Indirectly Driven High Yield Implosions and Phase Contrast Imaging the N+1 Shock

Iain Calum Freeman

Doctor of Philosophy.

University of York

Physics, Engineering and Technology

September 2025

Abstract

Implosion designs are presented which offer the potential to significantly increase the yield of indirect drive implosions relative to the N210808 National Ignition Facility (NIF) shot, which is used as a baseline. One-dimension (1D) radiation hydrodynamic simulations are presented in which the proposed modification increases the yield by a factor of 10, this is achieved by tuning the N+1 shock to reduce the shell adiabat and as a result increase areal density by a factor of 1.5. N+1 tuning alone increases the yield in simulations from 1.2MJ to 9.4MJ, combining N+1 tuning and the shock augmented ignition scheme gives a yield of 12MJ. Although the simulations presented are 1D, they also show equivalent or reduced in-flight aspect ratio (IFAR) and implosion velocity, indicating reduced susceptibility to 2D Rayleigh-Taylor instabilities relative to the base design of the N210808 NIF shot. Further simulations with alpha heating turned off also show improved areal density and IFAR. A pulse shaping sensitivity study is also presented indicating that this design is robust against the typical pulse shaping errors of a NIF-like laser system.

A synthetic diagnostic code (GREENER) is also presented, it was written to simulate X-ray phase contrast imaging (XPCI), a particularly useful diagnostic for examining shocks and interfaces. The synthetic diagnostic code was developed within the GEANT4 framework and is a Monte-Carlo based raytracer. It is shown to be able to produce synthetic images of 1D, 2D and 3D hydrodynamic simulations; importing complex geometries while preserving the required surface angles to allow refraction to be modelled accurately.

Finally, these threads are brought together by using the presented synthetic diagnostic code to design an experiment to image the N+1 shock on an indirect drive implosion, this experiment would verify the underpinning physics of the N+1 shock and pave the way for applying the N+1 tuning scheme to ignition scale experiments.

Contents

1 Introduction **18**

1.1 Motivation 18

1.2 Fusion Overview 21

1.3 Inertial Fusion Energy 25

2 Background **27**

2.1 Hydrodynamics and Shocks 27

2.1.1 Shock formation 28

2.1.2 Jump Conditions 29

2.1.3 Shock Hugoniots 30

2.1.4 Equation of State and Maximum Single Shock Compression 31

2.1.5 Strong Shocks Increase Entropy 32

2.1.6 Rarefactions 33

2.1.7 Shocks and Rarefactions at Interfaces 33

2.1.8 Mutli-Dimensional Instabilities 34

2.2 Implosion Design 35

2.2.1 Ablation, Direct Drive and Indirect Drive 35

2.2.2 Central Hotspot Ignition 38

2.2.3 Hydrodynamic Instabilities 39

2.2.4 Fusion Metrics 40

2.2.5 Shock Design 41

2.2.6 Fast Ignition 42

2.2.7 Shock Ignition 42

2.2.8 Shock Augmented Ignition 43

2.3 Xray Imaging 43
2.3.1 Complex Refractive Index 44
2.3.2 Absorption Imaging - Theory 45
2.3.3 Absorption Imaging - Applications in ICF 45
2.3.4 Phase Contrast Imaging - Theory 48
2.3.5 Phase Contrast Imaging - Applications in ICF 50
2.3.6 Self-emission Imaging 51
2.3.7 X-ray Sources 52
2.4 Shock Diagnostics: VISAR 53

3 Methods **54**
3.1 Imaging Simulation Codes 54
3.2 Experiment Simulation Codes 55
3.2.1 HYADES 57
3.3 Betatron Spectrum 59

4 GREENER **60**
4.1 Chapter Overview 60
4.2 Geometry Construction Methods 61
4.2.1 Voronoi Method 64
4.2.2 Isosurface Method 68
4.3 1D Problems: Spherical Capsules 69
4.4 2D Problem: Shocked Cylinder 71
4.5 3D Problem: Relativistic Transparency Experiment 76
4.6 Chapter Summary 80

5 N+1 Tuning Design **81**
5.1 Chapter Overview 81
5.2 Origin of the N+1 Shock: Rarefactions and Reflected Pressure Waves 82
5.3 Formation of the N+1 Shock 82
5.4 Justification for Tuning the N+1 Shock 84
5.5 Baseline Simulation 85
5.6 Peak Only Design 88

5.7 Timing Scan 91
5.8 Peak and Dip Design 94
5.9 N+1 Tuning and Shock Augmented Ignition 97
5.10 Alpha Heating-Off Simulations 98
5.11 Pulse Shape Sensitivity Study 99
5.12 Chapter Summary 103

