximum foam intensity was therefore used to calculate the foam shock temperature.

This approach has large uncertainties associated with it. The use of on-shot calibration meant that the quartz temperature had to be estimated, based on the measured quartz shock velocity and fits to previous data; this introduced error that propagated through to the foam temperature. The models to estimate the quartz and foam uncertainty also influence the final result. To capture all these contributors in the final uncertainty quantification, the same Monte Carlo analysis described in 7.2.3 was applied. The same quartz shock velocity samples used for the impedance matching were also used to perform 10,000 iterations of the temperature calculation. As uncertainties for the various fits used were not quantified, estimates had to be made. The quartz temperature fit was assumed to have an uncertainty of 8%, as this is the standard uncertainty for SOP temperature measurements at OMEGA [162] (where the data for the fit was obtained). The quartz reflectivity fit was assumed to have a 10% error, based on the size of the error bars in the published data [162]. The foam reflectivity was given a large error of 0.1, to account for the uncertainty regarding this material (and lack of published experimental data). The streak counts for the quartz and foam measurements were assumed to be based on a Poisson distribution, with the uncertainty appropriate for that distribution used (the square root of the number of counts). The uncertainty in all these models were included in the MC analysis. As described in that section, the quoted temperature for each shot is the 'exact' temperature (that obtained without sampling), while the error bars represent the standard deviation of the temperature distributions obtained using the MC methodology.

The results of the temperature measurements, compared to the theoretical models, are displayed in Figure 7.11. These results are discussed in Section 7.6.

7.3.2 Comparing SOP timings to VISAR timings

It was not possible to identify the breakout timings from each layer directly from the SOP data, as there was no clear change in signal corresponding to shock breakout from the quartz. However, the signal was expected to begin when the shock entered the quartz, and the foam signal was expected to peak at shock breakout from the

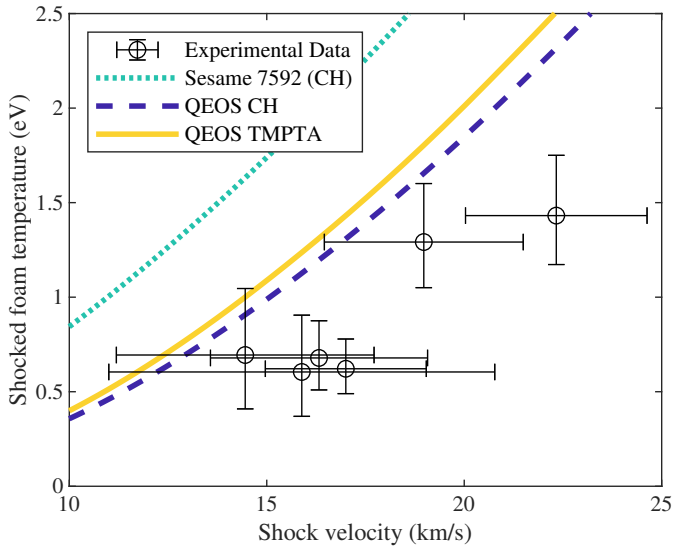


Figure 7.11: The calculated grey-body temperatures of the shocked foam, as a function of shock velocity, compared to the theoretical equation of state models.

foam. Based on this, was possible to determine a combined quartz/foam transit time that could be compared to the VISAR data.

This combined transit time for both the SOP and the VISAR data for all the valid shots is shown in Figure 7.12. The error bars on the SOP data are large; this is due to the large slit size used for the SOP (used to maximise signal strength), and the large sweep time. However, it can be seen that the agreement between the two diagnostics is good. This comparison serves two purposes. Firstly, it gives more confidence in the timing data for these two shots, as it is confirmed by two independent diagnostics. Secondly, it provides more confidence to the assumption that (despite the relatively low signal strength) the foam peak in the SOP data does in fact correspond to the foam breakout.

7.4 Analysis from additional diagnostics

Data was also obtained from the x-ray pinhole camera, fiducial, and photodiode traces of the laser pulse. Firstly, the fiducial allowed the laser switch-on time to be determined, which allowed the shock transit time through the combined ablator and gold layers to be estimated. This is discussed in Section 7.4.1. The pinhole camera allowed the laser spot size to be estimated, which allowed the assumption

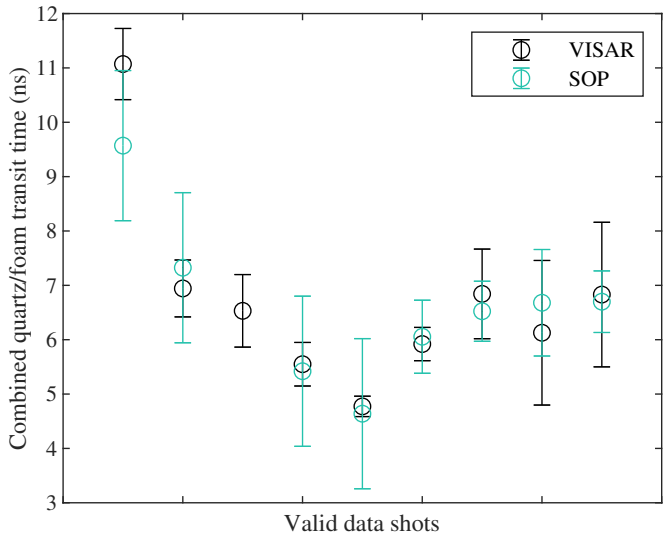


Figure 7.12: The combined quartz/foam transit time measured using the VISAR (black) and SOP (teal).

of a 400 µm spot to be tested - this is described in Section 7.4.2. Section 7.4.3 describes how the laser pulse length was also measured (using both the fiducial and photodiode traces). These two pieces of information allowed an estimate of laser intensity to be made; this intensity is then presented in Section 7.4.4, where the trend between laser intensity and shock velocity is also considered.

7.4.1 Using the fiducial to estimate laser turn-on time

As discussed in Section 6.2.5, a fiducial on one of the VISAR streak cameras displayed the Vulcan pulse, and two-minute shots were performed to allow the time delay between the fiducial and the main pulse to be calculated. Three such shots were performed, and the data from one of these is displayed in Figure 7.13. The time difference between the peak fiducial signal and the peak of the (spatially-averaged) Vulcan signal was calculated. This delay was averaged across the three shots to calculate the 'fiducial delay'.

The fiducial was not present for many of the early shots. For those where it was present the foot of the fiducial signal was identified, and the previously identified fiducial delay time was subtracted to find the time at which the Vulcan pulse was applied. This had an uncertainty associated with how accurately the fiducial delay

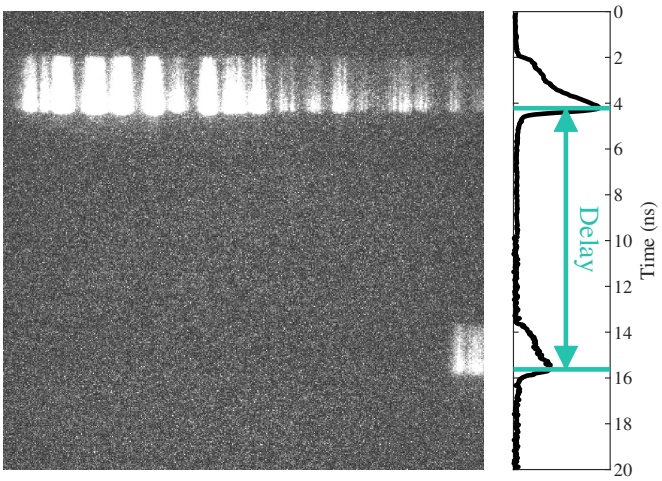


Figure 7.13: VISAR streak for a two-minute shot. The main Vulcan signal can be seen, along with the fiducial (on the bottom right of the image). On the right, a line-out is taken showing the average intensity of these signals (the fiducial is averaged only across the pixels where it appears, while the laser signal is averaged across the whole width of the image). The maximum intensity of each of these signals is identified (teal lines), and the delay time between them determined.

could be calculated from the two minute shots, and how accurately the foot of the fiducial could be identified. If the quartz shock entry signal could be identified in the same VISAR streak as the fiducial, it was therefore possible to calculate the ablator/gold transit time - the time between when the laser was first applied and the quartz entry signal. This is shown in Figure 7.14.

7.4.2 Using the pinhole camera to estimate spot size

The x-ray pinhole camera was used to estimate the laser spot size. This diagnostic was not particularly accurate for such measurements, and was therefore not intended to be used in the intensity calculation; rather, the spot size estimated from this diagnostic could be used to indicate whether the 400 µm spot that the phase plates were designed to create was in fact being achieved, or if the spot was substantially different from this. It could also be used to identify poor alignment between the six beams. An example pinhole camera image is displayed in Figure 7.15.

The pinhole image showed a bright disk corresponding to the laser spot, with

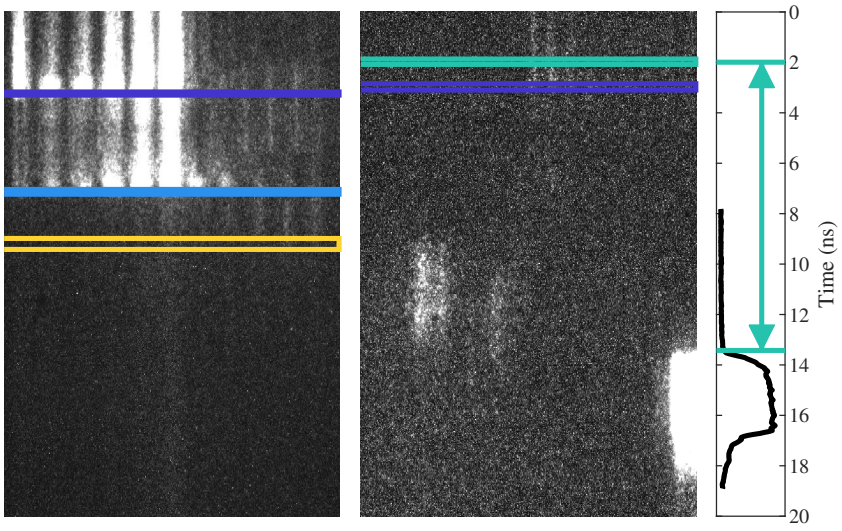


Figure 7.14: The two VISAR images for shot 47. The fiducial is present on the second VISAR image. The trace on the right shows the intensity of the fiducial; its start time is identified, and the delay time is used to estimate when the laser is first applied. This is plotted on the streak image (with uncertainty) in teal. The signal on this VISAR image is weak, and only the quartz entry signal (dark blue) is identified; this is sufficient to determine the ablator/gold transit time. The other signals (quartz breakout in light blue, foam breakout in orange) are identified from the other VISAR, where the signal strength is better. Shot 47 was used for this figure rather than shot 38, as shot 38 did not include a working fiducial.

a bright halo corresponding to emission from the copper shielding. The fact that the laser seems to be a single spot without noticeable structure suggests good alignment between the six Vulcan beams. As the pinhole camera was positioned above the target the image is distorted in the y-direction, but the camera was normal to the target horizontally and thus the scaling in the x-direction is expected to be accurate. A narrow (in y) region was initially selected, positioned so that the aperture diameter was at its maximum. The peak in emission either side of the aperture was identified. This can be seen on the trace below the image in Figure 7.15. The distance between these peaks corresponds to the diameter of the aperture in the copper shielding, which was known to be ~ 1.2 mm. This allowed a horizontal length scale in the image to be determined.

The trace above the image in Figure 7.15 shows the signal in another narrow horizontal selection of the image, located at the peak horizontal diameter of the

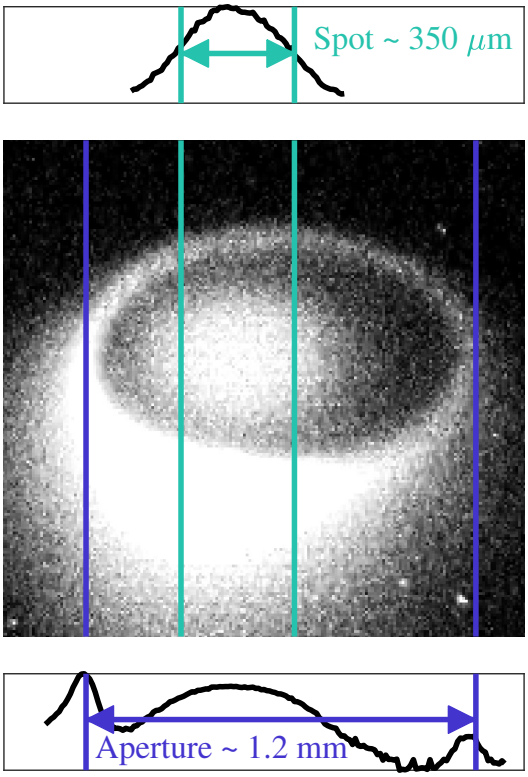


Figure 7.15: X-ray pinhole image for shot 38. A ring of emission is visible (corresponding to the copper shielding, with a known aperture of ∼ 1.2 mm diameter), and the laser spot is visible in the middle of this. The trace below the image shows the log intensity of the signal over a narrow y-range, selected to capture the maximum diameter of the aperture. The position of peak signal either side of the aperture is identified (dark blue lines); this is assumed to be 1.2 mm, providing a reference for the pixel scaling. The trace above the image shows the signal intensity over a narrow region corresponding to the maximum diameter position of the laser spot. The full width half maximum of the spot is then identified (teal lines); in this shot, this was found to be 350 µm.

laser spot. The full-width half-maximum (FWHM) of this signal was identified, and for this shot was measured (using the above scaling) as having a diameter of 350 µm. This method is not expected to give a high accuracy for a variety of reasons (the laser spot clearly extends beyond the FWHM, the diameter of the aperture is assumed etc.) - but the good agreement with the expected 400 µm suggests that the phase plates were producing the desired spot, and thus a spot diameter of 400 µm could reasonably be used for estimating the laser intensity. The pinhole data

showed similarly good agreement for all relevant shots.

7.4.3 Measuring laser pulse length using the fiducial and photodiode traces

The laser pulse length could be determined using two methods. Firstly, the photodiode traces were used. Traces were captured for each of the six beams, but these sometimes included noise/spurious signals. The time and amplitude scales of the traces were also independent, which meant it was not possible to know exactly how the beams varied in strength/timing. Nonetheless, a reasonable estimate of the pulse length was found by selecting the traces where a sensible signal was recorded, normalising the amplitude of each trace and centring it in time, and then averaging them to produce a combined signal. The pulse length was then measured as the average full-width half maximum of the individual traces. The six individual laser traces, and the average signal, are shown in Figure 7.16.

The fiducial was also used to measure the pulse length, on shots where this was present and of reasonable intensity. The intensity profile of the fiducial was extracted from the streak camera image, and also plotted as a function of time. The full-width half maximum of this signal was then also taken and used as a measure of the pulse length. This is also shown in Figure 7.16.

The agreement between these two measures was generally good. Figure 7.16 shows that averaging the traces gives a pulse with similar profile to the fiducial, suggesting that it is possible to use either method to give a reasonable estimate of the true profile. It was also found that these measured pulse lengths generally agreed well with the pulse length that was requested for the shot.

Where either the traces or fiducial were present for a shot, the measured pulse length was used in place of the requested pulse length for estimating uncertainty. On shots such as that in Figure 7.16 where both were present and the fiducial was of good signal strength the fiducial was used as the more accurate measure, since the fiducial was a measurement of the combined rather than individual beams. The exception to this was for some of the lower intensity shots, where the fiducial

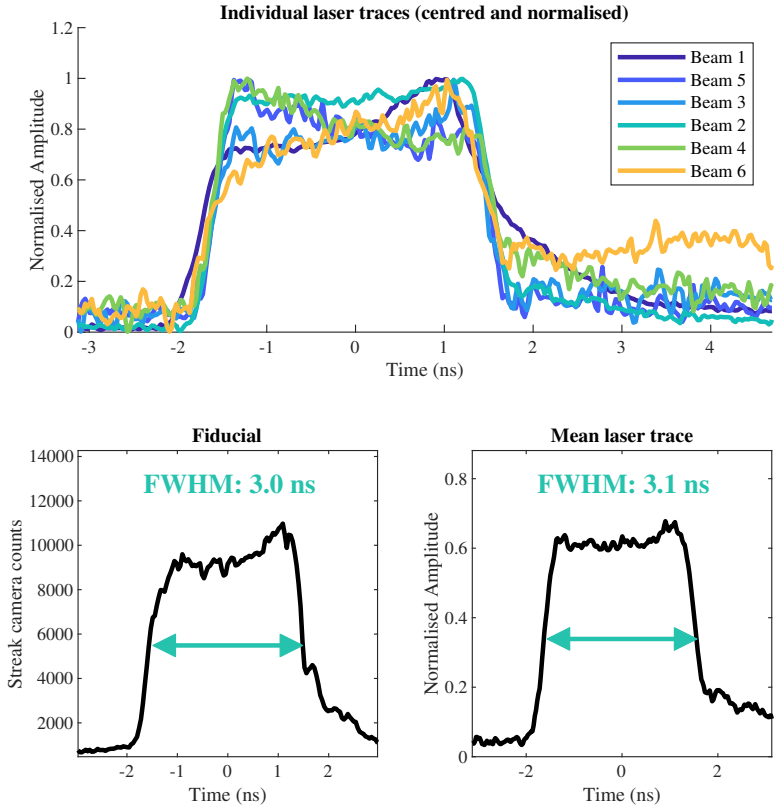


Figure 7.16: The laser pulse, as detected using the two methods, for shot 53. The top plot shows the individual traces for the 6 beams, each normalised and centred. The bottom right plot then averages these to approximate the overall pulse. The bottom left plot shows the intensity profile of the fiducial. The fiducial and average pulse traces show good agreement. The full width half maximum of the fiducial and the combined traces also show good agreement, and either can therefore be used to approximate the pulse length.

strength was much lower; in this case the traces gave a much more detailed and realistic pulse profile, and so these were used instead.

7.4.4 Estimating pulse intensity and comparison with shock strength

The laser energy for each beam was recorded using the on-shot calorimetry. This therefore allowed the laser intensity to be estimated, using the 400 µm spot diameter and the measured pulse lengths. This was performed for each shot.

Figure 7.17 plots the shock velocity achieved in both the quartz and foam for

each shot, as a function of this measured intensity. This provides a simple sanity check on the data; it is expected that the shock strength (and thus shock velocity) in both materials would increase with the laser intensity, and this is indeed observed. Linear trend lines have been included to demonstrate this for both materials. These trend lines are included simply to demonstrate that the shock velocity for each material does increase with intensity, and not to suggest that a linear relationship is expected (or provides the best fit) between these variables.

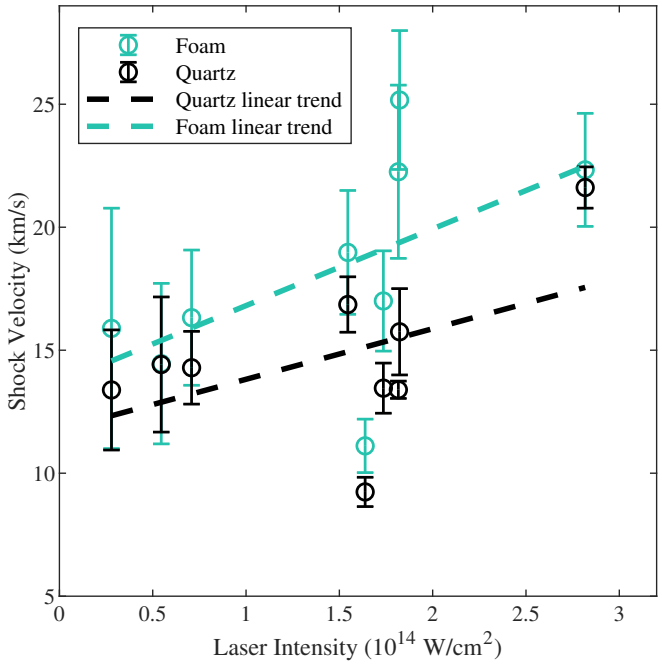


Figure 7.17: The measured shock velocity in the quartz (black) and foam (teal) as a function of intensity. The dashed lines show a linear fit to these two data sets, showing a general trend of increasing shock velocity with intensity. The solid lines show the results from pre-experiment 1D simulations.

Figure 7.17 shows an interesting contrast from the pre-experiment simulations (displayed in Figure 6.12). For a given intensity, the shock velocity measured in the experiment in both materials is significantly lower than the simulations suggest. While some difference from 1D simulation would be expected, this is a significant discrepancy. This adds to the discussion in Section 7.2.4, where it was observed that the shock pressures achieved in the quartz were significantly lower than had been predicted experimentally. This discrepancy is also seen in the comparison with post-experiment simulations in Section 7.7, and is explored further in Section 7.8.

7.5 Shot selection for final data set

With the main analysis of the data now performed, this final analysis section considers how decisions were made regarding the data set used for the final results. Subsection 7.5.1 describes artefacts seen in the VISAR and SOP data on certain shots, and how this is interpreted as potentially corresponding to a shock merger occurring in the quartz. Subsection 7.5.2 then outlines the confidence ranking procedure that was used to select shots for the final results, ensuring shots that returned low quality data or such artefacts were omitted.

7.5.1 Possible evidence of second shocks

A small number of shots showed unexpected additional changes in signal in the VISAR and SOP data. The VISAR data would sometimes display a sudden increase in quartz signal strength part way through the quartz transit time. In the SOP, there would sometimes appear to be distinct periods of emission in the quartz. These two behaviours were often correlated. VISAR and SOP data for a shot where this is observed are displayed in Figure 7.18.