6 N+1 Imaging Experiment Design 104
6.1 Chapter Overview 104
6.2 Implosion Design 105
6.3 Synthetic Images and a Figure of Merit 108
6.4 Interpolating Synthetic Images into a Time Series 111
6.5 Contrast Over Time 115
6.6 Chapter Summary 117

7 Conclusion 118

A Gain Requirements 135

B Tuning Reflected Shocks 137

C Convergence Testing GREENER 144

List of Tables

5.1 Table of implosion metrics, comparing simulations of N210808 before retuning the first 3 shocks, and after. Both simulations have the same flux limiters. ρR, T_{ion} and $P_{hotspot}$ are burn averaged. 87

5.2 Table of implosion metrics, directly comparing simulations with and without the N+1 Tuning modification. ρR, T_{ion} and $P_{hotspot}$ are burn averaged. . . . 91

5.3 Table of implosion metrics, directly comparing simulations with and without the N+1 Tuning modification. ρR, T_{ion} and $P_{hotspot}$ are burn averaged. . . . 99

5.4 Table of implosion metrics for alpha-off simulations, directly comparing simulations with and without the N+1 Tuning modification. ρR, T_{ion} and $P_{hotspot}$ are burn averaged. 100

List of Figures

1.1 A graph of binding energy per nucleon against number of nucleons in nucleus, with key isotopes labeled, note the scale is logarithmic up to mass number 50, and linear from then on [23]. 23
1.2 A graphs showing reaction cross section averaged over a Maxwell-Boltzmann distribution vs temperature for a few key fusion relevant reactions [24]. 24

2.1 An unperturbed 1D column of fluid (shown in orange), a plot of density against position is shown below the column of fluid, with a thin slice from x_1 to x_2 identified. 27
2.2 A 1D-like column of fluid with a pressure wave travelling through it from right to left, with a plot of the density against position shown below. 29
2.3 A 1D-like column of fluid with a shock travelling through it from right to left, with a plot of the density against position shown below. 29
2.4 Schematic showing a direct drive implosion (left) and an indirect drive implosion (right), with key features colour coded for clarity. In both only show a few laser beams, while in reality many more beams are required for even illumination of the capsule. 36
2.5 Diagram showing absorption imaging 46
2.6 Diagram showing refraction generating phase contrast effects 49

4.1 A diagram showing the different components of the GREENER code, starting with the output of a physics simulation, preprocessing it, using the geometry in the GREENER code, then post-processing the results. The authors contributions are shown in green, while modules written by others are shown in purple. 61

4.2 Synthetic image of a sphere (radius 0.5mm) constructed as a simple GEANT4 solid, backlit with fully columnated 4.5keV X-rays, simulated with 10 million rays. 62

4.3 Synthetic image of two spheres (radius 0.5mm) constructed from voxels, backlit with fully columnated 4.5keV X-rays, simulated with 10 million rays. Left has 20^3 grid points, right has 30^3 grid ponits 63

4.4 Diagram of a ray travelling through a refractive index gradient, showing how different discretisation grids give different refracted rays. 63

4.5 Two tangling mechansims in a Lagrangian mesh: Overlap, where a zone overlaps another zone, and bowtie, where the edges of a zone cross each other. . . 65

4.6 Voronoi diagram for a sphere, constructed in cylindrical coordinates. Blue points show the grid points used, black lines show the edges of each cell, and orange points show the grid points of the new mesh, which can be used to construct the geometry in GEANT4. Black dotted lines show edges of cells which are unbounded, these must either be closed or discarded. 66

4.7 Synthetic images of two spheres (radius 0.5mm) each constructed from a voronoi diagram generated from a Lagrangian grid, backlit with fully columnated 4.5keV X-rays, simulated with 10 million rays. Left has 6^2 grid points in the underlying 2D mesh, while right has 30^2. 67

4.8 A Voronoi diagram showing a grid of points experiencing a curved shocked from below. The blue dots represent the points on the mesh, with dashed lines showing the true curvature of surfaces and solid lines representing the voronoi surfaces. The central points of the shock have been deliberately spread in such a way that the voronoi surfaces become angled against the direction of curvature. 67

4.9 Synthetic images of two spheres (radius 0.5mm) each constructed from an isosurface generated from an Eulerian grid, backlit with fully columnated 4.5keV X-rays, simulated with 10 million rays. Left has 6^3 grid points, right has 30^3 grid ponits 69

4.10 A lineout of the capsule imaged in Dewald et al. [74], compared to a synthetic XPCI lineout produced by the GREENER code. Note that the synthetic code does not include the tungsten wire at the left edge of the experimental data. . 70