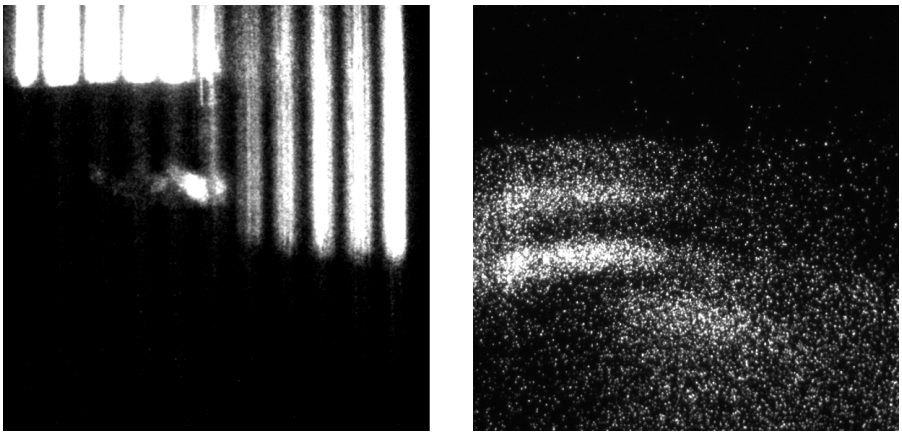


Figure 7.18: VISAR (left) and SOP (right) data for shot 62, displaying a possible second shock. A sudden increase in the intensity of the VISAR signal is seen as the shock transits the quartz, which could correspond to a higher pressure shock with enhanced reflectivity. In the SOP signal, two distinct bands of emission within the quartz are observed.

The most likely explanation for this is the formation of a second shock, which catches and merges with the primary shock within the quartz layer. This leads to

a higher shock pressure being achieved in the quartz later in time (and thus an increased reflectivity and large VISAR signal), and an increase in shock temperature (and thus a region of second emission). Such a behaviour means that the velocity of the shock front changes during the quartz transit, and thus the calculated average shock velocity is not physically meaningful. This then prevents an accurate impedance matching calculation from being performed. For this reason, any shots where such behaviour was observed were not included in the final results for this experiment.

The second shock phenomena was observed in the pre-experiment hydrodynamic simulations for higher intensities, and is discussed in Section 6.3.2. The increased quartz thickness used in the experiment (see Section 6.4.1) made this problem more likely, as did the higher intensities used (in an attempt to generate stronger shocks).

This effect was investigated in post-shot simulations, and is discussed in Section 7.7.3. These simulations suggested that the second shock could indeed be an issue in this regime, but it was not possible to predict/recreate this behaviour for specific shots. This is a difficult phenomena to simulate - the state behind the first shock is already complicated, and will have some error/uncertainty compared to the real experimental conditions, and thus describing the transit of a second shock through this uncertain state is likely to have much larger uncertainty associated with it. In addition, later sections discuss is a discrepancy between the shock strength in the quartz layer seen in the experiment relative to the simulations; such an effect would also impact any secondary shocks, leading to further uncertainty surrounding this effect. By omitting the shots where this behaviour is observed it is assumed that this effect has not had a significant impact on the experimental results; however, it is possible that in some shots the second shock would have undergone recombination in the foam, and this may not lead to an obvious signal.

7.5.2 Confidence ranking of shots

Only a fraction of the experimental shots resulted in VISAR images where all three necessary transitions could be confidently identified. Often it was not possible to

select a particular breakout time from the VISAR data, or it was unclear if there was really a change in signal there. To account for this, a confidence ranking of each shot was performed as the data was analysed.

Each shot was given a score between 0 and 10. A 10 meant that it was absolutely certain that the times identified corresponded to those signals. A 0 meant that there was no timings which could be determined. A range of factors contributed to this ranking. Agreement between the two VISARs resulted in higher confidence, while unexpected artefacts or behaviours in the signals led to lower confidence (as did unreasonable calculated values). This was an iterative process, which was returned to with later information (such as evidence of second shocks, or cross-comparison with SOP data). For the final results only data with a confidence value greater than 5 were used, which removed shots where there was low confidence that real behaviours were being measured.

7.6 Results and discussion

The foam shock states achieved for each shot in the final data set are displayed in Figure 7.19, in the (u_p, P) plane. These data points are compared to three theoretical Hugoniots generated using different equation of state models; the SESAME EOS for CH [176], as well as a quotidian EOS (QEOS) model for both CH and TMPTA [13], generated using the `HUGONIOT` utility packaged with `HYADES` and provided by Cascade Inc. In each case, the Hugoniots are generated for the homogeneous material at the foam density, rather than for a foam material. The QEOS models demonstrate that for the compression behaviour there is no noticeable difference between CH and TMPTA, and thus it is a reasonable comparison to use the CH SESAME table (there is no TMPTA SESAME table to be compared to).

The data is well described by all three of the models, although the large error bars unfortunately prevent a conclusion being drawn about which model provides the best fit. As is expected, the error is dominated by the experimental uncertainty in the quartz and foam shock velocities (with the Hugoniot error for instance contributing very little). It can be seen that the highest pressure point is not

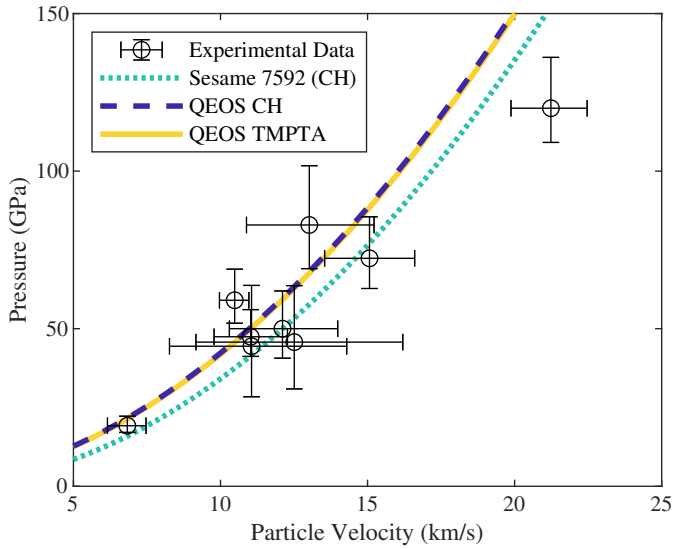


Figure 7.19: The foam Hugoniot data calculated from the impedance matching calculation, compared to Hugoniots generated from theoretical equation of state models.

well described by any of the three models, and is at significantly lower pressure than these models predict. This could be indicative that the models describe the material less well at high pressure; however, it is unwise to make any conclusions about this based on a single data point.

Overall, this suggests that in the pressure range displayed in Figure 7.19, the three theoretical models used here are sufficient to describe the material to within the experimental uncertainty. This is notable due to the fact that they are based on the homogeneous material; and this therefore suggests that this approximation is a valid one to this level of uncertainty for this material. This is a useful result, as it means that such foams should therefore be able to be simulated to reasonable accuracies using these models, rather than requiring foam specific versions.

The grey-body foam temperature is displayed in Figure 7.20. There are fewer data points on this plot, as the SOP did not return good data for all the shots shown in Figure 7.19. Here, it is clear that the theoretical models do not well describe the experimental data; all three theoretical models pass above all of the data points, and many of the data points are outside of the error of these predictions. This is despite the large error bars, which originate from the large experimental uncertainty in the shock velocities, along with the large uncertainty in the foam reflectivity model.

However, there is not sufficient confidence in this temperature data to make firm conclusions about the accuracy of the theoretical models. As noted previously, the SOP signal in this experiment was very weak; it was often difficult to detect any foam signal, and the streak camera had to be operated on maximum gain. As such, this temperature data should be used as an initial indication of this only; it may well be that there is a temperature difference from the models, but this needs to be explored with more accurate temperature measurements before this could be concluded with any degree of certainty[7].

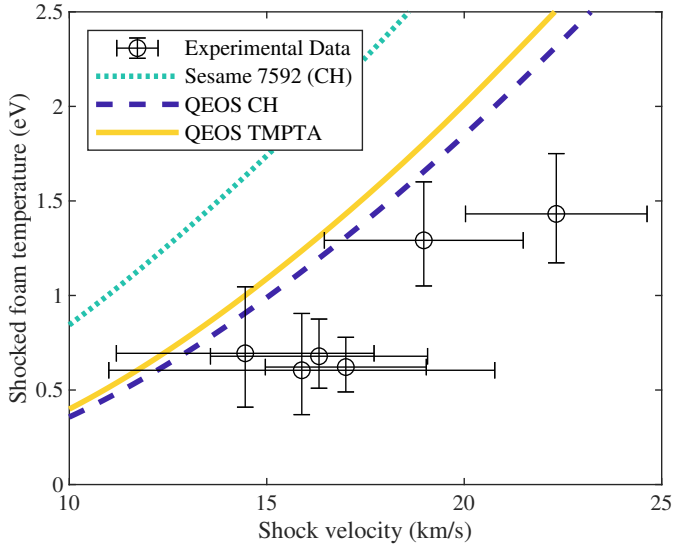


Figure 7.20: The calculated grey-body temperatures of the shocked foam, as a function of shock velocity, compared to the theoretical equation of state models.

These experimental results can be compared with those obtained in previous experiments, looking at similar materials. In particular, the data here is compared to the results of explosively driven experiments performed at Los Alamos on CH foams [177], absolute Hugoniot measurements of CH foam performed using the Naval Research Laboratory's Nike laser facility [178], and impedance matched TMPTA-plastic foam experiments performed at the LULI facility at Ecole Polytechnique [179]. These experiments explored foams at a range of densities, which leads to an interesting comparison - as seen in Figure 7.21.

[7]It is worth noting that if a low foam intensity was measured (such as if the signal was too low to be accurately detected), this would also result in an underestimate of the foam temperature.

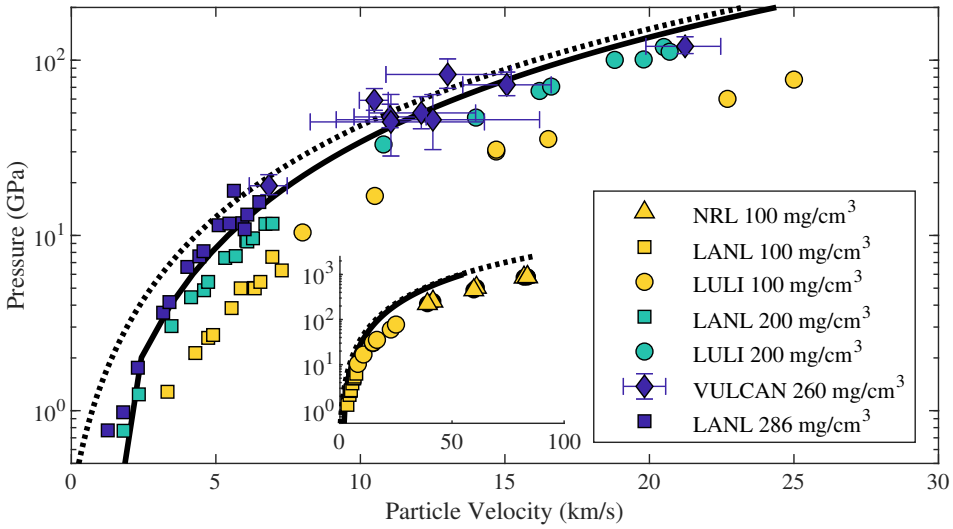


Figure 7.21: A comparison of the experimental data points with CH and TMPTA data from previous experiments. The different colours represent different initial densities, while the different symbols represent the facility they were performed at. The NRL data has a much larger particle velocity than the other experiments, and thus is not visible on the main plot; the inset therefore shows the 100 mg cm^{-3} data over a wider velocity range to include these results.

In a plot of (u_p, P), as seen in Figure 7.21, the Hugoniot dependence on density is observed. For higher initial density foams (darker colours in the figure), the experimentally achieved pressure is higher. There are three rough density groupings in the figure, and as expected the data forms three curves based on this. This trend is consistent between data sets and facilities, and is seen across the full pressure range investigated.

The SESAME and QEOS Hugoniots for CH used in Figure 7.19 are also repeated here. It is seen that the LANL data for a similar initial density (slightly higher at 286 mg cm^{-3}) continues the trend seen in the data from this experiment at lower pressures, effectively extending the pressure range of the data set to cover two orders of magnitude. It is also clear that both experimental models capture the rough trend of the data over this range, but the SESAME model appears to match the LANL data at low pressure more closely than the QEOS model, potentially suggesting that the SESAME model better describes these foams over the wider pressure range.

The same data is also considered in terms of compression ratio vs pressure in

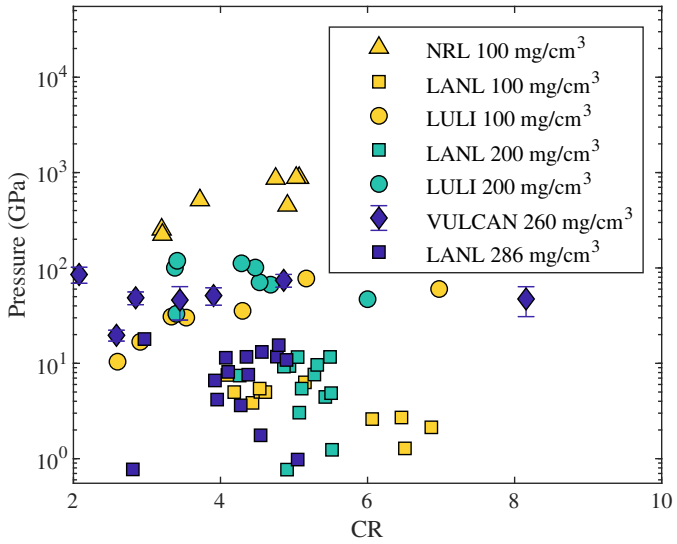


Figure 7.22: A comparison between the different experiments, but in terms of the compression ratio. Horizontal error bars have been omitted as they cover the full horizontal range of the plot. No trends are obvious given the spread of the data.

Figure 7.22, where compression ratio is the density of the shocked material divided by its initial density. However, density is particularly sensitive to the uncertainty in the data, leading to large relative errors in this value [180]. For this reason error bars are not displayed on this plot for clarity, as they cover the full horizontal range. This issue also affects the other experiments, and makes discrimination between models impossible even with the addition of the new data set.

7.7 Comparison with simulation

Post-experiment simulations were also performed in a variety of 1D and 2D codes in order to compare with and help interpret the experimental results. These simulations used the real measured target dimensions. A realistic laser profile was used; the intensity profile of the fiducial was normalised, and scaled to give the appropriate laser intensity (as seen in Section 7.4.3, the fiducial trace was seen to show good agreement with the photodiode traces and thus gave a reasonable approximation of the pulse). Such simulations were performed in the codes `HYADES` (by the author), `HELIOS` (separately by both the author, and by Artem Martynenko

and Paul Neumayer), `MULTI` (by Artem Martynenko and Paul Neumayer), and `FLASH` (by Piotr Rączka).

An example 2D `FLASH` simulation (performed by Piotr Rączka) is displayed in Figure 7.23. The 2D `FLASH` simulations are unique as the only ones to include the 2D step structure, and were thus used to assess the impact of the step on shock propagation. In Figure 7.23 (a), the 2D shock is displayed at four snapshots in time. The step structure is seen to have no significant effect on the shock; the target has a planar shock front across the majority of the laser spot, and this is not significantly disrupted by the presence of the step. The shock propagation as a function of time at a particular horizontal position within the target is displayed in Figure 7.23 (b), allowing the shock transit times through the target to be determined.

Three key features of the post-experiment simulations are discussed in this section. First, the shock propagation times are compared between experiment and simulation. Second, the level of shock stability is investigated. Finally, the shock merger phenomenon is discussed.

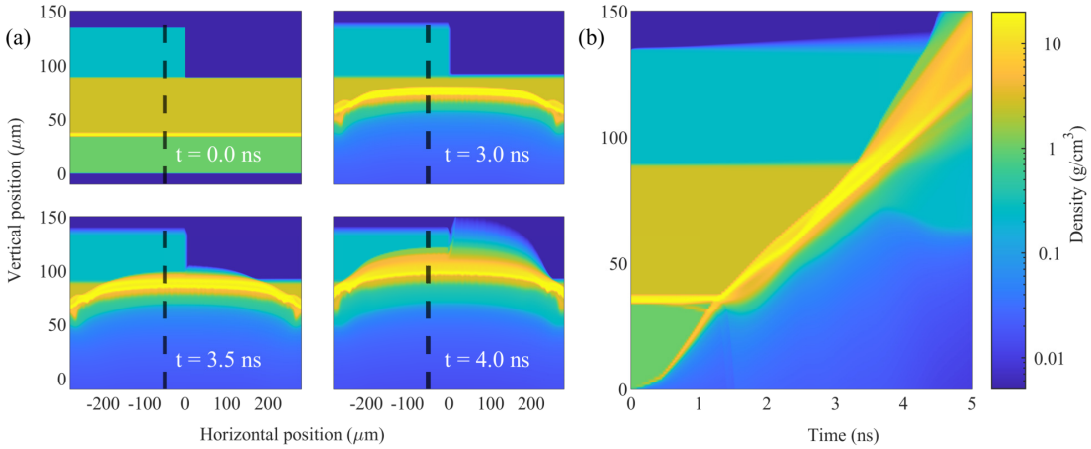


Figure 7.23: 2D `FLASH` simulation (performed by Piotr Rączka) using the measured target dimensions and temporal laser profile for shot 47, with a peak simulated laser intensity of 1.8×10^{14} W/cm^2. (a) shows four snapshots of the shock propagating through the target in 2D. It can be seen that the shock propagation through the foam is largely undisturbed by the step structure. (b) shows the shock propagation vs time at a single horizontal position (50 µm within the target, indicated by the dashed black line in (a)). The same colour scale (representing density) is used in all plots.

7.7.1 Shock propagation time

Simulations were performed in all of the aforementioned codes for shot 47 with a range of laser intensities and thus shock strengths. The same realistic laser profile was used, but scaled to different maximum intensities. For each simulation, the transit time of the shock through the ablator/gold and quartz layers was measured. These timings were compared to the experimentally measured transit times for these layers (determined from the VISAR for the quartz, and from the fiducial for the ablator/gold); this can be seen in Figure 7.24. The figure shows good agreement between the different codes (both 1D and 2D) for these propagation times.

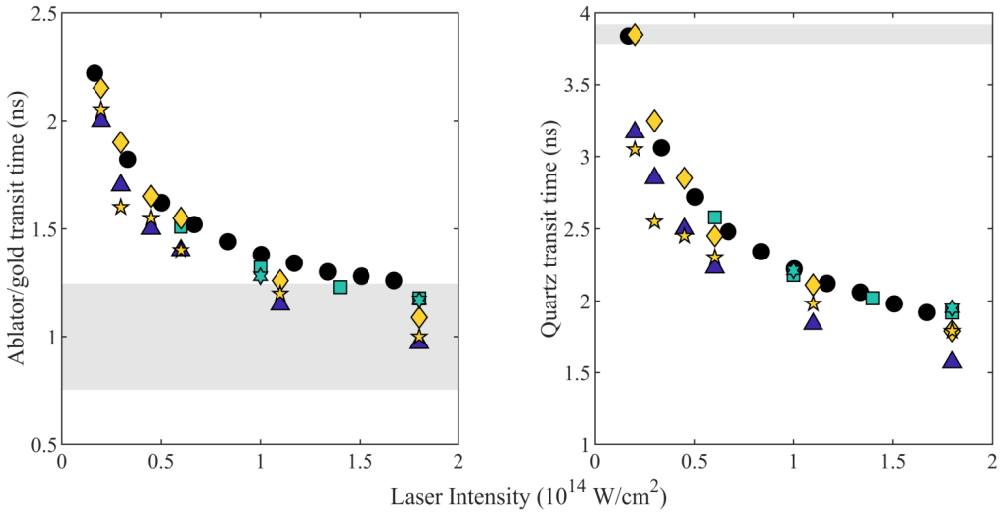


Figure 7.24: Simulated shock transit times through the combined ablator/gold layers (left) and the quartz layer (right) for shot 47. This was performed in four different 1D radiation hydrodynamics codes: `HYADES` (black circles), `FLASH` (teal squares), `MULTI` (yellow diamonds) and `HELIOS` (blue triangles). 2D simulations were also performed in `MULTI` (yellow stars) and `FLASH` (teal hexagrams). On each plot, the grey shaded region corresponds to the experimentally measured value (bounded by the uncertainty). The agreement between codes is good, but in all cases it is not possible to match the experimentally measured times in both ablator/gold and quartz layers at any one intensity. `HYADES` simulations here were performed by the author, while those in other codes were performed by collaborators.

It is clear from the figure that it is not possible to match transit times in the two materials in a single simulation (i.e. for a single laser intensity). This was found to be the case for all the targets. Matching the experimentally measured ablator/quartz transit times requires a high intensity, comparable to the $\sim 1.8 \times 10^{14}$ W/cm^2

intensity used in the real shot. However, matching the quartz transit time requires this simulated intensity to be almost an order of magnitude lower. This is a significant discrepancy, and suggests that something is causing the shock to be significantly weaker in the quartz than in previous layers in the target, and than would be expected based on the simulations.

It was noted throughout the analysis that the quartz shock strength was weaker than expected. This analysis shows that the shock is not weaker across the whole target (i.e. that it is not due to, for instance, a lower than expected laser intensity - something also ruled out by the analysis in Section 7.4), but rather due to a decrease in shock strength within the target. This weaker quartz shock also helps provide an explanation for the lack of VISAR fringe curvature, and further supports the claim made in Section 7.2.4 that this was due to an insufficient quartz shock pressure. This discrepancy between the shock strength in the ablator/gold and the quartz is further discussed and investigated in Section 7.8, where a possible mechanism is proposed and explored.

It is important to note at this point that this effect does not affect the accuracy of the impedance matching calculation, and thus does not undermine the previously presented results. This calculation simply requires that a shock passes through the quartz and into the foam, and that it is reasonably consistent across these two layers; the fact that the shock was previously stronger in the ablator/gold layers is not significant. The shock velocity measurements are not made until the shock has entered the quartz, and thus the shock behaviour in the previous layers is not captured[8].

7.7.2 Shock stability

The simulations also allowed the stability of the shock to be investigated. The impedance matching calculation requires accurate knowledge of the shock either side of the interface - the shock velocity immediately before it exits the quartz,

[8]Of course, the exception here would be if whatever is responsible for this decrease in shock strength led to other effects, such as significant shock decay beyond that seen in the simulation. These second order effects could influence the impedance matching calculation. However, no evidence for such effects was observed.

and immediately after it enters the foam. Yet in this experiment average shock velocities for these two materials were used[9], and thus the calculation therefore requires that the shock is reasonably stable (without significant decay) within these two layers. Any deviation from a stable shock would introduce a systematic error.

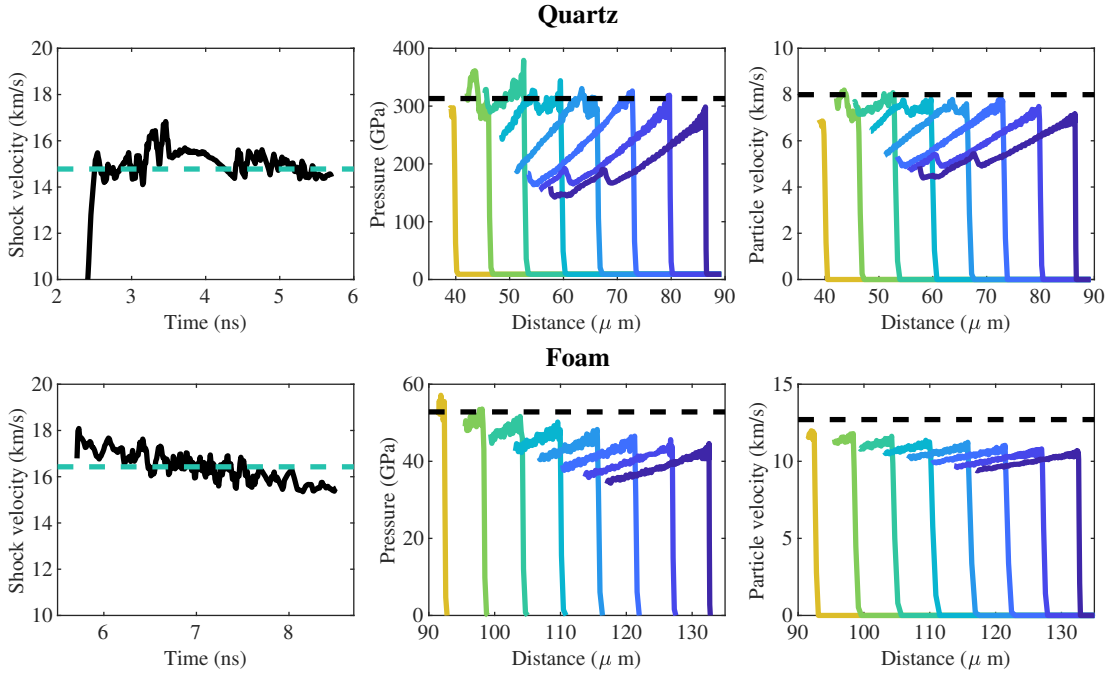


Figure 7.25: Simulated shock velocity, pressure and particle velocity at the shock front as the shock propagates through the quartz and foam layers, from a 1D `HELIOS` simulation (performed by the author) of shot 47 with a simulated laser intensity of 1.5×10^{13} W/cm^2. The shock velocity is displayed as a function of time, while for pressure and particle velocity, profiles of the shock front are provided at eight equally spaced time intervals (later times are further to the right, and are plotted in darker colour). The shock transit times through the quartz and the foam were used to calculate average shock velocities (the dashed horizontal teal lines). The same analysis procedure as used on the experimental data was then used to calculate the pressure and particle velocity in the two materials (the dashed horizontal black lines).

Figure 7.25 shows the simulated shock velocity, pressure, and particle velocity in the foam from a 1D `HELIOS` simulation, again using the realistic laser profile. While the shock is relatively stable, some decay is seen in both materials. This is most significant in the foam shock pressure. The dashed teal line in the shock velocity plots indicates the average shock velocity which would be calculated for

[9]Had the VISAR data displayed sufficient fringe curvature for the quartz shock velocity to be determined as a function of time, the shock stability would be known and the shock velocity could be determined just prior to shock breakout.

each layer based on the shock propagation time, representing the value which would be measured experimentally. The impedance matching calculation was then performed for these two average velocity measurements, and the results of this are represented by the dashed black lines in the pressure and particle velocity plots; these represent the values that would be calculated based on these experimental results. Ideally, these should correspond to the shock state just before the shock leaves the quartz (the darkest trace in the quartz plots) and just after it enters the foam (the lightest foam trace). In the case of Figure 7.25 it is seen that the shock decay does indeed lead to some discrepancy, but this is relatively small and the calculated values are a reasonable approximation of these shock states. Indeed, this suggests that the use of average shock velocities in this experiment was sufficient to give calculated shock states average to within a systematic error (based on the shock decay) of around 10%, which is reasonable and typically less than the random error represented by the error bars in the experimental data.

7.7.3 Shock merger

A second potential systematic error was the second shock which originates in the gold layer, reflects from the ablation front, and potentially merges with the primary shock within the quartz or foam layers. The experiment was designed to mitigate this phenomenon through choice of target dimensions and laser intensities, as discussed in Section 6.3.2. However, the delivered targets were often thicker than originally proposed, and the low shock strength led to larger laser intensities being frequently used. Both of these effects made the second shock more probable. Section 7.5.1 described how some of the experimental shots appeared to indicate that this shock merger had occurred in the quartz, and were thus not included in the final data set.

Simulations were performed for some of the shots where a second shock was seen to occur in the data, using the real laser profiles; however, these did not recreate the observed behaviour. Meanwhile, many of the higher intensity simulations of shots which did not appear to show this effect did demonstrate it. This can be seen in Figure 7.26, which shows the propagation of shots through the Lagrangian

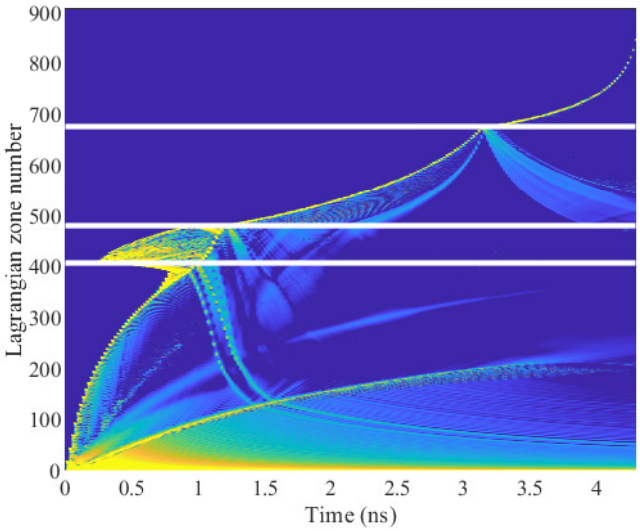


Figure 7.26: Shock propagation plot (log derivative of pressure) from a 1D HYADES simulation of the same shot (47) and laser intensity as in Figure 7.23. It can be seen that a second shock (originating from the gold layer) catches and overtakes the primary shock just before it reaches the foam layer (this is different to the FLASH simulation, where this occurs at an earlier time). The figure is plotted as a function of Lagrangian simulation zone number (rather than radius), so that the zone material boundaries are stationary and the shock propagation can be seen more easily. The material interfaces are indicated by the white horizontal lines.

simulation zones of the target in a 1D HYADES simulation. It is clearly seen here that the second shock merges with the first just as it crosses from the quartz into the foam. This effect was found to be different between codes; the 2D FLASH simulation shown in Figure 7.25 of the same shot shows the merger occurring earlier in the quartz, while in a 1D HELIOS simulation with the same settings the merger does not occur until the shock breaks out from the rear of the target.

This difference between codes demonstrates the difficulty in accurately simulating this effect. The second shock has a particularly large uncertainty associated with it, as it is propagating through material which has already been shocked. Any uncertainty/difference between codes in the first shock will result in a different shock state for this first material, which will then be magnified further with the transit of the second shock. In addition to this, Section 7.7.1 already noted that the simulations cannot accurately describe the shock propagation through the target in a single simulation, as there is an unexplained decrease in shock strength between the ablator/gold and the quartz. This effect would also be expected to have a

significant (but unknown) impact on the second shock, which would also not be included in the simulations. Given this, it is no surprise that it is not possible to simulate accurately whether this effect will or will not occur for a given shot. It is expected that by identifying shots which show behaviour linked to this effect as described in Section 7.5.1, this shock merger was prevented from have a significant impact on the experimental results (although it is worth noting that it may not be possible to identify cases where the shock merger occurred within the foam layer).

7.8 Investigation of the shock strength discrepancy

Throughout the analysis section it was noted that the measured shock strength was significantly lower than predicted in the simulations performed prior to the experiment. This was further highlighted by the post-experiment simulations, which also suggested that the shock strength in the ablator/gold layers was much closer to the simulated results than was the case for the quartz - suggesting that something was leading to a decrease in shock strength between these layers. This section presents some further investigation of this discrepancy, and proposes a potential explanation.

7.8.1 Analysis of non-data targets

As discussed in Section 6.4.1, in some of the early targets the gold/CH coating delaminated from the quartz, and had to be glued back on. The resulting glue layer between the quartz and gold layers prevented the average quartz transit time from being accurately measured, which meant that the quartz shock state could not be determined and thus impedance matching could not be performed. These targets were shot early on as set-up targets while first testing the system.

A disproportionate number of these glue targets displayed fringe curvature compared to the data targets. There were five glue targets shot, and three of these displayed clear evidence of curvature. Of the other two, one displayed no clear signal at all, and in the other it was not possible to determine if there was curvature or not. This is a much better success rate than for targets without

glue - thirty eight of these targets were shot in total, and only seven displayed any signs of possible fringe motion.

As discussed in Section 7.2.4, fringe curvature in this experiment generally correlated to a higher quartz pressure, and this suggests that the glue targets were far more likely to access higher pressures in this experiment than those without glue. As the targets were otherwise alike, this suggests that gluing the gold and quartz together tended to improve the shock strength in the quartz - which in turn suggests that the decrease in shock strength was likely occurring at the interface between these two materials.

7.8.2 Proposed explanation for the discrepancy

The key points are therefore as follows:

- The shock decreases significantly in strength between the ablator/gold layers and the quartz layer, compared to predictions from 1D and 2D simulations.
- This effect is not seen between the quartz and foam layers, suggesting it is not a simple shock decay.
- Gluing the gold to the quartz appears to significantly improve the shock strength in the quartz.

The proposed explanation for these observations is a poor contact between the gold and quartz layers, which led to a loss of shock strength as the shock crosses this interface. This was likely due to a partial delamination between these two layers, resulting in gaps between the two materials. This was something that could well have happened, given that full delamination of these layers was observed in a small number of cases. Gluing the gold on to the quartz likely prevented such delamination from occurring, which explains why adding a glue layer seemed to enable higher shock pressures in the quartz. This is the working assumption for why this discrepancy was observed, as no other explanation was found that could explain all the aforementioned evidence.

7.8.3 Simulation of a gap between layers

Efforts were made to simulate the impact of a gap between these layers. HYADES simulations were performed using a low gas fill between the layers, or a HYADES 'vacuum' zone[10], but neither resulted in any reduction in shock strength in the quartz, nor a significant increase in shock transit time. Further simulations performed by collaborators using MULTI and FLASH also failed to find a substantial change in behaviour when a gap was included.

However, it is possible that such an effect cannot be accurately simulated in one-dimensional hydrodynamics codes. There is no obvious energy loss mechanism in 1D, and so once the gold material crosses the gap, the shock would continue undisturbed. This is demonstrated by the following simple model. When the shock reaches the rear target surface, it undergoes free-surface expansion at roughly twice the shock velocity [49]. If this is approximated as a uniform density plate at the shocked density (a very basic assumption, since in fact the expanding gold will have a strong density gradient), then the impact of the gold and quartz is considered to be a standard plate-flier impedance match experiment (a commonly discussed set-up [49]). The solution to this interaction is a shock wave in the quartz with the exact same strength as if there had been no gap.

However, it is suggested that in a real interaction, there are likely substantially more losses. The shock is not perfectly planar or uniform, and this non-linearity is likely increased as the material crosses the gap. There could also be lateral energy transfer during this process. The process of shock-breakout is likely not perfectly clean and ideal, and all these factors potentially contribute to a reduction in the shock strength.

This was investigated in a series of 2D simulations, performed by Piotr Rączka in a customised version of FLASH (since FLASH does not natively support non-linear boundaries). Three different gaps were simulated with varying degrees of non-uniformity, as seen in Figure 7.27. The level of non-uniformity in the gap does appear to correlate with an increased shock propagation time. However,

[10] A Lagrangian zone which closes once compressed.

this is a very small increase, on the order of tens of picoseconds, and is thus significantly smaller than the nanosecond level discrepancy seen previously. As such, the simulated behaviour clearly not explain this effect. However, it is worth noting that a real gap would likely be highly non-linear, and thus it is possible that this effect could be much greater in practice.

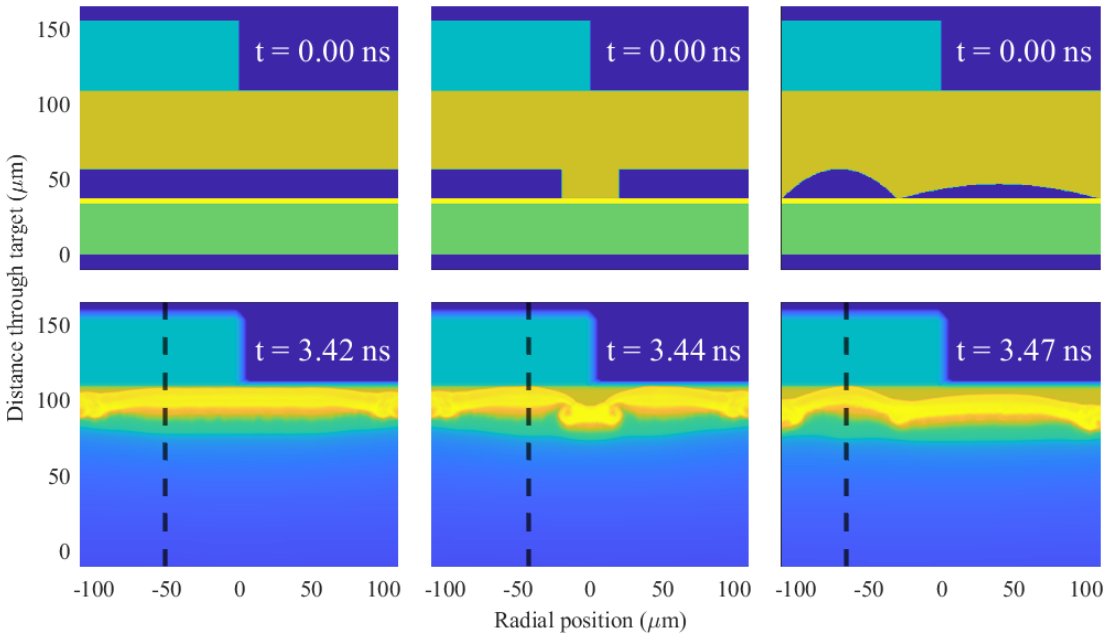


Figure 7.27: 2D simulations of three different gap structures in `FLASH`, performed by Piotr Rączka. The bottom row of plots show the time at which the shock front first breaks out from the quartz, while the dashed black line indicates the position this occurs at. Increasing the non-linearity of this gap structure is shown to increase the propagation time, albeit only by a few tens of picoseconds.

7.9 Suggested improvements

Were this experiment to be repeated, there are a number of improvements which could be made that would improve the quality of the results.

7.9.1 Target improvements

A number of improvements could be made to the target design. Firstly, having the originally designed target dimensions would hopefully help avoid the second shock, and having thinner layers (i.e. the quartz layer was intended to be 40 µm as opposed

to 50 µm) would result in less shock decay and mean average velocity measurements were therefore more accurate. It would also be useful to have enhanced target metrology to probe the existence of gaps and delamination in the target, to help identify whether this was indeed present and leading to reductions in shock strength.

In addition to this, changes to the materials could be made to help make a more robust target. The quartz could be replaced with a reference material where the shock becomes opaque at lower pressures, which would enable the VISAR to be used as intended to determine shock velocity for the lower pressures seen in this experiment. In addition, the gold layer could be replaced with additional CH doped with iodine or bromine [181]. This doped plastic could prevent preheating without the need for the gold - a high impedance material which leads to the generation of the second shock. In addition, this might result in a better contact with the quartz, and thus help with the delamination issue.

7.9.2 Equipment

Improvements to some of the equipment and diagnostics would also aid the experiment. The VISARs and SOP both suffered significantly with signal strength, and a large factor was the streak cameras available. If three cameras with the more sensitive S20 streak tubes were obtained and used, these diagnostics would likely perform significantly better. The probe laser worked unreliably throughout the experiment and resulted in many wasted shots; a probe laser that worked reliably and well would be a significant benefit if this was to be repeated. A power meter or photodiode could be permanently placed to measure the probe laser light reflected from the input beam-splitter; this would provide a measure of the probe laser power (and possibly even pointing) on each shot, which would enable variations in this probe laser to be diagnosed with greater ease.

The SOP had to be built from scratch, and this required it to be used with on-shot calibration. This is likely unavoidable at a facility such as the CLF. However, using a calibrated SOP (such as that at OMEGA) would be a significant advantage

to this type of work, as the well characterised diagnostic would result in significantly less experimental uncertainty (as well as likely being better set-up and aligned).

7.9.3 Experimental layout

The initial plan to inject the probe laser through the VISAR beam-splitter required the diagnostics to be located near the probe laser, which meant placing them on the south side of the chamber. This was a large distance from the target position, and needed the 5 lens optical relay to transmit the captured light with low losses. However, injecting the probe laser inside the target chamber (as was done in the final set-up) meant the diagnostics could be placed on the north side instead.

Doing so significantly reduces the distance from the target, which in turn allows the optical relay to be simplified to a 3 lens set-up. This would make the alignment faster and easier, and would likely also result in fewer losses and thus stronger signals. A new suggested set-up has been created; this is displayed in schematic form in Figure 7.28, and a possible spatial arrangement compatible with this is shown in Figure 7.29. This again produced minimal losses in the ray-tracing code, and results in similar magnifications for all diagnostics (M=44.9 for the VISAR, and M=8.9 for the SOP). The probe laser could also be delivered into the chamber through an optical fibre; this would again decrease losses, and improve ease of set-up.

7.10 Conclusion

The results of the experiment have been analysed. It was found that, for the 20 - 120 GPa pressure range considered, the compression behaviour of the 260 $\mathrm{mg\,cm^{-3}}$ TMPTA foam was well described using QEOS and SESAME equation of state models for CH plastic. These equation of state models assume the foam to be a low density homogenous plastic, and this agreement thus suggests this approximation is sufficient to obtain a reasonable description of the foam under these conditions. While the equation of state models over-estimated the experimentally measured shock temperatures, a lack of confidence in these measurements due to poor signal strength means that further data is required to make firm conclusions on this topic.

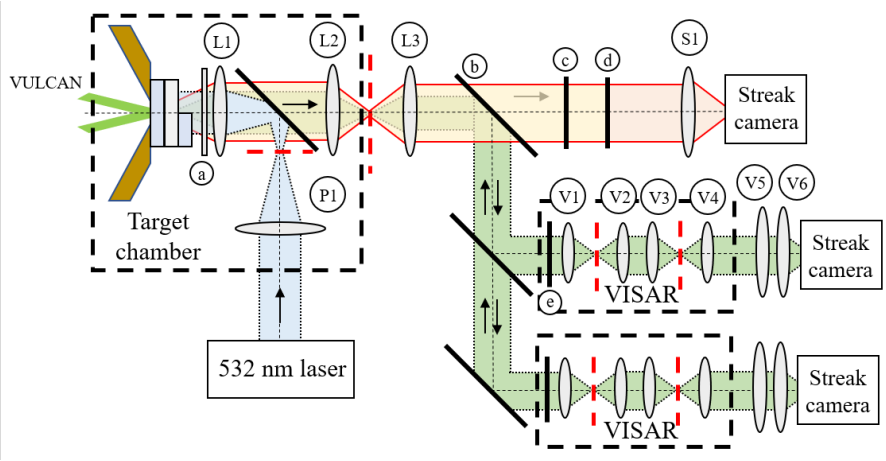


Figure 7.28: Schematic of the suggested new set-up. All labelled components are as for the previous set-ups, except the focal lengths of some lenses have changed: L2 is now 1000 mm, L3 is 300 mm, S1 is 400 mm, and the first lens in each VISAR (V1) is now 500 mm. This schematic does not have the beam collimated as it leaves the chamber, but it would be possible to include this if required.

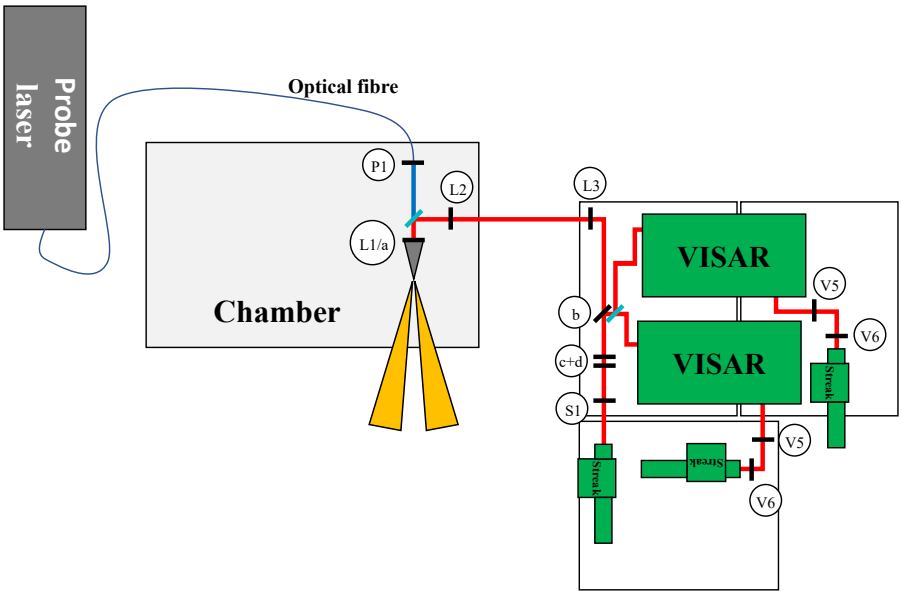


Figure 7.29: A suggested spatial arrangement for the new set-up displayed in Figure 7.28, with labels as in the previous figures. This set-up gives appropriate optical path lengths to minimise losses, while satisfying all other spatial requirements. No full CAD model of this set-up has been produced; this is not accurate to scale, and the lengths in this image are thus approximate.

It was observed that the shock strength in the quartz was significantly lower than expected, based on both simulations and shock-propagation times through previous layers of the multi-layer target. After assessing the evidence, it was suggested that a partial delamination of the quartz from the previous layers is a likely explanation for this. Further comparison was made of the experimental results to simulations in a variety of 1D and 2D codes, and a range of potential improvements that could be implemented in a follow on experiment were suggested.

8

Summary and future work

Contents

8.1 Summary **204**
8.2 Future work **205**

8.1 Summary

This thesis presents a number of new results that are relevant to foam capsules for direct-drive ICF. Firstly, the performance of low convergence ratio implosions of wetted-foam capsules was explored. An interesting new 'low instability' regime was identified, defined through restricting convergence ratio and other key implosion parameters, where instability growth is expected to be minimal. This regime was explored in a large simulation campaign, and promising performance was identified within it. This also established the necessary code and methods to allow other parameter spaces and implosion types to be explored in this way.

This simulation work was then expanded in a variety of different directions. The simulation campaigns were repeated to explore the fusion performance of alternative laser drivers, including higher frequency and novel 'two-colour' implosions. It was demonstrated that moving to higher frequencies allows higher gains to be

generated at low energy, therefore allowing gains closer to that required for an IFE reactor to be achieved even within the 'low instability' regime. Further simulations explored the potential of auxiliary heating of implosions using electron beams for improving fusion performance. It was found that this technique significantly improves the yield of sub-ignition implosions, although the potential of such a scheme is heavily dependent upon how efficiently such beams of electrons are generated and deposited into the plasma.

In order to facilitate future experimental verification of these results and the low instability nature of the regime, surrogate 'hydrodynamic-equivalent' capsules were produced. These demonstrate similar hydrodynamic behaviour to the main implosions but are fielded at room temperature, permitting easier experiments to allow the regime to experimentally validated. Such capsules replace the wetted-foam layer with a 'dry' (without DT-wetting) foam of equivalent overall density (i.e. a higher foam density, to compensate for the lack of DT).

Foams are poorly described by existing equation of state models, and this presents a limit on the accuracy of simulations of these 'hydrodynamic-equivalent' capsules, and wetted-foam capsules in general. To help address this, an experiment was performed at Vulcan to measure principal Hugoniot data of TMPTA foam at 260 $\mathrm{mg\,cm^{-3}}$. Vulcan was used to drive a shock wave through a multi-layer target containing α-quartz (a reference material) and the TMPTA foam, and shock breakout times in the quartz and foam were measured using VISAR, while SOP was used to measure shock temperatures. The results suggested that in the pressure range of 20 - 120 GPa probed in the experiment, the foam is reasonably well described by equation of state models for low density homogeneous plastic. This is encouraging for the accuracy of simulations of the hydrodynamic-equivalent targets, which are based on such models.

8.2 Future work

There are a few research directions suggested by this work. The simulation results in the first chapter suggest promising performance in a regime where

instability is expected to be low, and it would be desirable to test these assumptions experimentally. The hydrodynamic-equivalent capsules were proposed for this purpose, and an experiment should be proposed based on these capsules to validate this regime. This requires further simulation work in higher dimensions, which is being performed through collaboration with the University of Rochester.

There are other routes that could also be taken forward. The auxiliary heating results discussed in Chapter 5 show significant promise, but much more work is required. Firstly, further simulations should be performed where particle-in-cell or Vlasov codes are used to simulate the actual heating mechanism, the results of which are then fed into a fluid code. Work is underway in our group to develop this. Further research is required into electron beams to better understand the efficiency with which these can be created and delivered to the hotspot. Ultimately, it would also be interesting to conduct an experiment where such beams are used to heat a plasma so that this process can be demonstrated experimentally.

Other follow-on experiments could also be conducted following the TMPTA shock compression experiment. While this was an investigation of a dry foam material, there is an open research question regarding the equation of state of DT-wetted-foam. This is a far more challenging experiment to conduct (particularly because it would need to be performed at cryogenic temperatures), but there is significant interest in this topic. Such an experiment should therefore be performed as a high priority.

Finally, the simulation platform developed and used to explore the low instability regime demonstrated a lot of potential to effectively explore fusion performance in different parameter spaces. It could easily be applied to a range of other types of implosion (as demonstrated with the high frequency and two-colour campaigns), but it is currently labour-intensive. This type of optimisation procedure is well suited to machine learning applications, however, and this would remove the need for a user to manually run the simulations and significantly speed up how quickly results could be obtained. It would therefore be valuable to implement machine learning to run this optimisation, and this is again something that is currently being investigated.

Appendices

A

Benchmarking HYADES

Contents

A.1 Direct-drive OMEGA shots 208
A.2 Polar direct-drive NIF shot 211
A.3 Conclusions 213

In this appendix, a series of benchmarking simulations are presented. HYADES simulations were performed (using the same settings as used for the main simulation campaigns in this thesis) of experimental shots performed at OMEGA and the NIF, in order to assess the accuracy of the simulations. This benchmarking was not intended to be exhaustive, but to give some indication of whether the settings used in these simulations were sufficient to describe real experimental data.

A.1 Direct-drive OMEGA shots

The first benchmarking simulations were of 10 calibration shots, performed in a direct-drive configuration on OMEGA. These shots were performed as part of a study of diffusion-dominated mixing [182]. The 10 shots simulated here used capsules containing different HT vapour mixtures, surrounded by either a pure CH shell or a CH shell with 2% D. The shells used in these shots were homogeneous, and

did not have the multi-layer structure described in [182]. Each shot was simulated using the same settings as described for the main simulation campaign, including the 0.8 multiplier applied to the input laser energy to account for laser losses. The simulation results were then compared to the absorbed laser energy, DT and TT yields, along with the bang time and burn width for the three experimental shots where this was recorded.

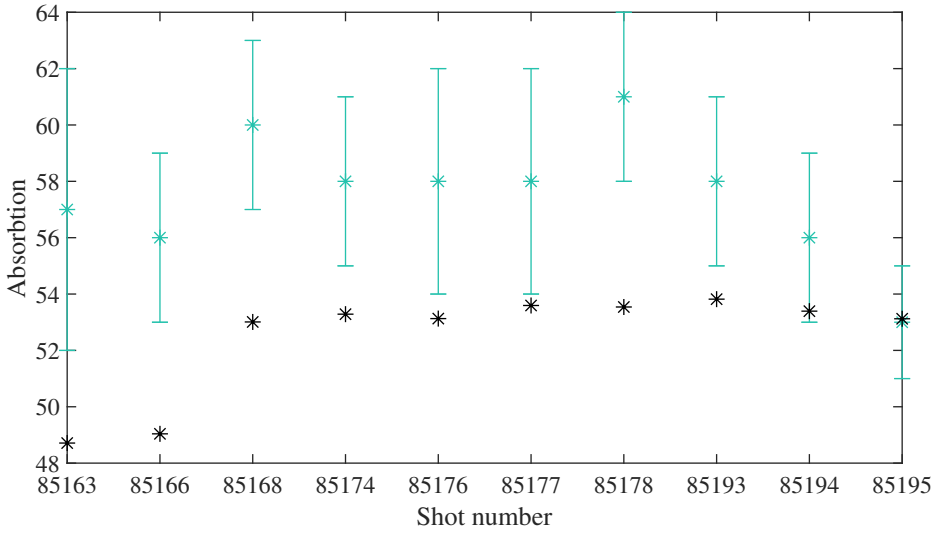


Figure A.1: Percentage absorption of the laser energy measured in simulation (black) and experiment (teal). The simulated data is obtained by measuring the absorption within the simulation, and multiplying this value by the 0.8 input power multiplier used to account for laser losses not described by the code.

The simulated and experimental fractional laser absorption for each of the ten shots is displayed in Figure A.1. In almost all cases, the simulated absorption underestimates that seen in the experiment, and lies just outside the error. This could be improved by increasing the value of the CBET multiplier. However, the discrepancy is relatively small (at most 6%), and so given the 1D nature of these simulations this level of agreement is likely sufficient (especially as underestimating this value means that the simulations are more pessimistic, rather than exaggerating performance).

Figure A.2 shows the DT and TT neutron yields, normalised to the experimental performance (for this reason, the experimental data points all have a value of 1; they are included simply to show the error). The normalised DT yield has a mean of 1.14 and varies between approximately 0.7 and 1.7, while the normalised TT

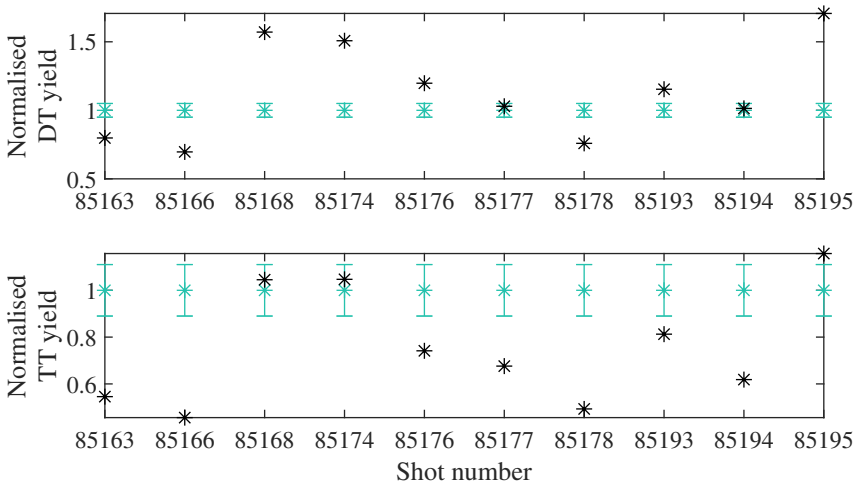


Figure A.2: DT and TT yields for the different OMEGA shots, normalised to the experimental data. The black points are the normalised simulated data, while the teal points with a consistent value of 1 are the normalised experimental data, included to show the error bars.

yield has a mean of 0.76 and varies between approximately 0.4 and 1.2. While there is some clear variation in these values (which should be expected, given the 1D nature of these simulations), it is clear that generally the simulations are providing a reasonable prediction of the experimental performance. It is notable that the simulated TT yield significantly underestimates the experimental yield, while the DT yield is a slight overestimate. This could possibly be indicative of kinetic effects such as fuel stratification [183]; such effects are not included in hydrodynamics codes, and it is known that neglecting them can lead to an underestimation of TT yields [184].

Figure A.3 displays the bang time and burn width (both normalised to the experimental values for each shot). The normalised simulated bang time varies between 0.97 and 1.05, showing close agreement with the experimental value. The normalised simulated burn width however varies between 0.4 and 0.5, and therefore significantly underestimates the experimental results. It is possible this could be due to a difference in definition of burn width; the FWHM of the neutron production rate was used for the simulated results, but details were not provided on how this was calculated from the experimental data.

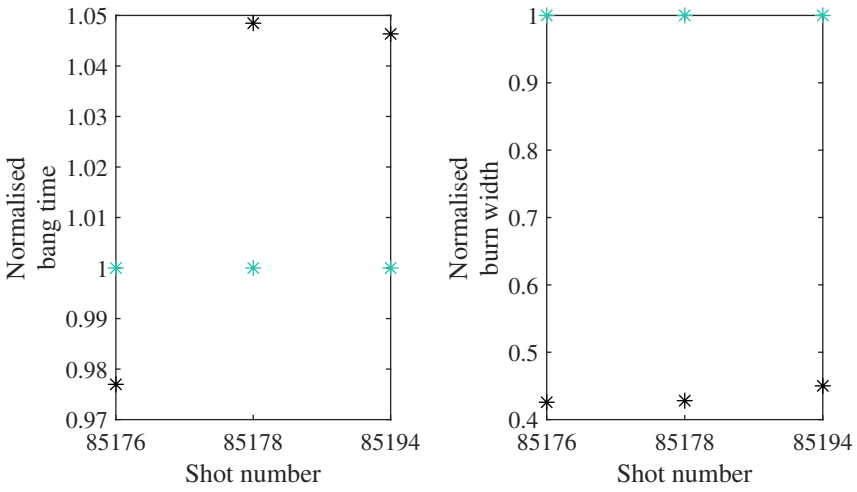


Figure A.3: Bang time and burn width normalised to the experimental value. The black points represent the normalised simulated data, while the teal points with a consistent value of 1 represent the normalised experimental data.

Overall, these results suggest that the HYADES simulations, using these settings and a laser absorption multiplier of 0.8, were sufficient to provide a reasonable estimate of experimental performance in these OMEGA shots. The TT yield and burn width were underestimated, but such parameters are not particularly important for the work discussed in this thesis. The simulations were generally able to match the bang time, laser absorption, and DT yield relatively closely; while there was some variation in the DT yield (and the laser absorption was a slight underestimate), this is to be expected given the 1D nature of the code.

A.2 Polar direct-drive NIF shot

Further benchmarking was performed by simulating the NIF shot N190227, based on data provided by Dr Brian Haines [185]. This was a polar direct-drive shot, and so tested how well such simulations describe a shot in such a configuration. The capsule was filled with DT vapour at a pressure 6000 torr (64% D, 36% T), with a 25 µm CH ablator and an outer diameter of 3943 µm. The laser ramped linearly from 0 to 400 TW over 1.5 ns, remained at this power for 2 ns, and then decreased to 0 TW over 0.5 ns. It produced a DT yield of 1.11×10^{16} neutrons (no error was

given for this value), a bang time of 4.51 ± 0.03 ns, and a burn width of 468 ± 50 ps. The same simulation settings were used as for the main simulation campaigns and OMEGA benchmarking, but a range of different input power multipliers were used (in expectation that the PDD configuration would result in lower laser absorption).

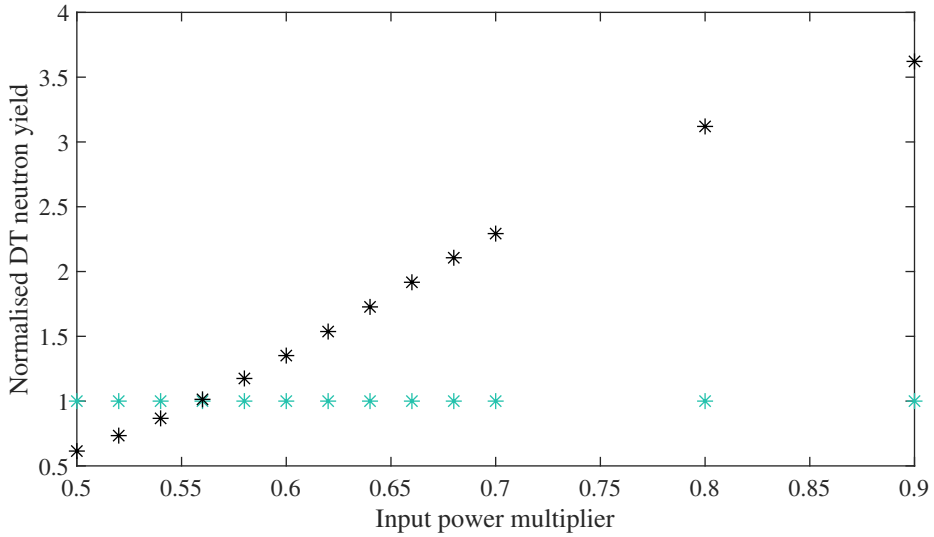


Figure A.4: DT yield for NIF shot N190227 simulated using a range of input power multipliers, normalised to the experimental value of 1.11×10^{16}. The black points are the normalised simulated data, while the teal points with a consistent value of 1 are the normalised experimental data.

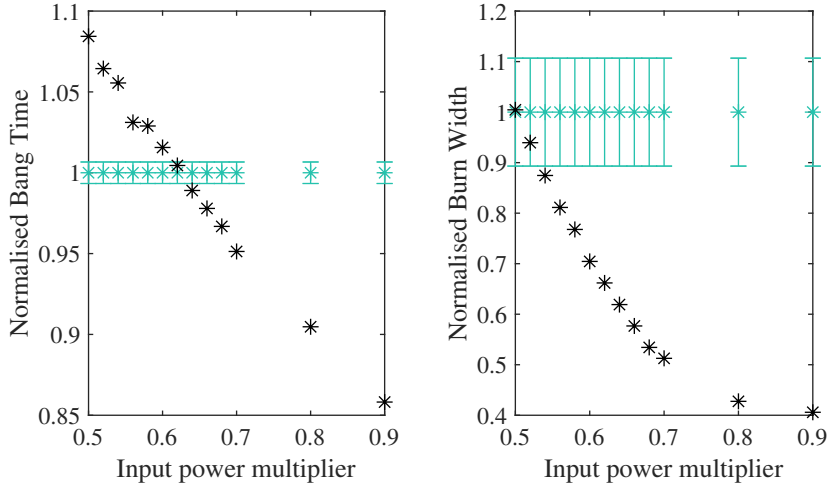


Figure A.5: Bang time and burn width for NIF shot N190227 simulated using a range of input power multipliers, normalised to the experimental values of 4.51 ± 0.03 ns and 468 ± 50 ps respectively. The black points are the normalised simulated data, while the teal points with a consistent value of 1 are the normalised experimental data.

The simulated and experimental DT yield, normalised in each case to the experimental value, is displayed as a function of the laser power multiplier in Figure A.4. The best agreement is observed for multipliers between 0.54 and 0.58, while the multiplier of 0.8 (used for the OMEGA symmetric direct-drive benchmarking) results in a DT yield three times higher than measured experimentally. The normalised bang-time and burn widths are seen in Figure A.5. The bang-time shows much better agreement for all input multipliers, but shows the closest agreement (between 1.05 and 0.95) for multipliers ranging between 0.56 to 0.7. The best agreement is observed for a multiplier of 0.62. The burn width is to within the experimental error only for the lowest multipliers used here of 0.5 and 0.52 (although agreement would also likely be observed at slightly lower values as well). For a multiplier of 0.6 the burn width is approximately 0.7 times that of the experimental value, whilst for a multiplier of 0.8 this falls to around 0.4.

There is no single multiplier which perfectly describes all three data sets, but reasonable agreement for all is seen for input multipliers of 0.55 to 0.6. This is significantly lower than the input multiplier of 0.8 which was sufficient to well describe the OMEGA shots. This difference is to be expected based on the PDD configuration; PDD is well known to suffer from decreased absorption compared to spherically symmetric direct-drive shots, and is particularly susceptible to cross-beam energy transfer. Overall, these simulations of NIF shot N190227 suggest that the simulations presented here provide a reasonable estimate of polar direct-drive experimental performance, but that a lower laser input multiplier is required to account for the reduced laser absorption. These additional losses mean that when simulating a PDD configuration, either reduced performance should be expected, or the laser input power is required to be around 30% higher to achieve the same performance.

A.3 Conclusions

The results in this section show that spherically symmetric direct-drive OMEGA experimental shots are described with reasonable accuracy using `HYADES` with the

settings used in this thesis, including a laser input power modifier of 0.8. However, simulating a polar direct-drive NIF shot required a lower multiplier in the range of 0.55 - 0.6, with a multiplier of 0.8 significantly overestimating the fusion yield. This is likely due to the reduced laser absorption in polar direct-drive. This suggests that simulations performed in HYADES with these settings and a 0.8 input multiplier provide reasonable estimates of performance for spherical direct-drive, but that less agreement is expected if a PDD experimental configuration is used.

The benchmarking here is not exhaustive. A relatively small range of capsules and implosions have been considered, and the polar direct-drive results in particular are based on only a single shot; more simulations of experimental data would be required to make firmer conclusions. However, these initial results are sufficient to indicate that the simulation models and settings used throughout this thesis are reasonable, and are expected to provide a reasonable level of accuracy.

B
Key code

Contents

B.1 Capsule meshing 215
B.1.1 Theory 216
B.2 Ray tracing 220
B.3 VISAR analysis 221

B.1 Capsule meshing

To represent the target in `HYADES` in any of the three available geometries (planar, cylindrical or spherical), a mesh must be created which breaks the target into a large number of smaller zones. The necessary quantities (density, temperature, pressure etc.) are calculated at the centre point of each zone, and the mesh size therefore defines both the resolution and accuracy of the simulation.

The mesh command in the `HYADES` input deck will create a mesh between two specified positions R_1 and R_2, with given number of zones/mesh points n, and a fixed ratio between successive mesh thicknesses α. This is applied to different regions of the capsule to create meshes with different ratios between points.

Typically, there are a few requirements that the mesh needs to satisfy:

- A particular resolution / zone size at the target edges (in spherical, this is at the centre and outer edge of the capsule).

- A small (typically <2%) mass difference between neighbouring zones.

- A sufficiently low number of zones, so that the simulation runs in an appropriate time-frame.

Maintaining a small mass difference between neighbouring zones becomes challenging at the interface between materials of different density. If a low density material is in contact with a high density one, then the zone size in the high density region is required to be much smaller. This often leads to a very large number of mesh points, which makes the simulation unfeasible. Therefore, effort has to be made to design a mesh which satisfies this, while also increasing the zone size where available to reduce the overall number of zones.

The meshing script used for this work was based on the below principles. The mesh was broken up into multiple regions (at least one for every material layer present). Appropriate zone thicknesses were chosen for the edges of these regions, and the theory below was used to find the appropriate number of zones n and scaling α to satisfy these constraints.

B.1.1 Theory

In this section, the theory used to identify the number of zones n and scaling α is described. This was derived from first principles for this purpose. To do this, four values are required: the start and end position of the region R_1, R_2 and the desired thickness of the first and last region t_1, t_2. This provides two constraints; the end zones must have these thicknesses, and the sum of all the zone thicknesses in this region also equal the overall size of the region. This leads to two simultaneous equations which can be solved for n and α.

Derivation

Let $i = 1...N$ be the zone index, where $i = 1$ is the first zone with radius R_1 and thickness t_1 and $i = n$ is the last zone with radius R_n and thickness t_n. α is the ratio between the thickness of successive zones in this region, such that

$$\frac{t_i}{t_{i-1}} = \alpha \ . \tag{B.1}$$

The thickness of zone i can therefore be described with relation to t_1 as

$$t_i = t_1 \cdot \alpha^{i-1} \ . \tag{B.2}$$

The ratio of the thicknesses of the first and last zone can therefore be expressed

$$\frac{t_n}{t_1} = \alpha^{n-1} \ , \tag{B.3}$$

which can be rearranged to give

$$\alpha = (\frac{t_n}{t_1})^{1/(n-1)} \ . \tag{B.4}$$

As both t_1 and t_n are known, this is a first simultaneous equation in terms of α and n.

The known endpoints of the region are used to define the region length. It is clear that the sum of all the zone thicknesses must equal this, and thus

$$L = R_n - R_1 = t_1 \cdot \sum_{i=1}^{n} \alpha^{i-1} = t_1 \cdot \sum_{i=0}^{n-1} \alpha^i \ . \tag{B.5}$$

The thickness of the first zone, t_1, is therefore expressed in terms of the number of zones and the thickness ratio,

$$t_1 = \frac{R_n - R_1}{\sum_{i=0}^{n-1} \alpha^i} \ . \tag{B.6}$$

Applying the standard solution for the geometric sequence,

$$\sum_{i=0}^{n-1} \alpha^i = \frac{1 - \alpha^n}{1 - \alpha} \ , \tag{B.7}$$

gives a final expression for this known thickness of

$$t_1 = (\frac{1 - \alpha}{1 - \alpha^n}) \cdot (R_n - R_1) \ . \tag{B.8}$$

This is rearranged to obtain an expression for the number of zones n in terms only of α and known quantities,

$$n = \log_\alpha(1 - \frac{(1-\alpha)\cdot(R_n - R_1)}{t_1}) \,, \tag{B.9}$$

the second simultaneous equation.

The pair of simultaneous Equations B.4 and B.9 are then solved numerically to return the number of zones n and thickness ratio α required to describe the region with end positions R_1 and R_n and thicknesses t_1 and t_n.

Masses are then also be calculated for each mesh zone, by calculating the zone size and multiplying by the density of the relevant material. For a planar geometry, the zone size is proportional to the zone thickness. However, for spherical geometries, the volume of zone i is proportional to $4\pi R_i^2 \cdot t_i$ (since each zone is actually a thin shell of a sphere).

Implementation

This theory was implemented in a capsule meshing script across multiple regions. The zones must be small at the capsule centre and edges (for accuracy), and larger away from these edges (for speed), while maintaining mass matching at the boundaries. Generally, this problem therefore became setting one region per material layer, and choosing appropriate boundary thicknesses so that 1) there was a valid solution to Equations B.4 and B.4, and 2) neighbouring zones changed mass at an appropriate rate, while requiring as few zones as possible.

The meshing was done as follows. The vapour region was meshed with 150 zones, with a constant thickness ratio. These occur at small radii, so mass matching in this region is impossible. The thickness of the first wetted-foam zone is calculated so that it perfectly mass-matches the last vapour zone. The zone thickness then increases in the foam until it reaches the CH boundary. The CH and foam are again mass-matched, and the CH zone thickness decreases until it reaches the specified thickness at the capsule edge. This set-up has one 'free' parameter - the thickness at the foam/CH boundary.

This foam/CH thickness was varied by trial and error for a single capsule until an appropriate meshing was obtained which satisfied the mass-matching criteria using a low number of zones. A scaling relation for this thickness was then found semi-empirically (some terms were found based on physical arguments, and others based on trial and error), which allowed this thickness to be scaled for variations on the capsule dimensions while still maintaining comparable mass-differences between the ice layers.

The final meshing script used for the optimisation campaign actually used two regions in the CH layer. As this layer had a much higher density than the foam, there was often a large number of zones in this material with a relatively small mass difference. The CH layer was therefore split in two; the first layer maintained a relatively large CH zone size to reduce the required number of zones, while the second layer aimed for the maximum permissible 2% mass matching criteria to bring the zone mass down to the size required at the outer edge. This worked well for the majority of capsule designs considered, but failed in cases where the CH layer was thinner; in these cases, the three layer set-up discussed above was used.

Overall, this meshing script was sufficient for the optimisation campaign to be conducted. It allowed meshes to be automatically produced for a given capsule, and was robust enough to do this over the range of different capsules typically considered. This was essential for running large parameter scans where 100 - 200 different capsules were being simulated at a single time. The meshing script would fail or perform poorly if the capsule began to vary too far from the configuration it was designed for (for instance, if the layer thicknesses began to vary too drastically). Such capsules were found to be in a region of poor performance, where either the shell was very thin (and thus the IFAR and implosion velocity criteria were not met), or very thick (where the ion temperature was too low for significant fusion to occur), and thus this was not a concern. Further development of this code could allow such capsules to be investigated, if this was necessary.

B.2 Ray tracing

Under standard geometric optics, a ray R is described by a vector describing its position y and angle θ as

$$R = \begin{bmatrix} y \\ \theta \end{bmatrix}.$$

Optical components are described as matrices, which are then applied to the ray vector to evolve it as it moves through the optical system [186]. For instance, the matrix describing free space propagation over a distance d is

$$M_{\text{free-space}} = \begin{bmatrix} 1 & d \\ 0 & 1 \end{bmatrix},$$

while the matrix describing the effect of a lens with focal length f is

$$M_{\text{lens}} = \begin{bmatrix} 1 & 0 \\ -\frac{1}{f} & 1 \end{bmatrix}.$$

These matrices are applied in the order that the ray passes through them; if an initial ray R_i passes through component 1 and then component 2, then the final ray R_f is given by $R_f = M_2 \cdot M_1 \cdot R_i$.

The ray tracing code produced for this work allows the user to specify the radius of the target being imaged, the number of rays to be simulated, and the optical set-up consisting of lenses, free space, and 'non-focussing components'. Non-focussing components are used to specify components that do not affect the beam propagation, such as filters and beam-splitters. Any number of components can be specified.

The code then generates the requested number of rays at the target. The position of each ray on the target is generated at random. For each position, the angular limits for which the ray is captured by the first lens are calculated, and an angle for the ray is generated at random within this range (this avoids simulating rays which do not enter the lens set-up). Propagation through each stage of the set-up is then performed on all the rays using the relevant matrix, and the ray position and angle after that stage calculated. At each component position (lens or 'non-focussing'), the code checks to see if the ray radius is within the component radius. If not, the ray is considered 'lost'. The code then calculates the fraction of

rays that make it to the end of the set-up, and uses this to calculate the overall transmission. The transmission as a function of ray position at the target is also calculated in order to assess the level of vignetting present in the set-up.

The code also recognises the image planes in the set-up. It indicates the position of these, and the magnification of the image at that point. A plot is produced showing the ray propagation through the system, the different components, and the position of the image planes - as seen in Figure 6.10.

B.3 VISAR analysis

This section describes how the velocity information encoded in the VISAR fringe pattern is obtained. This approach is based on the method first outlined by Takeda *et al.* [187], and well outlined in [157]. A step-by-step explanation is also presented by [188].

The signal recorded on the VISAR streak camera $S(x,t)$ is described as

$$S(x,t) = B(x,t) + A(x,t)\cos[\phi(x,t) + 2\pi f_0 x + k\epsilon] \, . \tag{B.10}$$

This equation is in agreement with the pattern described in Equation 6.5, but with a few minor adaptations. Firstly, a background term $B(x,t)$ and amplitude term $A(x,t)$ have been included (this was neglected in Equation 6.5, which describes the functional form of an ideal signal). In addition, the linear ramp $kx\sin\theta$ has been expressed as $2\pi f_0 x$ where f_0 is referred to as the carrier frequency (the frequency of the fringes). In this equation, $\phi(x,t)$ is used to refer to the phase difference containing the data (kd in Equation 6.5, where d is now a function of position and time), and as before, $k\epsilon$ is the constant phase offset introduced by the interferometer. This is rewritten in exponential form as

$$S(x,t) = B(x,t) + C(x,t)\exp(i2\pi f_0 x + ik\epsilon) + c.c. \, , \tag{B.11}$$

where

$$C(x,t) = \frac{1}{2}A(x,t)\exp(i\phi(x,t)) \, . \tag{B.12}$$

To analyse this, a 1D Fourier transform is performed of the image, giving data of the form

$$s(f,t) = b(f,t) + c(f - f_0, t) + c^*(f + f_0, t) \; . \tag{B.13}$$

In the Fourier spectrum, this is displayed as 3 peaks: one at $f = 0$, and one each at $f = \pm f_0$. The image is filtered so that only the top or bottom node remains,

$$s(f,t) = c(f - f_0, t) \; . \tag{B.14}$$

The data is then shifted by f_0 to place this peak at $f = 0$, which removes the carrier frequency (the spatial fringes) from the data,

$$s(f,t) = c(f,t) \; . \tag{B.15}$$

An inverse Fourier transform is performed, giving

$$\begin{aligned} D(x,t) = C(x,t) = \frac{1}{2} A(x,t) \exp(i\phi(x,t) + ik\epsilon) \\ = \frac{1}{2} A(x,t)[\cos(\phi + k\epsilon) + i \sin(\phi + k\epsilon)]. \end{aligned} \tag{B.16}$$

Looking at the real (Re) and imaginary (Im) parts of $D(x,t)$, it is clear that

$$\frac{\mathrm{Im}[D]}{\mathrm{Re}[D]} = \tan(\phi + k\epsilon) \; , \tag{B.17}$$

and therefore that the phase difference is obtained from

$$\phi + k\epsilon = \tan^{-1}\left(\frac{\mathrm{Im}[D]}{\mathrm{Re}[D]}\right) \; . \tag{B.18}$$

The resulting data is wrapped between $-\pi$ and π. This is unwrapped to give a continuous curve without discontinuities using a phase unwrapping algorithm, and the constant phase offset $k\epsilon$ is removed to obtain the phase difference $\phi(x,t)$ (this is obtained by performing the same analysis procedure on a reference image obtained from a stationary target, where $\phi = 0$). Finally, $\phi(x,t)$ is converted into the surface velocity $v(x,t)$ by scaling via the VPF (since by definition, the VPF is the velocity change that occurs each complete phase shift - i.e., every phase change of 2π):

$$v(x,t) = \frac{\mathrm{VPF}}{2\pi} \cdot \phi(x,t) \; . \tag{B.19}$$

C

Experiment photos and equipment list

Contents

C.1 Final experimental results 223
C.2 Equipment list 225
C.3 Experimental set-up images 226
C.4 Target images 233

This appendix provides further reference material for the TMPTA foam experiment. A table is given with a full description of the key optics used in the experiment. Then, a series of photographs are provided showing different aspects of the experiment set-up. Full CAD models and technical drawings (produced by Michael Woodward from UKRI-STFC) are also displayed.

C.1 Final experimental results

Table C.1 provides the key experimental data for each of the shots in the final data set.

Table C.1: Principal Hugoniot data for the compressed foam. U_s^{quartz} and U_s^{foam} are the measured quartz and foam shock velocities. The quartz particle velocity, pressure, and temperature (u_p^{quartz}, P^{quartz} and T^{quartz}) are not directly measured, and are instead calculated from the shock velocity using known relations (the Hugoniot [166], and a measured relationship between U_s^{quartz} and T^{quartz} [162]). u_p^{foam} and P^{foam} are the foam particle velocity and pressure, determined from the impedance matching calculation, while T^{foam} is the foam grey-body temperature calculated in the SOP analysis. Other than U_s^{quartz} and U_s^{foam}, where the quoted error is the experimental uncertainty, the errors given for each value are the standard deviation of a distribution for that value obtained from 10,000 random Monte-Carlo iterations. The SOP did not work correctly on all shots, hence the missing T^{foam} values.

U_s^{quartz} (km/s)	u_p^{quartz} (km/s)	P^{quartz} (GPa)	T^{quartz} (eV)	U_s^{foam} (km/s)	u_p^{foam} (km/s)	P^{foam} (GPa)	T^{foam} (eV)
9.2 ± 0.6	$4.3 \,^{+\,0.4}_{-\,0.4}$	$106 \,^{+\,17}_{-\,16}$	-	11.1 ± 1.1	$6.8 \,^{+\,0.6}_{-\,0.7}$	$19.2 \,^{+\,3.0}_{-\,2.1}$	-
13.4 ± 2.4	$7.0 \,^{+\,1.8}_{-\,1.6}$	$249 \,^{+\,123}_{-\,93}$	$0.96 \,^{+\,0.59}_{-\,0.40}$	15.9 ± 4.9	$11.1 \,^{+\,3.2}_{-\,2.8}$	$44.4 \,^{+\,19.3}_{-\,16.0}$	$0.60 \,^{+\,0.30}_{-\,0.23}$
13.4 ± 0.4	$7.0 \,^{+\,0.2}_{-\,0.2}$	$250 \,^{+\,16}_{-\,15}$	-	22.3 ± 3.5	$10.5 \,^{+\,0.5}_{-\,0.5}$	$59.0 \,^{+\,9.9}_{-\,7.3}$	-
13.5 ± 1.0	$7.1 \,^{+\,0.7}_{-\,0.7}$	$253 \,^{+\,47}_{-\,42}$	$0.98 \,^{+\,0.21}_{-\,0.18}$	17.0 ± 2.0	$11.0 \,^{+\,1.2}_{-\,1.2}$	$47.4 \,^{+\,8.6}_{-\,6.2}$	$0.62 \,^{+\,0.16}_{-\,0.13}$
14.3 ± 1.5	$7.7 \,^{+\,1.0}_{-\,1.0}$	$290 \,^{+\,73}_{-\,64}$	$1.14 \,^{+\,0.35}_{-\,0.28}$	16.3 ± 2.7	$12.1 \,^{+\,1.9}_{-\,1.8}$	$50.0 \,^{+\,12.0}_{-\,9.4}$	$0.68 \,^{+\,0.20}_{-\,0.17}$
14.4 ± 2.7	$7.7 \,^{+\,2.0}_{-\,1.9}$	$296 \,^{+\,150}_{-\,117}$	$1.17 \,^{+\,0.76}_{-\,0.53}$	14.5 ± 3.3	$12.5 \,^{+\,3.7}_{-\,3.3}$	$45.7 \,^{+\,17.9}_{-\,14.8}$	$0.69 \,^{+\,0.35}_{-\,0.28}$
15.7 ± 1.8	$8.7 \,^{+\,1.3}_{-\,1.2}$	$362 \,^{+\,102}_{-\,87}$	-	25.2 ± 2.8	$13.0 \,^{+\,2.2}_{-\,2.1}$	$82.9 \,^{+\,18.8}_{-\,13.9}$	-
16.9 ± 1.1	$9.5 \,^{+\,0.8}_{-\,0.8}$	$423 \,^{+\,69}_{-\,62}$	$1.79 \,^{+\,0.36}_{-\,0.31}$	19.0 ± 2.5	$15.1 \,^{+\,1.5}_{-\,1.5}$	$72.3 \,^{+\,13.2}_{-\,9.6}$	$1.29 \,^{+\,0.31}_{-\,0.24}$
21.6 ± 0.8	$13.1 \,^{+\,0.7}_{-\,0.7}$	$749 \,^{+\,69}_{-\,66}$	$3.62 \,^{+\,0.42}_{-\,0.40}$	22.3 ± 2.3	$21.2 \,^{+\,1.2}_{-\,1.4}$	$120.0 \,^{+\,16.1}_{-\,10.9}$	$1.43 \,^{+\,0.32}_{-\,0.26}$

C.2 Equipment list

Table C.2 provides a list of the key components used in the experiment.

Table C.2: Equipment list for the key equipment and components used in the experiment.

Label	Description	Supplier and model
L1	150 mm achromatic doublet (2 inch)	OptoSigma EU, DLB-50.8-150PM
L2	700 mm achromatic doublet (2 inch)	OptoSigma EU, DLB-50.8-700PM
L3	500 mm achromatic doublet (2 inch)	OptoSigma EU, DLB-50.8-500PM
L4	700 mm achromatic doublet (2 inch)	OptoSigma EU, DLB-50.8-700PM
L5	250 mm achromatic doublet (2 inch)	OptoSigma EU, DLB-50.8-250PM
S1	300 mm achromatic doublet (2 inch)	Thorlabs, ACT508-300-A
V5	300 mm plano-convex lens (2 inch)	Thorlabs, LA1256-A
V6	200 mm cylindrical lens	Thorlabs, LJ1653L1-A
b	Dichroic mirror (2 inch, 505nm cut-off)	Thorlabs, DMSP505L
c	500 nm short-pass filter (1 inch)	Thorlabs, FES0500
d	400 nm long-pass filter (2 inch)	Thorlabs, FGL400S
e	532 nm laser-line filter (1 inch, 532 $\pm$ 0.2 nm, FWHM = 1 $\pm$ 0.2 nm)	Thorlabs, FL532-1
VISAR 1 streak	Hamamatsu C7700 high dynamic range streak camera with S20 tube	Hamamatsu C7700
VISAR 2 streak	Hamamatsu C7700 high dynamic range streak camera with S1 tube	Hamamatsu C7700
SOP streak	Hamamatsu C5680 streak camera with S1 tube	Hamamatsu C5680
VISAR laser	Quantel YG980 20ns SLM custom laser	Quantel YG980

C.3 Experimental set-up images

Figures C.1 to C.9 are photographs of the real experimental setup. They have been annotated to show the beam paths, and the key components have been labelled. Figures C.10 and C.11 are technical drawings of the final setup, corresponding to the CAD model in Figure 6.18. This model (and these drawings) were produced post-experiment by Michael Woodward of UKRI-STFC, and describes the real positions and spacing of the components in the experiment.

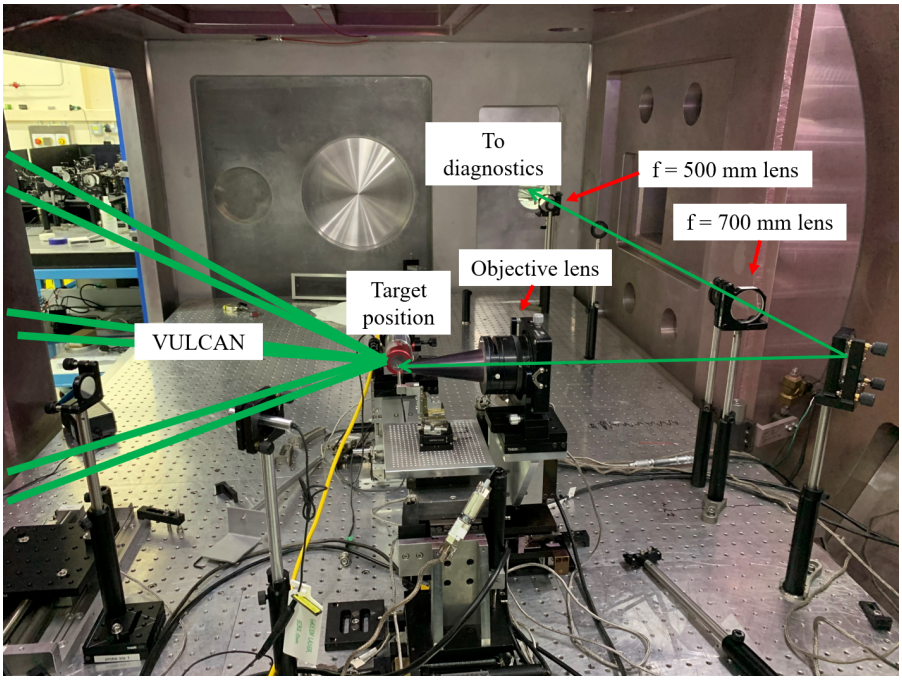


Figure C.1: The target chamber for the initial set-up. The probe laser (thin green line) enters from the chamber window, and passes through the optics to the target. The reflected probe laser light and the self emission are captured by the objective lens, and passes through the same optics and out of the chamber window.

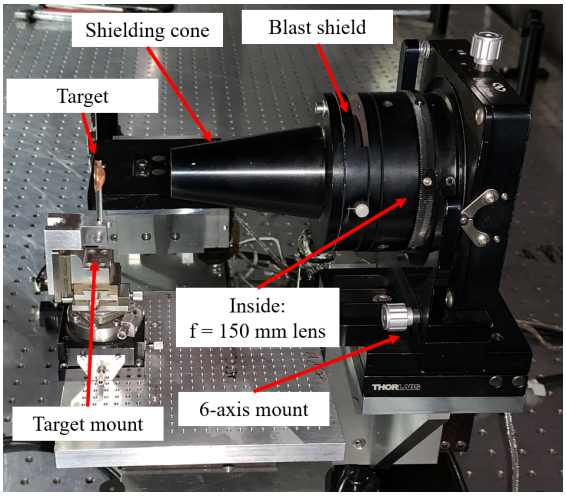


Figure C.2: A close-up image showing the target and objective lens.

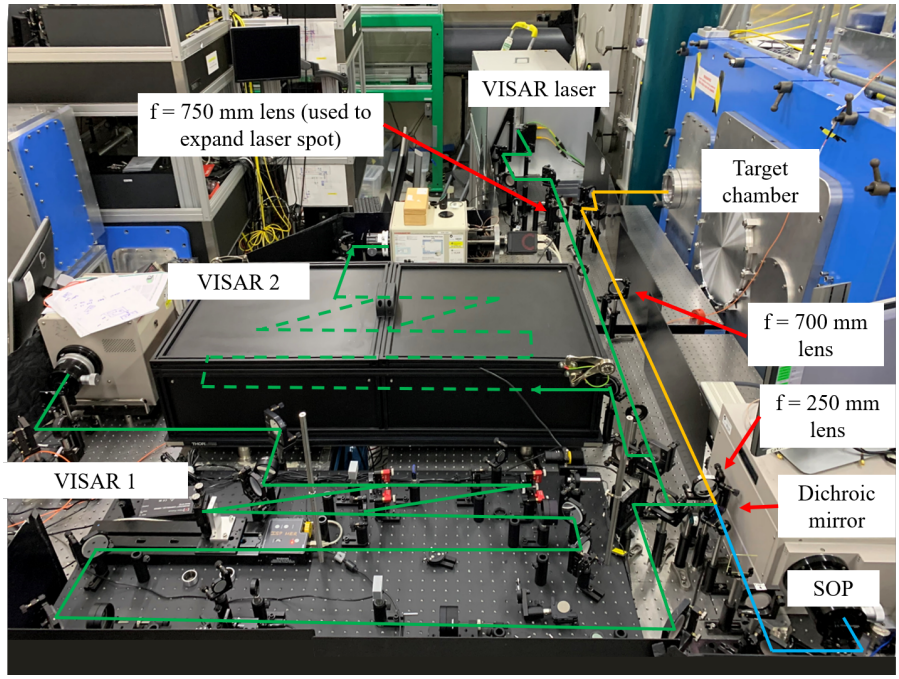


Figure C.3: The initial set-up outside the target chamber. The light leaving the chamber is represented by the orange line. The dichroic mirror separates this into the self-emission (blue) and the reflected probe laser (green). The self-emission continues onwards to the SOP streak, while the reflected probe laser goes to the two VISARs. The VISAR (probe) laser travels into the chamber through the same optical system in the opposite direction, entering the set-up thought the rear of the beam-splitter by the second VISAR. A 750 mm lens is also displayed in front of the VISAR laser which is absent from the schematics - this was used to try and decollimate the probe laser slightly before injection, in an attempt to improve the spot size at the target before the change to the updated set-up.

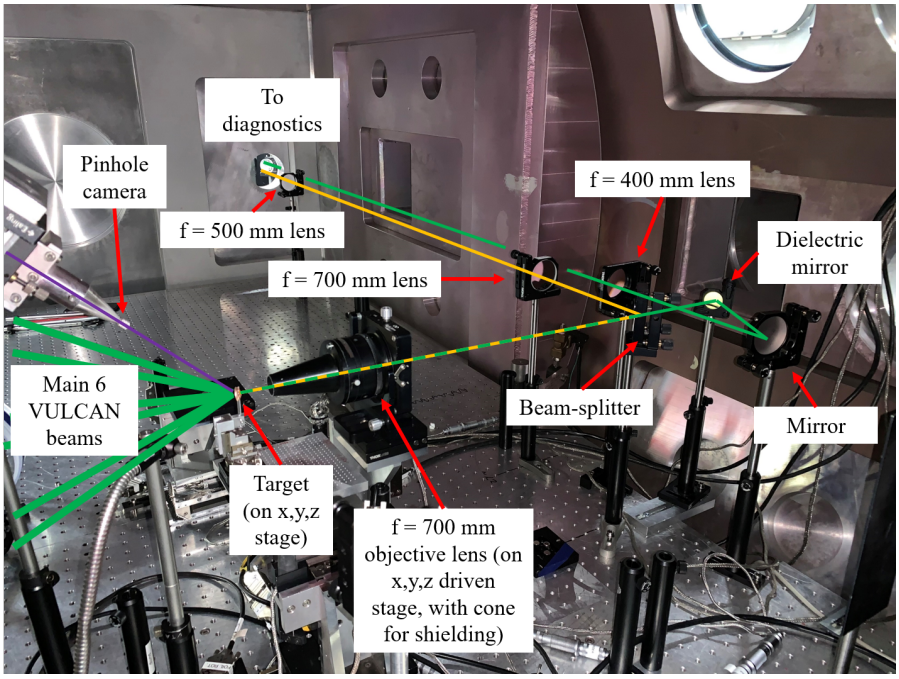


Figure C.4: The target chamber for the final set-up. The probe laser beam path (green) is now separate from the imaging beam path (orange). The first mirror has been replaced with a beam-splitter, allowing the probe laser to be injected into the main path just prior to the objective lens (the 400 mm lens in the new probe beam path forms a lens pair with the objective, such that the objective lens now collimates the probe beam). The objective lens captures the reflected probe light, plus the self-emission, and this is transmitted down the orange beam path and out of the chamber as in the initial set-up. The pinhole camera is also shown in this image, with its line of sight represented by the purple line.

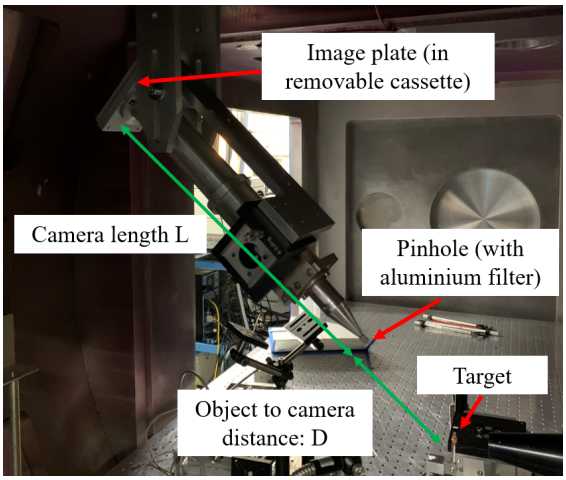


Figure C.5: The pinhole camera, in position. The image plate is placed at the rear of the camera. A thin piece of aluminium is placed on the front of the pinhole to provide filtering of low energy x-rays.

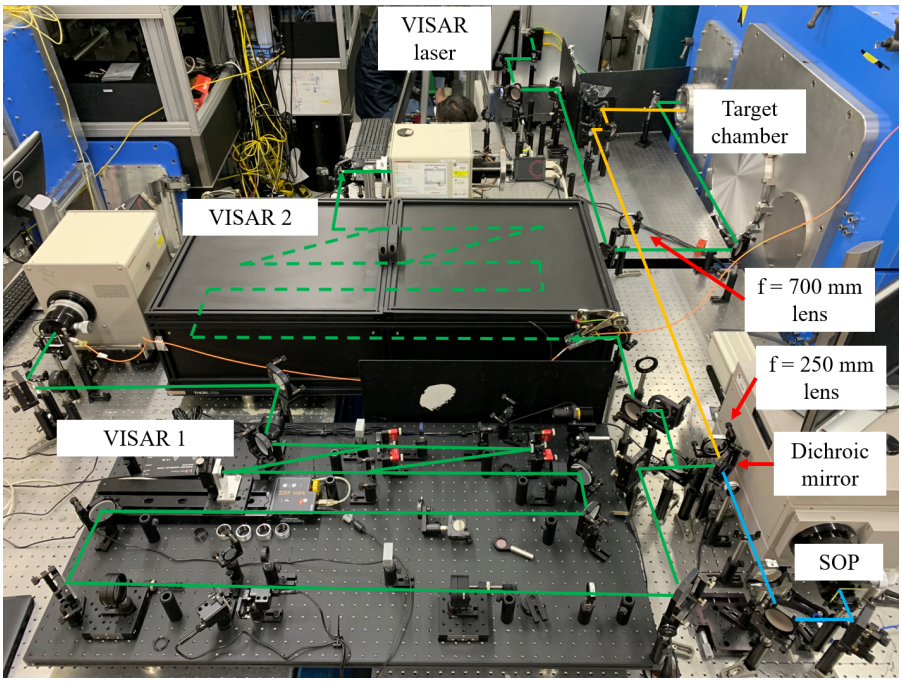


Figure C.6: The final set-up outside of the chamber. The imaging beam lines are unchanged from before, except that the beam-splitter before the second VISAR has been changed to a mirror. The probe laser follows a new route into the target chamber.

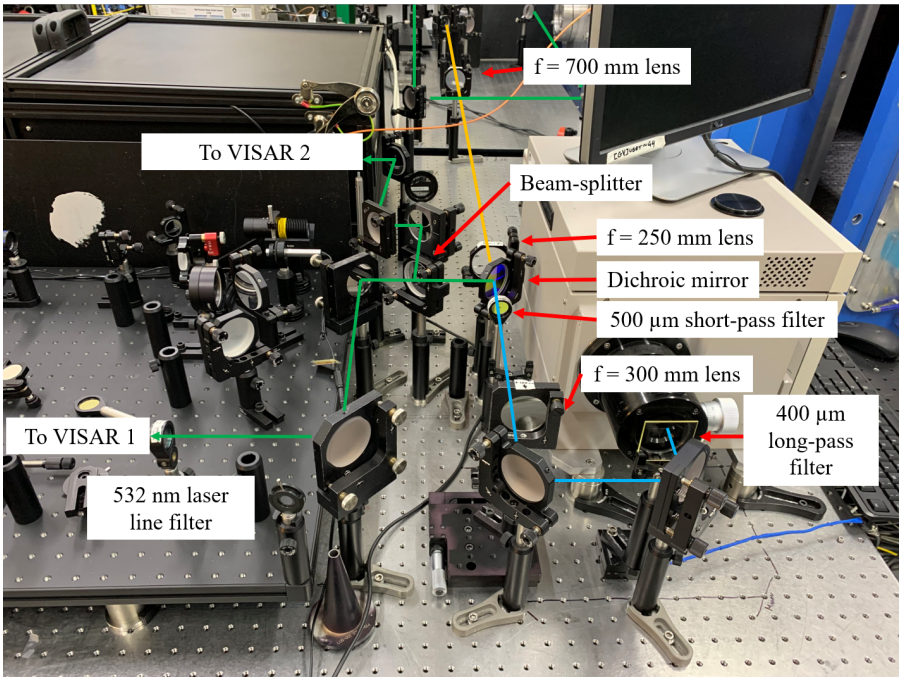


Figure C.7: A close-up image of the components around the dichroic mirror, showing the filters in position. The 400 µm long-pass filter is in a slightly different position to the main schematic; this change in position has no effect.

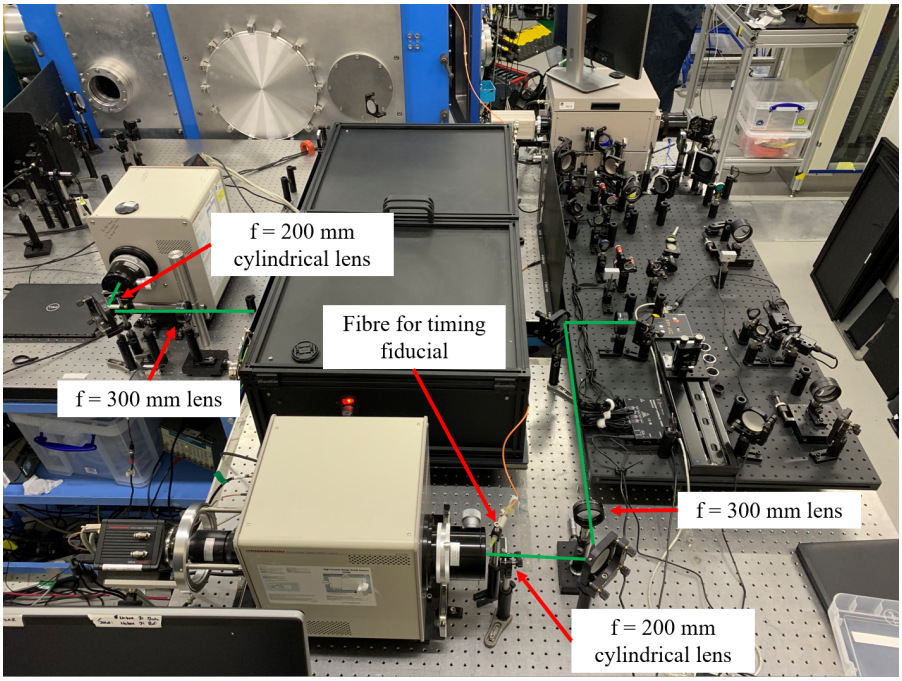


Figure C.8: A close-up showing the beam path as the imaging lines leave the VISARs and are transmitted to the streak cameras. The fiducial captures light leaking through one of the Vulcan mirrors, and shines this on the side of one of the streak cameras.

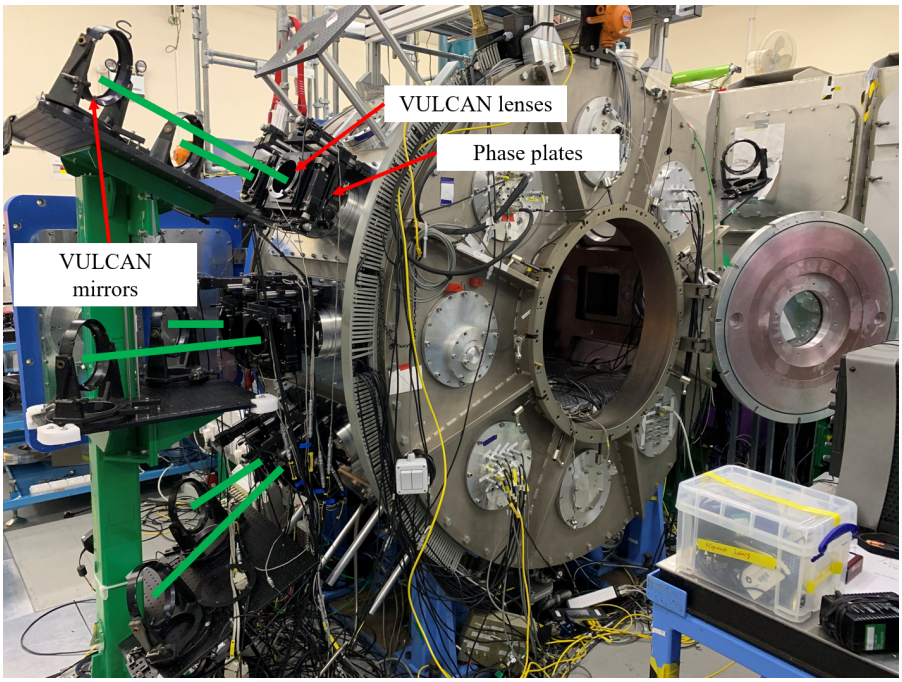


Figure C.9: The Vulcan target chamber, showing the final Vulcan mirrors, lenses, and phase plates. These beams enter through six ports on the chamber and irradiate the target. The two chamber doors are sealed shut so that the chamber can be pumped down to vacuum. A small fraction of the Vulcan light leaks through the mirrors; small photodiodes were placed behind these to capture this signal and thus record the pulse shape.

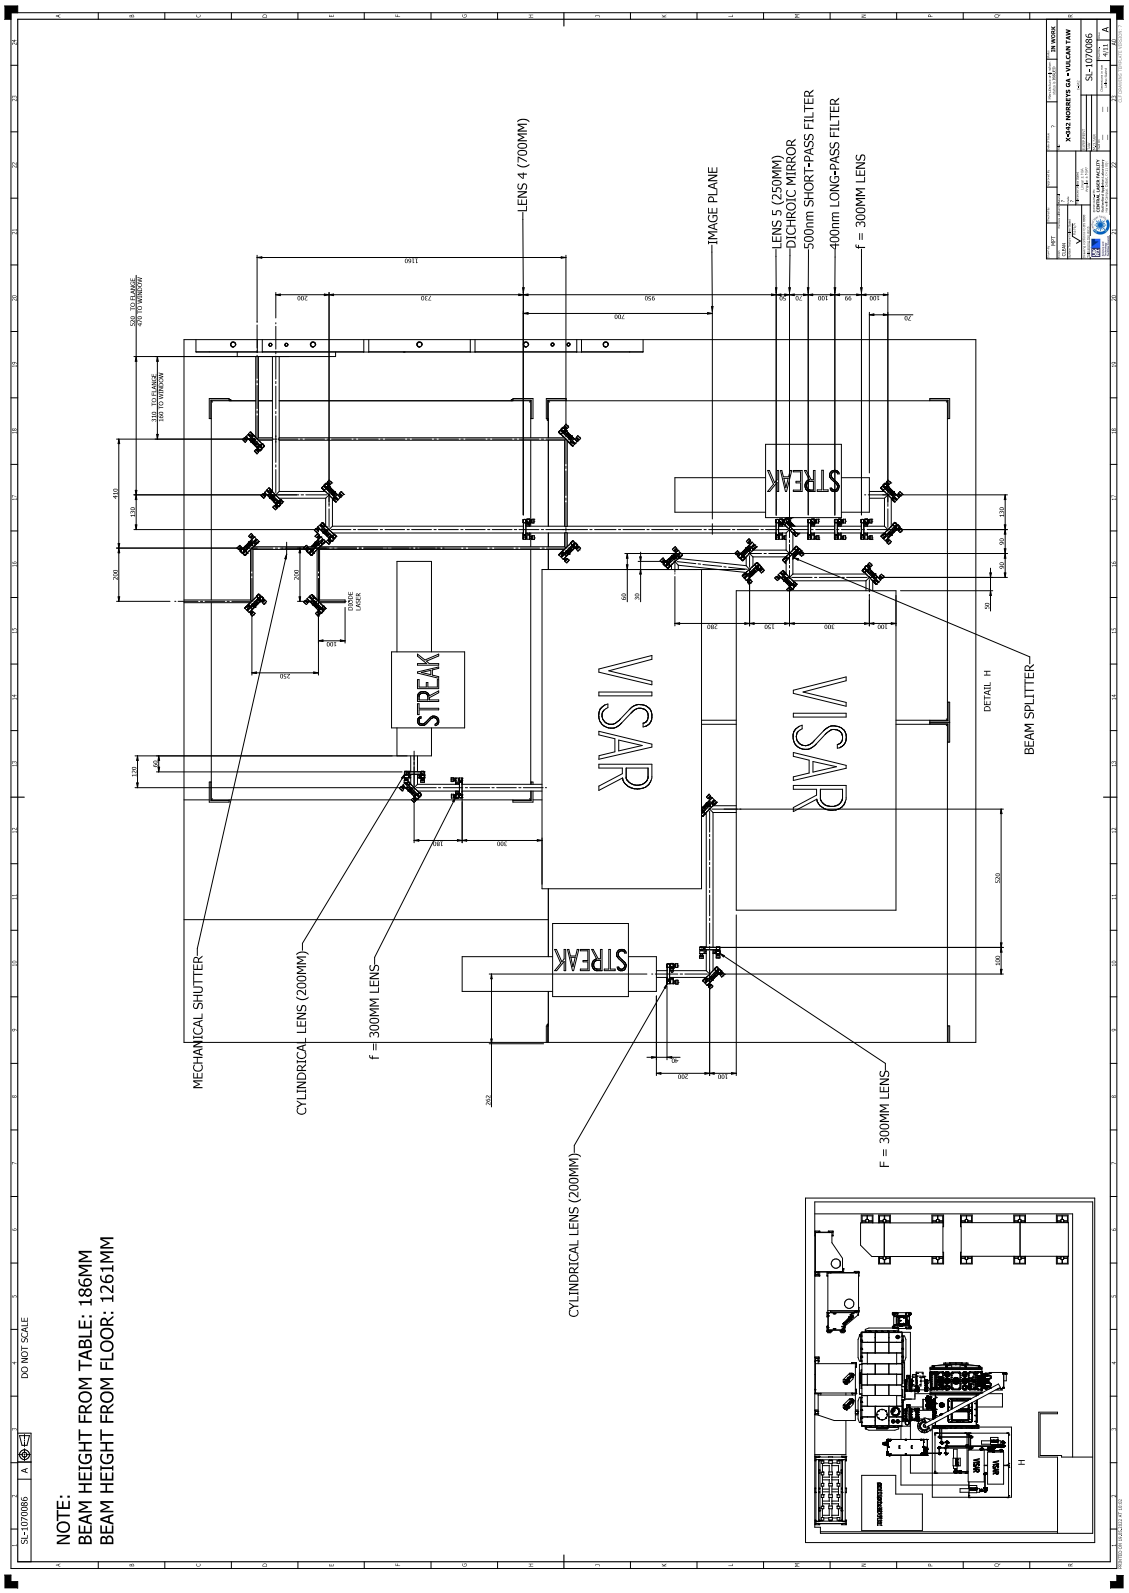


Figure C.10: Engineering drawing showing the optical tables for the final set-up, with full measurements.

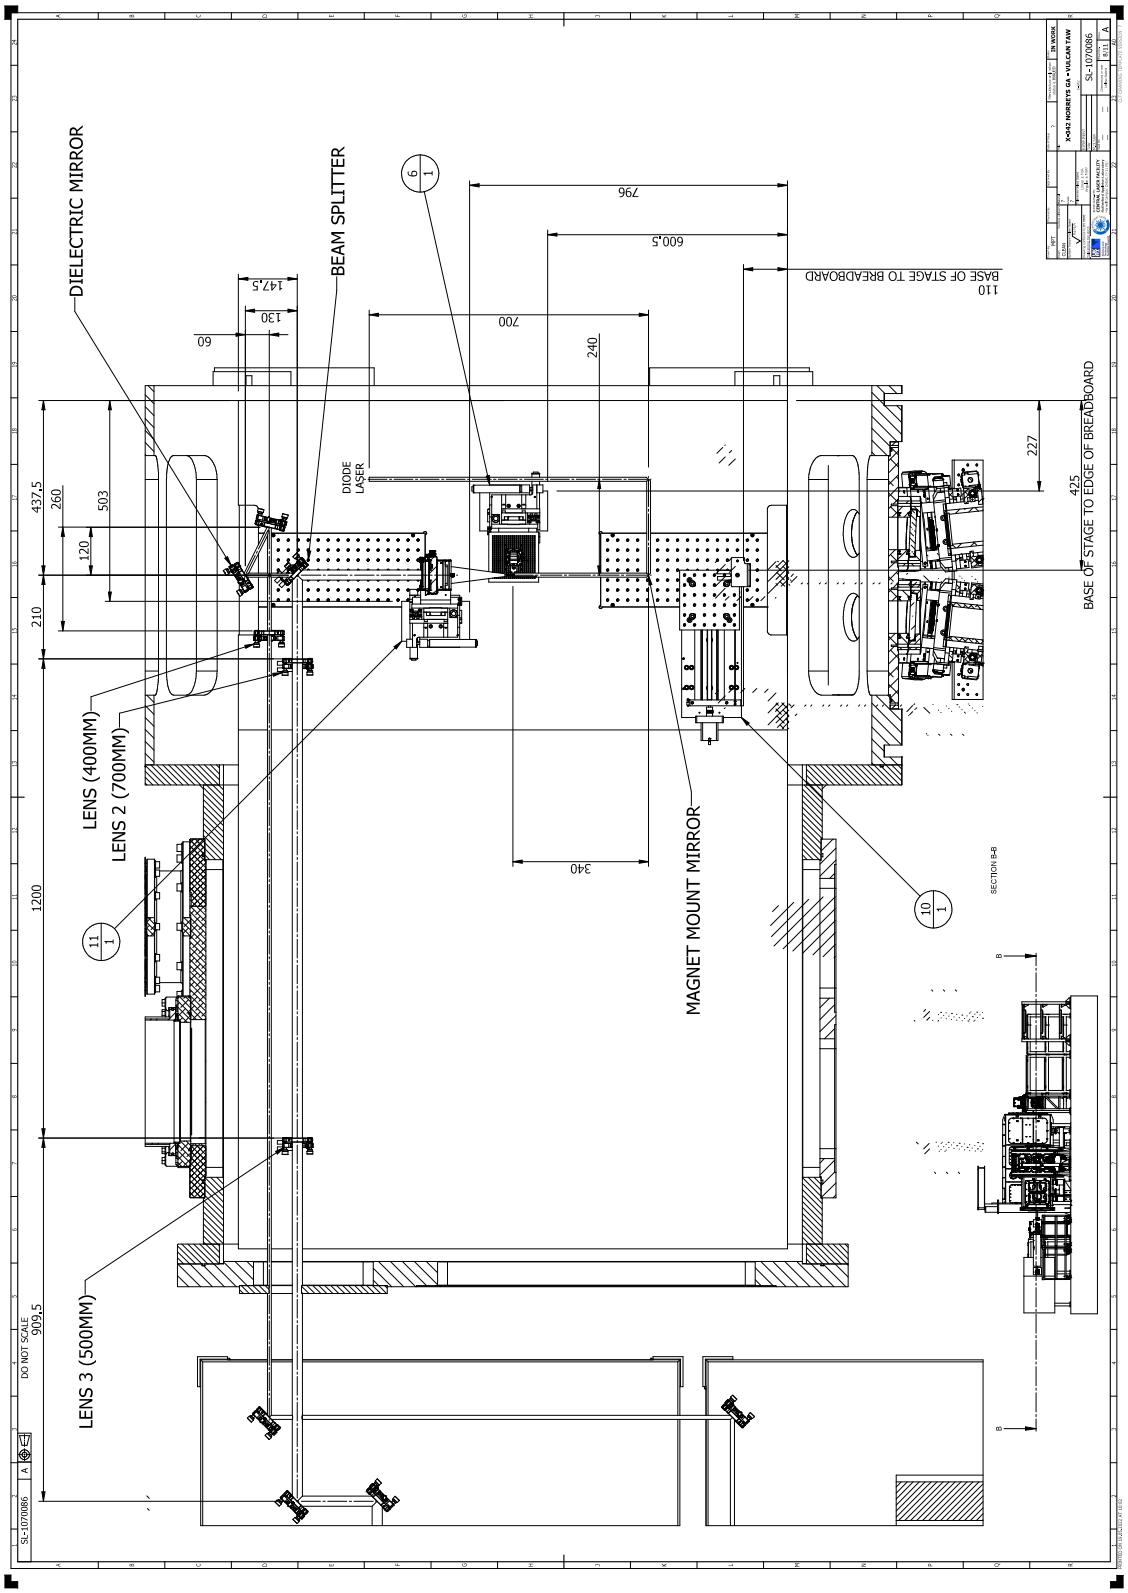


Figure C.11: Engineering drawing showing the target chamber for the final set-up, with full measurements.

C.4 Target images

The following images show the targets used in this experiment. With the exception of Figure C.12, these were provided by Chris Spindloe and CLF target fabrication.



Figure C.12: A photograph of the target. Half of the front surface is covered by foam (white), while the other half is quartz (as the quartz is transparent, this half appears gold due to the gold layer underneath).

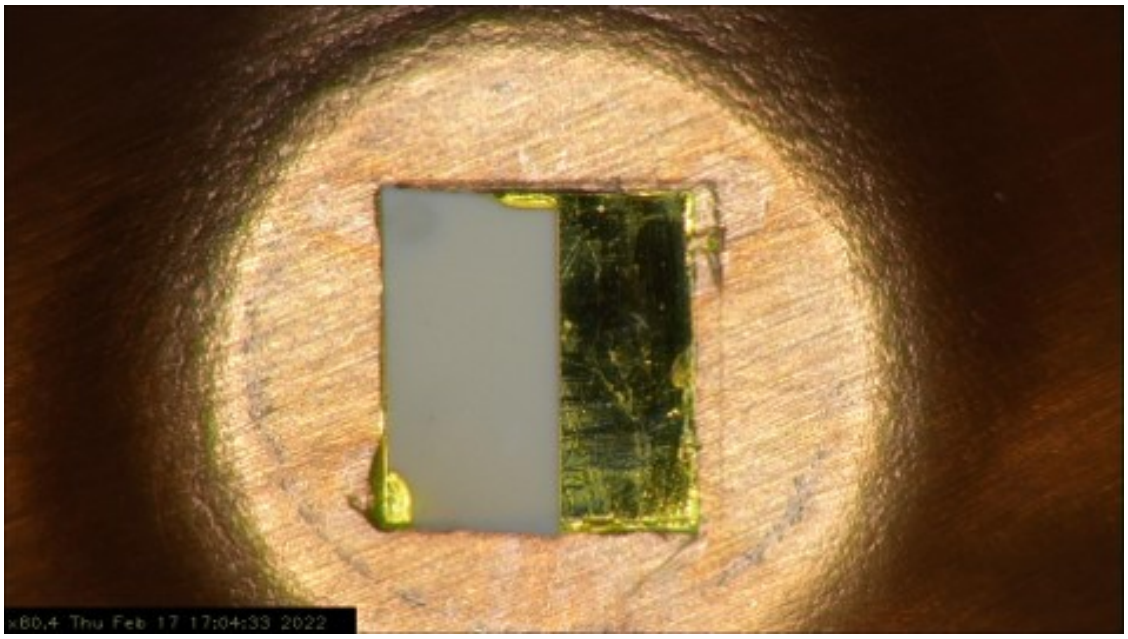


Figure C.13: Microscope image of a target, showing the step between the foam and quartz.

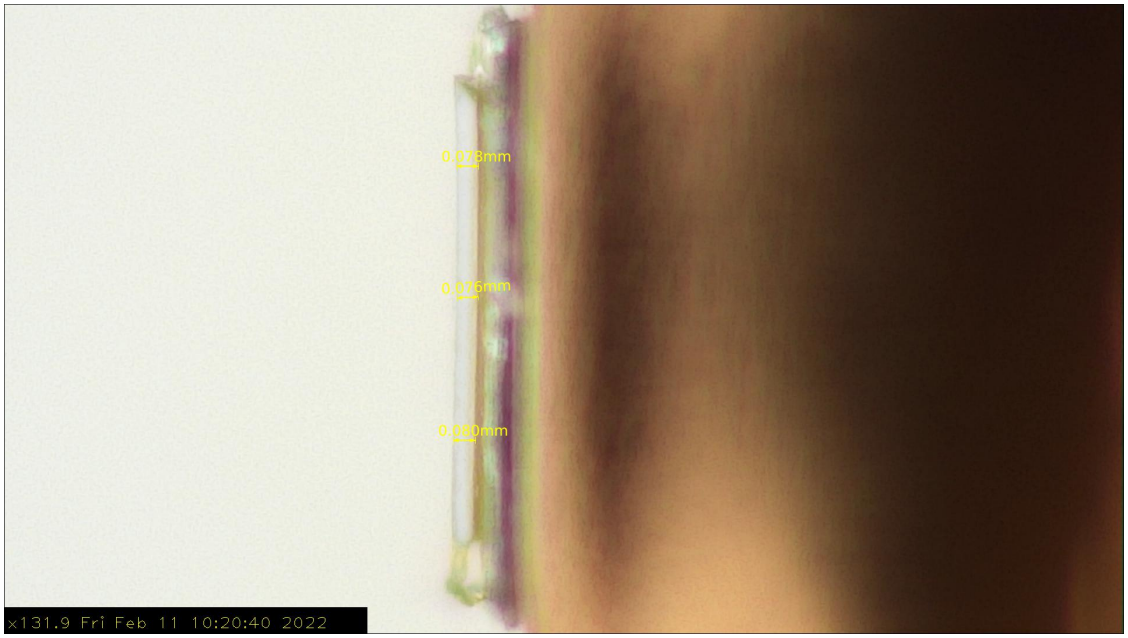


Figure C.14: Microscope image of the target, side on. The foam has been measured at three positions, providing the foam thickness used for the average velocity calculation. The variation between different points across a range of these images was used to estimate the uncertainty in this measurement.

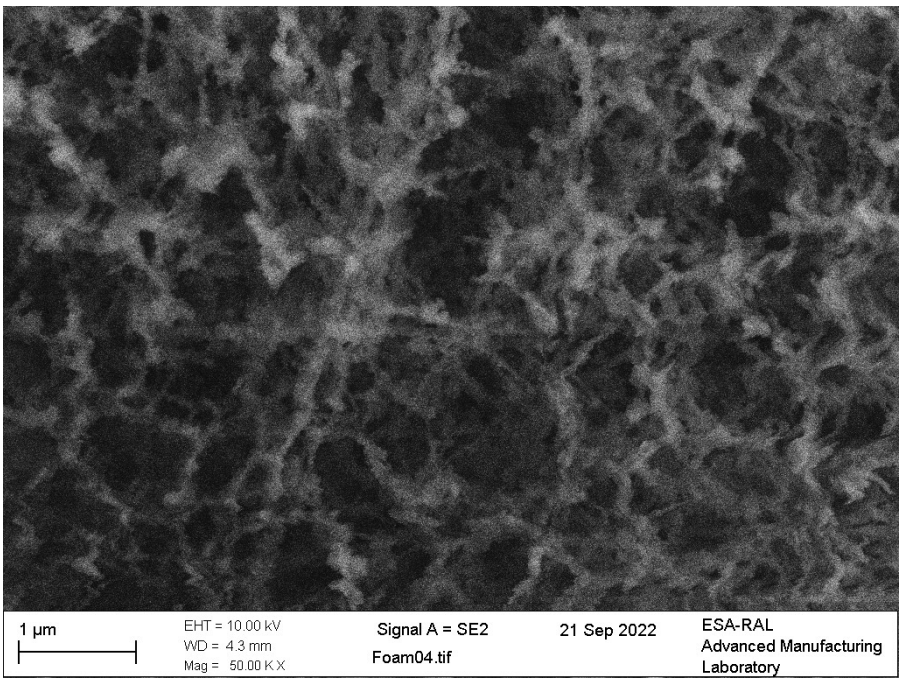


Figure C.15: A scanning electron microscope image of the foam material. Fibres of plastic polymer can be seen. A number of these were performed, and allowed the pore size to be estimated as roughly 1 µm.

References

[1] E. M. Campbell et al., "Laser-direct-drive program: Promise, challenge, and path forward", Matter and Radiation at Extremes **2**, 37–54 (2017).

[2] G. R. Tynan and A. Abdulla, "How might controlled fusion fit into the emerging low-carbon energy system of the mid-twenty-first century?", Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences **378**, 20200009 (2020),

[3] K. Gi, F. Sano, K. Akimoto, R. Hiwatari, and K. Tobita, "Potential contribution of fusion power generation to low-carbon development under the Paris Agreement and associated uncertainties", Energy Strategy Reviews **27**, 100432 (2020).

[4] J. H. Nuckolls, "Grand challenges of inertial fusion energy", Journal of Physics: Conference Series **244**, 012007 (2010).

[5] V. N. Goncharov et al., "Novel Hot-Spot Ignition Designs for Inertial Confinement Fusion with Liquid-Deuterium-Tritium Spheres", Physical Review Letters **125**, 065001 (2020).

[6] R. E. Olson et al., "A polar direct drive liquid deuterium-tritium wetted foam target concept for inertial confinement fusion", Physics of Plasmas **28**, 122704 (2021),

[7] R. E. Olson et al., "First Liquid Layer Inertial Confinement Fusion Implosions at the National Ignition Facility", Physical Review Letters **117**, 245001 (2016).

[8] F. F. Chen, *Introduction to Plasma Physics and Controlled Fusion*, Third edition (Springer International Publishing, Cham, 2016).

[9] J. Colvin and J. Larsen, *Extreme Physics*, First edition (Cambridge University Press, Cambridge, 2013),

[10] A. Piel, *Plasma physics: an introduction to laboratory, space, and fusion plasmas*, Second edition, Graduate texts in physics (Springer, Cham, 2017).

[11] G. Belmont, L. Rezeau, C. Riconda, and A. Zaslavsky, *Introduction to Plasma Physics*, First edition (Elsevier, Oxford, 2019),

[12] J. T. Larsen and S. M. Lane, "HYADES-A plasma hydrodynamics code for dense plasma studies", Journal of Quantitative Spectroscopy and Radiative Transfer **51**, 179–186 (1994).

[13] R. M. More, K. H. Warren, D. A. Young, and G. B. Zimmerman, "A new quotidian equation of state (QEOS) for hot dense matter", Physics of Fluids **31**, 3059–3078 (1988),

[14] S. P. Lyon and J. Johnson, "SESAME tables", in *Los Alamos National Laboratory Report No. LA-UR-92-3407* (1995).

[15] S. Atzeni and J. Meyer-ter-Vehn, *The Physics of Inertial Fusion*, First edition, International series of monographs on physics (Oxford University Press, Oxford, 2004).

[16] 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**, 092705 (2013),

[17] J. D. Lawson, "Some criteria for a power producing thermonuclear reactor", Proceedings of the Physical Society. Section B **70**, 6–10 (1957).

[18] B. Cheng, T. J. Kwan, Y. M. Wang, and S. H. Batha, "On thermonuclear ignition criterion at the National Ignition Facility", Physics of Plasmas **21**, 102707 (2014),

[19] G. S. Fraley, E. J. Linnebur, R. J. Mason, and R. L. Morse, "Thermonuclear burn characteristics of compressed deuterium-tritium microspheres", Physics of Fluids **17**, 474–489 (1974),

[20] M. Tabak et al., "Ignition and high gain with ultrapowerful lasers", Physics of Plasmas **1**, 1626 (1994).

[21] M. Tabak et al., "Review of progress in Fast Ignition", Physics of Plasmas **12**, 057305 (2005),

[22] M. Tabak, D. Hinkel, S. Atzeni, E. M. Campbell, and K. Tanaka, "Fast ignition: Overview and background", Fusion Science and Technology **49**, 254–277 (2006),

[23] R. Kodama et al., "Fast heating scalable to laser fusion ignition", Nature **418**, 933 (2002),

[24] Y. Kitagawa et al., "Progress of fast ignitor studies and Petawatt laser construction at Osaka University", Physics of Plasmas **9**, 2202 (2002),

[25] P. Norreys et al., "Fast electron energy transport in solid density and compressed plasma", Nuclear Fusion **54**, 054004 (2014).

[26] R. Betti et al., "Shock ignition of thermonuclear fuel with high areal density", Physical Review Letters **98**, 155001 (2007).

[27] R. H. H. Scott et al., "Shock-Augmented Ignition Approach to Laser Inertial Fusion", Physical Review Letters **129**, 195001 (2022).

[28] S. Skupsky et al., "Polar direct drive on the national ignition facility", Physics of Plasmas **11**, 2763 (2004),

[29] R. S. Craxton et al., "Polar direct drive: Proof-of-principle experiments on OMEGA and prospects for ignition on the National Ignition Facility", Physics of Plasmas **12**, 056304 (2005),

[30] T. J. Collins and J. A. Marozas, "Mitigation of cross-beam energy transfer in ignition-scale polar-direct-drive target designs for the National Ignition Facility", Physics of Plasmas **25**, 072706 (2018),

[31] M. Hohenberger et al., "Polar-direct-drive experiments on the National Ignition Facility", Physics of Plasmas **22**, 056308 (2015),

[32] J. D. Lindl et al., "Review of the national ignition campaign 2009-2012", Physics of Plasmas **21**, 020501 (2014),

[33] O. A. Hurricane et al., "The high-foot implosion campaign on the National Ignition Facility", Physics of Plasmas **21**, 056314 (2014),

[34] V. A. Smalyuk et al., "Review of hydrodynamic instability experiments in inertially confined fusion implosions on National Ignition Facility", Plasma Physics and Controlled Fusion **62**, 014007 (2020),

[35] T. Ma et al., "Onset of hydrodynamic Mix in high-velocity, highly compressed inertial confinement fusion implosions", Physical Review Letters **111**, 085004 (2013).

[36] H. Takabe and T. Ishii, "Effect of nonuniform implosion on high-gain intertial confinement fusion targets", Japanese Journal of Applied Physics **32**, 5675 (1993).

[37] H. Takabe, K. Mima, L. Montierth, and R. L. Morse, "Self-consistent growth rate of the Rayleigh-Taylor instability in an ablatively accelerating plasma", Physics of Fluids **28**, 3676 (1985),

[38] R. Betti, V. N. Goncharov, R. L. McCrory, and C. P. Verdon, "Growth rates of the ablative Rayleigh-Taylor instability in inertial confinement fusion", Physics of Plasmas **5**, 1446 (1998),

[39] B. M. Haines et al., "Robustness to hydrodynamic instabilities in indirectly driven layered capsule implosions", Physics of Plasmas **26**, 012707 (2019),

[40] L. O. Silva and R. Bingham, "Theory of Underdense Laser-Plasma Interactions with Photon Kinetic Theory", in *Laser-plasma interactions and applications*, edited by P. McKenna, D. Neely, R. Bingham, and D. A. Jaroszynski, First edition, Scottish Graduate Series (Springer International Publishing, Heidelberg, 2013), pp. 3–18,

[41] R. S. Craxton et al., "Direct-drive inertial confinement fusion: A review", Physics of Plasmas **22**, 110501 (2015).

[42] M. J. Rosenberg et al., "Origins and Scaling of Hot-Electron Preheat in Ignition-Scale Direct-Drive Inertial Confinement Fusion Experiments", Physical Review Letters **120**, 055001 (2018).

[43] B. Yaakobi, C. Stoeckl, T. Boehly, D. D. Meyerhofer, and W. Seka, "Measurement of preheat due to fast electrons in laser implosions", Physics of Plasmas **7**, 3714 (2000).

[44] I. V. Igumenshchev et al., "Crossed-beam energy transfer in direct-drive implosions", Physics of Plasmas **19**, 056314 (2012),

[45] P. Michel et al., "Tuning the implosion symmetry of ICF targets via controlled crossed-beam energy transfer", Physical Review Letters **102**, 025004 (2009).

[46] R. K. Follett et al., "Suppressing Two-Plasmon Decay with Laser Frequency Detuning", Physical Review Letters **120**, 135005 (2018).

[47] E. M. Campbell et al., "Direct-drive laser fusion: Status, plans and future: Direct-Drive Laser Fusion", Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences **379**, 20200011 (2021).

[48] Y. B. Zel'dovich and Y. P. Raizer, *Physics of Shock Waves and High-Temperature Hydrodynamic Phenomena*, edited by W. D. Hayes and R. F. Probstein, Dover edition (Dover, New York, 2002).

[49] J. W. Forbes, *Shock wave compression of condensed matter: A Primer*, First edition (Springer Berlin Heidelberg, 2012).

[50] L. Davison, *Fundamentals of Shock Wave Propagation in Solids*, First edition, Shock Wave and High Pressure Phenomena (Springer, Berlin Heidelberg, 2008).

[51] A. B. Zylstra et al., "Variable convergence liquid layer implosions on the National Ignition Facility", Physics of Plasmas **25**, 056304 (2018),

[52] B. M. Haines et al., "The effects of convergence ratio on the implosion behavior of DT layered inertial confinement fusion capsules", Physics of Plasmas **24**, 072709 (2017),

[53] J. D. Lindl et al., "The physics basis for ignition using indirect-drive targets on the National Ignition Facility", Physics of Plasmas **11**, 339–491 (2004),

[54] V. N. Goncharov, Personal Communication, Sept. 2020.

[55] R. Betti and C. Zhou, "High-density and high-ρR fuel assembly for fast-ignition inertial confinement fusion", Physics of Plasmas **12**, 110702 (2005).

[56] H. Abu-Shawareb, "Lawson Criterion for Ignition Exceeded in an Inertial Fusion Experiment", Physical Review Letters **129**, 075001 (2022).

[57] P. B. Radha et al., "Triple-picket warm plastic-shell implosions on OMEGA", Physics of Plasmas **18**, 012705 (2011),

[58] 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**, 051001 (2011),

[59] H. S. Park et al., "High-adiabat high-foot inertial confinement fusion implosion experiments on the national ignition facility", Physical Review Letters **112**, 055001 (2014).

[60] B. M. Haines, Personal Communication, 2023.

[61] B. M. Haines et al., "A mechanism for reduced compression in indirectly driven layered capsule implosions", Physics of Plasmas **29**, 042704 (2022),

[62] National Research Council (NRC) Committee for the Review of the Department of Energy's Inertial Confinement Fusion Program (The National Academies Press., Washington DC, 1997).

[63] A. R. Christopherson et al., "Theory of ignition and burn propagation in inertial fusion implosions", Physics of Plasmas **27**, 052708 (2020),

[64] Y. Kato et al., "Implosion experiments with uniformity-improved GEKKO XII: Overview", AIP Conference Proceedings **369**, 101–107 (1996),

[65] H. Nishimura et al., "Indirect-direct hybrid target experiments with the GEKKO XII laser", Nuclear Fusion **40**, 547 (2000).

[66] D. D. Meyerhofer et al., "Core performance and mix in direct-drive spherical implosions with high uniformity", Physics of Plasmas **8**, 2251 (2001),

[67] C. K. Li et al., "Effects of fuel-shell mix upon direct-drive, spherical implosions on OMEGA", Physical Review Letters **89**, 165002 (2002).

[68] S. Le Pape et al., "Observation of a reflected shock in an indirectly driven spherical implosion at the National Ignition Facility", Physical Review Letters **112**, 225002 (2014).

[69] S. F. Khan et al., "Symmetry tuning of a near one-dimensional 2-shock platform for code validation at the National Ignition Facility", Physics of Plasmas **23**, 042708 (2016),

[70] M. H. Key et al., "Study of ablatively imploded spherical shells", Physical Review Letters **45**, 1801–1804 (1980).

[71] S. P. Regan et al., "Demonstration of Fuel Hot-Spot Pressure in Excess of 50 Gbar for Direct-Drive, Layered Deuterium-Tritium Implosions on OMEGA", Physical Review Letters **117**, 025001 (2016),

[72] C. A. Williams, R. Betti, V. Gopalaswamy, and A. Lees, "High yields in direct-drive inertial confinement fusion using thin-ice DT liner targets", Physics of Plasmas **28**, 122708 (2021),

[73] R. E. Olson et al., "Wetted foam liquid fuel ICF target experiments", Journal of Physics: Conference Series **717**, 012042 (2016).

[74] R. C. Shah et al., "Systematic Fuel Cavity Asymmetries in Directly Driven Inertial Confinement Fusion Implosions", Physical Review Letters **118**, 135001 (2017).

[75] M. Gatu Johnson et al., "Impact of stalk on directly driven inertial confinement fusion implosions", Physics of Plasmas **27**, 032704 (2020),

[76] A. B. Zylstra et al., "Experimental achievement and signatures of ignition at the National Ignition Facility", Physical Review E - Statistical, Nonlinear, and Soft Matter Physics **106**, 025202 (2022).

[77] D. S. Clark et al., "Capsule physics comparison of National Ignition Facility implosion designs using plastic, high density carbon, and beryllium ablators", Physics of Plasmas **25**, 032703 (2018),

[78] A. J. Mackinnon et al., "High-density carbon ablator experiments on the National Ignition Facility", Physics of Plasmas **21**, 056318 (2014),

[79] C. R. Weber et al., "Improving ICF implosion performance with alternative capsule supports", Physics of Plasmas **24**, 056302 (2017),

[80] B. A. Hammel et al., "A "polar contact" tent for reduced perturbation and improved performance of NIF ignition capsules", Physics of Plasmas **25**, 082714 (2018),

[81] C. R. Weber et al., "Mixing in ICF implosions on the National Ignition Facility caused by the fill-tube", Physics of Plasmas **27**, 032703 (2020),

[82] A. Pak et al., "Impact of Localized Radiative Loss on Inertial Confinement Fusion Implosions", Physical Review Letters **124**, 145001 (2020).

[83] J. Lindl, "Development of the indirect-drive approach to inertial confinement fusion and the target physics basis for ignition and gain", Physics of Plasmas **2**, 3933–4024 (1995),

[84] 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**, 1906–1918 (2003),

[85] B. Cheng et al., "Ignition and pusher adiabat", Plasma Physics and Controlled Fusion **60**, 074011 (2018),

[86] A. L. Kritcher et al., "Design of an inertial fusion experiment exceeding the Lawson criterion for ignition", Physical Review E - Statistical, Nonlinear, and Soft Matter Physics **106**, 025201 (2022).

[87] D. A. Callahan et al., "Higher velocity, high-foot implosions on the National Ignition Facility laser", Physics of Plasmas **22**, 056314 (2015).

[88] V. N. Goncharov, "Cryogenic Deuterium and Deuterium-Tritium Direct–Drive Implosions on Omega", in *Laser-plasma interactions and applications*, edited by P. McKenna, D. Neely, R. Bingham, and D. A. Jaroszynski, First edition, Scottish Graduate Series (Springer International Publishing, Heidelberg, 2013), pp. 135–183,

[89] V. N. Goncharov et al., "Improving the hot-spot pressure and demonstrating ignition hydrodynamic equivalence in cryogenic deuterium-tritium implosions on OMEGA", Physics of Plasmas **21**, 056315 (2014),

[90] R. Betti and O. A. Hurricane, "Inertial-Confinement fusion with lasers", Nature Physics **12**, 435 (2016),

[91] H. F. Robey et al., "The effect of laser pulse shape variations on the adiabat of NIF capsule implosions", Physics of Plasmas **20**, 052707 (2013),

[92] O. L. Landen et al., "Fuel convergence sensitivity in indirect drive implosions", Physics of Plasmas **28**, 042705 (2021).

[93] A. Lees et al., "Experimentally Inferred Fusion Yield Dependencies of OMEGA Inertial Confinement Fusion Implosions", Physical Review Letters **127**, 105001 (2021).

[94] D. S. Montgomery, "Two decades of progress in understanding and control of laser plasma instabilities in indirect drive inertial fusion", Physics of Plasmas **23**, 055601 (2016),

[95] A. A. Solodov et al., "Hot-electron preheat and mitigation in polar-direct-drive experiments at the National Ignition Facility", Physical Review E - Statistical, Nonlinear, and Soft Matter Physics **106**, 055204 (2022).

[96] 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**, 062706 (2019),

[97] R. Nora et al., "Theory of hydro-equivalent ignition for inertial fusion and its applications to OMEGA and the National Ignition Facility", Physics of Plasmas **21**, 056316 (2014),

[98] B. M. Haines and A. Zylstra, Personal Communication, June 2020.

[99] R. Betti et al., "Alpha Heating and Burning Plasmas in Inertial Confinement Fusion", Physical Review Letters **114**, 255003 (2015).

[100] R. Betti, *Tutorial on the Physics of Inertial Confinement Fusion*, 2011, URL: `https://fire.pppl.gov/IFE_NAS3_ICF_tutorial_Betti.pdf`.

[101] R. E. Olson, "Concept for increased neutron yield and potential ICF ignition at the NIF", in Proceedings of 62nd Annual Meeting APS Division of Plasma Physics, Talk GO09.00005, No. LA- UR-20-28870 (2020),

[102] 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**, 381 (2006).

[103] C. Weber et al., "Hydrodynamic instability measurements in DT-layered ICF capsules using the layered-HGR platform", Journal of Physics: Conference Series **717**, 012057 (2016).

[104] 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**, 20200031 (2020).

[105] J. D. Sethian et al., "The science and technologies for fusion energy with lasers and direct-drive targets", IEEE Transactions on Plasma Science **38**, 690–703 (2010).

[106] E. J. Divall et al., "Titania—a 10^{20} W cm^{-2} ultraviolet laser", Journal of Modern Optics **43**, 1025–1033 (1996),

[107] V. D. Zvorykin et al., "Physical and technological issues of KrF laser drivers for inertial fusion energy", Journal de Physique IV France **133**, 567–573 (2006).

[108] S. P. Obenschain et al., "The Nike KrF laser facility: Performance and initial target experiments", Physics of Plasmas **3**, 2098–2107 (1996).

[109] V. N. Goncharov et al., "Demonstrating ignition hydrodynamic equivalence in direct-drive cryogenic implosions on OMEGA", Journal of Physics: Conference Series **717**, 012008 (2016).

[110] A. R. Christopherson et al., "A comprehensive alpha-heating model for inertial confinement fusion", Physics of Plasmas **25**, 012703 (2018),

[111] D. M. Kehne et al., "Implementation of focal zooming on the Nike KrF laser", Review of Scientific Instruments **84**, 013509 (2013),

[112] D. Eimerl et al., "StarDriver: A flexible laser driver for inertial confinement fusion and high energy density physics", Journal of Fusion Energy **33**, 476–488 (2014).

[113] S. Obenschain et al., "High-energy krypton fluoride lasers for inertial fusion", Applied Optics **54**, F103 (2015),

[114] N. Ratan et al., "Dense plasma heating by crossing relativistic electron beams", Physical Review E - Statistical, Nonlinear, and Soft Matter Physics **95**, 013211 (2017).

[115] J. J. Lee et al., "Toward more robust ignition of inertial fusion targets", Physics of Plasmas **30**, 022702 (2023),

[116] A. J. Kemp et al., "Laser-plasma interactions for fast ignition", Nuclear Fusion **54**, 054002 (2014).

[117] S. C. Wilks and W. L. Kruer, "Absorption of ultrashort, ultra-intense laser light by solids and overdense plasmas", IEEE Journal of Quantum Electronics **33**, 1954 (1997).

[118] T. Ma et al., "Hot electron temperature and coupling efficiency scaling with prepulse for cone-guided fast ignition", Physical Review Letters **108**, 115004 (2012).

[119] A. J. Kemp, Y. Sentoku, and M. Tabak, "Hot-electron energy coupling in ultraintense laser-matter interaction", Physical Review E - Statistical, Nonlinear, and Soft Matter Physics **79**, 066406 (2009).

[120] 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**, 072711 (2012),

[121] J. Tonge et al., "A simulation study of fast ignition with ultrahigh intensity lasers", Physics of Plasmas **16**, 056311 (2009),

[122] P. A. Norreys et al., "Preparations for a European R&D roadmap for an inertial fusion demo reactor", Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences **379**, 20200005 (2021).

[123] A. B. Zylstra et al., "Record Energetics for an Inertial Fusion Implosion at NIF", Physical Review Letters **126**, 025001 (2021).

[124] B. Cheng, P. A. Bradley, S. M. Finnegan, and C. A. Thomas, "Fundamental factors affecting thermonuclear ignition", Nuclear Fusion **61**, 096010 (2021),

[125] D. S. Clark and M. Tabak, "A self-similar isochoric implosion for fast ignition", Nuclear Fusion **47**, 1147–1156 (2007),

[126] R. K. Kirkwood et al., "Plasma-based beam combiner for very high fluence and energy", Nature Physics **14**, 80–84 (2018),

[127] R. K. Kirkwood et al., "A plasma amplifier to combine multiple beams at NIF", Physics of Plasmas **25**, 056701 (2018),

[128] R. K. Kirkwood, Personal Communication, Sept. 2021.

[129] P. Nicolaï et al., "Experimental evidence of foam homogenization", Physics of Plasmas **19**, 113105 (2012).

[130] S. Depierreux et al., "Laser smoothing and imprint reduction with a foam layer in the multikilojoule regime", Physical Review Letters **102**, 195005 (2009).

[131] 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**, 082710 (2018),

[132] S. E. Bodner et al., "Direct-drive laser fusion: Status and prospects", Physics of Plasmas **5**, 1901 (1998),

[133] T. J. Collins, A. Poludnenko, A. Cunningham, and A. Frank, "Shock propagation in deuterium-tritium-saturated foam", Physics of Plasmas **12**, 062705 (2005),

[134] W. Shang, R. Yu, W. Zhang, and J. Yang, "Optimization of x-ray emission from under-critical CH foam coated gold targets by laser irradiation", Nuclear Fusion **56**, 086002 (2016).

[135] G. Hazak, A. L. Velikovich, J. H. Gardner, and J. P. Dahlburg, "Shock propagation in a low-density foam filled with fluid", Physics of Plasmas **5**, 4357 (1998),

[136] 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**, e40 (2021),

[137] M. Cipriani et al., "Laser-supported hydrothermal wave in low-dense porous substance", Laser and Particle Beams **36**, 121–128 (2018),

[138] K. Falk et al., "Temperature measurements of shocked silica aerogel foam", Physical Review E - Statistical, Nonlinear, and Soft Matter Physics **90**, 033107 (2014).

[139] K. Falk et al., "Experimental observation of elevated heating in dynamically compressed CH foam", Plasma Physics and Controlled Fusion **62**, 074001 (2020),

[140] I. N. Ross et al., "Vulcan—A Versatile High-Power Glass Laser for Multiuser Experiments", IEEE Journal of Quantum Electronics **17**, 1653–1661 (1981).

[141] E. Springate, *Laser amplification*, 2010, URL: `https://www.clf.stfc.ac.uk/Pages/Laser-amplification.aspx`.

[142] D. Strickland and G. Mourou, "Compression of amplified chirped optical pulses", Optics Communications **56**, 219–221 (1985).

[143] I. N. Ross, P. Matousek, M. Towrie, A. J. Langley, and J. L. Collier, "The prospects for ultrashort pulse duration and ultrahigh intensity using optical parametric chirped pulse amplifiers", Optics Communications **144**, 125–133 (1997).

[144] C. Hernandez-Gomez et al., "Vulcan petawatt-operation and development", Journal de Physique IV France **133**, 555 (2006).

[145] I. Musgrave et al., "Review of laser diagnostics at the Vulcan laser facility", High Power Laser Science and Engineering **3**, e26 (2015).

[146] C. N. Danson et al., "Vulcan petawatt - An ultra-high-intensity interaction facility", Nuclear Fusion **44**, S239 (2004).

[147] L. M. Barker, "The development of the VISAR, and its use in shock compression science", AIP Conference Proceedings **505**, 11 (2000),

[148] L. M. Barker and R. E. Hollenbach, "Laser interferometer for measuring high velocities of any reflecting surface", Journal of Applied Physics **43**, 4669 (1972).

[149] L. M. Barker, "Fine Structure of Compressive and Release Wave Shapes in Aluminium Measured by the Velocity Interferometer Technique", in *Behaviour of dense media under high dynamic pressures* (Gordon and Breach, New York, 1968), p.483,

[150] B. D. Hammel, *Line-VISAR Theory*, 2018, URL: `http://www.bdhammel.com/line-visar-theory/`.

[151] R. Clifton, "Analysis of the laser velocity interferometer", Journal of Applied Physics **41**, 5335 (1970).

[152] D. R. Goosman, "Analysis of the laser velocity interferometer", Journal of Applied Physics **46**, 3516–3524 (1975).

[153] J. R. J. Bennett, *The VISAR. A Simple Analysis* (Science and Technology Facilities Council Rutherford Appleton Laboratory, Harwell Oxford, 2016), URL: https://epubs.stfc.ac.uk/manifestation/29837366/RAL-TR-2016-013.pdf.

[154] E. Hecht, *Optics*, Fifth edition (Pearson, Harlow, England, 2017).

[155] L. M. Barker and K. W. Schuler, "Correction to the velocity-per-fringe relationship for the VISAR interferometer", Journal of Applied Physics **45**, 3692–3693 (1974).

[156] L. M. Barker and R. E. Hollenbach, "Shock-wave studies of PMMA, fused silica, and sapphire", Journal of Applied Physics **41**, 4208–4226 (1970).

[157] P. M. Celliers et al., "Line-imaging velocimeter for shock diagnostics at the OMEGA laser facility", Review of Scientific Instruments **75**, 4916–4929 (2004),

[158] C. A. Bolme and K. J. Ramos, "Line-imaging velocimetry for observing spatially heterogeneous mechanical and chemical responses in plastic bonded explosives during impact", Review of Scientific Instruments **84**, 083903 (2013),

[159] D. G. Hicks et al., "Dissociation of liquid silica at high pressures and temperatures", Physical Review Letters **97**, 025502 (2006).

[160] J. E. Miller et al., "Streaked optical pyrometer system for laser-driven shock-wave experiments on OMEGA", Review of Scientific Instruments **78**, 034903 (2007),

[161] M. C. Gregor et al., "Absolute calibration of the OMEGA streaked optical pyrometer for temperature measurements of compressed materials", Review of Scientific Instruments **87**, 114903 (2016),

[162] M. Millot et al., "Shock compression of stishovite and melting of silica at planetary interior conditions", Science **347**, 418–420 (2015),

[163] T. Bonnet et al., "Response functions of imaging plates to photons, electrons and 4He particles", Review of Scientific Instruments **84**, 103510 (2013),

[164] P. R. Bolton et al., "Instrumentation for diagnostics and control of laser-accelerated proton (ion) beams", Physica Medica **30**, 255–270 (2014).

[165] M. D. Knudson and M. P. Desjarlais, "Shock compression of quartz to 1.6 TPa: Redefining a pressure standard", Physical Review Letters **103**, 225501 (2009).

[166] 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 - Condensed Matter and Materials Physics **88**, 184107 (2013).

[167] D. G. Hicks et al., "High-precision measurements of the diamond Hugoniot in and above the melt region", Physical Review B - Condensed Matter and Materials Physics **78**, 174102 (2008).

[168] D. G. Hicks et al., "Shock compression of quartz in the high-pressure fluid regime", Physics of Plasmas **12**, 082702 (2005),

[169] R. F. Trunin, "Shock compressibility of condensed materials in strong shock waves generated by underground nuclear explosions", Physics - Uspekhi **37**, 1123 (1994),

[170] S. Root, R. J. Magyar, J. H. Carpenter, D. L. Hanson, and T. R. Mattsson, "Shock compression of a fifth period element: Liquid xenon to 840 GPa", Physical Review Letters **105**, 085501 (2010).

[171] S. Root et al., "Shock compression of hydrocarbon foam to 200 GPa: Experiments, atomistic simulations, and mesoscale hydrodynamic modeling", Journal of Applied Physics **114**, 103502 (2013),

[172] C. A. McCoy et al., "Shock-wave equation-of-state measurements in fused silica up to 1600 GPa", Journal of Applied Physics **119**, 215901 (2016),

[173] G. A. Holton, *Value-at-risk: theory and practice*, Second edition (Academic Press, London, 2003).

[174] S. X. Hu, T. R. Boehly, and L. A. Collins, "Properties of warm dense polystyrene plasmas along the principal Hugoniot", Physical Review E - Statistical, Nonlinear, and Soft Matter Physics **89**, 063104 (2014).

[175] S. X. Hu, L. A. Collins, J. P. Colgan, V. N. Goncharov, and D. P. Kilcrease, "Optical properties of highly compressed polystyrene: An ab initio study", Physical Review B - Condensed Matter and Materials Physics **96**, 144203 (2017).

[176] S. P. Lyon and J. Johnson, "SESAME 7592", in *Los Alamos National Laboratory Report No. LA-UR-92-3407* (1995).

[177] S. P. Marsh, ed., *LASL Shock Hugoniot Data*, Los Alamos series on dynamic material properties (University of California Press, Berkeley/Los Angeles, California, 1980), p. 150,

[178] Y. Aglitskiy et al., "Absolute Hugoniot measurements for CH foams in the 2-9 Mbar range", Physics of Plasmas **25**, 032705 (2018),

[179] M. Koenig et al., "Equation of state data experiments for plastic foams using smoothed laser beams", Physics of Plasmas **6**, 3296 (1999),

[180] S. Le Pape et al., "Density measurement of shock compressed foam using two-dimensional x-ray radiography", Review of Scientific Instruments **79**, 106104 (2008),

[181] T. R. Desjardins, C. A. Di Stefano, E. C. Merritt, K. A. Flippo, and F. W. Doss, "Reducing direct drive preheat with dopants", High Energy Density Physics **39**, 100937 (2021).

[182] A. B. Zylstra et al., "Diffusion-dominated mixing in moderate convergence implosions", Physical Review E - Statistical, Nonlinear, and Soft Matter Physics **97**, 061201 (2018),

[183] C. Bellei et al., "Species separation and kinetic effects in collisional plasma shocks", Physics of Plasmas **21**, 056310 (2014),

[184] D. T. Casey et al., "Evidence for stratification of deuterium-tritium fuel in inertial confinement fusion implosions", Physical Review Letters **108**, 075002 (2012).

[185] J. A. Marozas, "NIF Polar-Drive High DT-Yield Exploder-Pusher Designs Modeled Using Pump Depletion in DRACO", in Proceedings of 61st Annual Meeting APS Division of Plasma Physics, Talk YO5.00010 (2019),

[186] A. Gerrard and G. M. Burch, *Introduction to Matrix Methods in Optics.* First edition (Wiley, New York, 1975),

[187] M. Takeda, H. Ina, and S. Kobayashi, "Fourier-transform method of fringe-pattern analysis for computer-based topography and interferometry", Journal of the Optical Society of America **72**, 156–160 (1982),

[188] B. D. Hammel, *Line-VISAR Analysis*, 2017, URL: http://www.bdhammel.com/line-visar-analysis/.