4.11 The weighted spectrum used to approximate the betatron source in Fourmaux et al. [121] using the method outlined in section 3.3. 72
4.12 A comparison of the experimental image for capsule 1 in Fourmaux et al.[121], to synthetic images produced by GREENER. The left synthetic image has refractive effects enabled, while the right synthetic image is an absorption only image. 72
4.13 A comparison of the lineout for capsule 1 in Fourmaux et al.[121], to synthetic lineouts produced by GREENER. 73
4.14 Schematic of a cross section of the cylindrical target used, showing the direction of laser incidence. 73
4.15 The experimental image produced on shot 69, with key features labeled. . . . 74
4.16 A synthetic image of a simulation of the experiment (left), and the 2D DUED density profile used to produce the synthetic image (right). 75
4.17 The experimental image produced on shot 69 (left) compared to the synthetic image produced from a simulation (right). 75
4.18 Simulations of proposed relativistic transparency experiment, for 3 possible densities, at timesteps once structures are well developed. 77
4.19 The weighted spectrum used to approximate the betatron source expected for the relativistic transparency experiment 77
4.20 Absorption only synthetic X-ray image of an epoch simulation of the relativistic transparency experiment, no features can be seen without XPCI. 78
4.21 Synthetic X-ray phase contrast images of the relativistic transparency experiment, showing different possible source offset distances. 79
4.22 Synthetic X-ray phase contrast image of the relativistic transparency experiment, with a source offset of 1m, and artificially high source brightness. . . . 80

5.1 A Lagrangian shock plot with the main features, shocks and rarefactions labeled. This shows the first three shocks as they propagate through the ablator and ice into the gas, as well as the rarefactions they create, then the reflected shocks which form from those rarefactions, until finally the reflected shocks merge to form the N+1 shock. 84

5.2 A plot of density, pressure and adiabat against radius in the lab frame, for the N210808 shot at a time chosen to show the N+1 shock travelling through the ice layer (t=6.4ns), where the increase in local adiabat due to the N+1 shock is clear. 85

5.3 A Lagrangian shock plot, with shocks shown in black, combined with a Lagrangian frame plot of the cell adiabat, shown as spectral colours. This shows the significant increase in adiabat precipitated by the N+1 shock merger . . . 86

5.4 A convergence test for the number of radiation groups to be used in hyades simulations, with the hotspot ion temperature at bang time, maximum areal density, and yield plotted for each number of radiation groups up to 70, after 70 radiation groups, only 84 and 140 groups are plotted. 88

5.5 The experimental data from N210808 for the radiation temperature and M-band fraction, shown along side the radiation temperature profile and M-band fraction used as the baseline for the simulations. 89

5.6 A Lagrangian shock plot for the result found by constrained bayesian optimisation, showing N+1 Tuning has been effective, as the N+1 shock mergers have been delayed until closer to the ice-gas iterface than in the baseline case. 90

5.7 A plot of yield vs time of peak maximum 92

5.8 A Lagrangian shock plot for a peak midpoint time (4.35ns) which is too early to tune the N+1 shock. The peak strengthens the first shock reflection too much (resulting in an early merger at 1017μm in the lagrangian frame). . . . 93

5.9 A Lagrangian shock plot for a peak midpoint time (4.95ns) which is relatively well timed to tune the N+1 shock. All 3 reflected shocks hit the ice-gas interface almost sequentially. 93

5.10 A Lagrangian shock plot for a peak midpoint time (5.55ns) which is too late to tune the N+1 shock. The early reflected shocks are able to merge with the later reflected shock, forming a strong shock in the ice layer (with shock mergers at 1017μm and 999μm in the lagrangian frame). 94

5.11 A Lagrangian shock plot for the result found by constrained optimisation of a peak and dip intorduced to the radiation temperature profile. Showing a slight improvement to the timing of the reflected shocks over a peak only design. . . 95

5.12 Diagram showing the 3 variables used to parameterise the peak-dip pulse shape. The baseline is shown in blue, with the modification shown in black, arrows show how the modification is changed with each variable. The pulse shape was smoothed before being passed to HYADES. 95
5.13 Heatmap showing the yield produced from HYADES simulations with different pulse shapes as a funciton of the slope gradient and midpoint, for a peak-to-dip time of 1ns. The left heatmap is without hohlraum constraints, and the right with hohlraum constraints. 96
5.14 Heatmap showing the yield produced from HYADES simulations with different pulse shapes as a funciton of the slope gradient and midpoint, for a peak-to-dip time of 0.8ns. The left heatmap is without hohlraum constraints, and the right with hohlraum constraints. 97
5.15 Diagram showing the many variables used to parameterise the peak-sai pulse shape. The baseline is shown in blue, with the modification shown in black, each variable used is labeled. The pulse shape was smoothed before being passed to HYADES. 97
5.16 A Lagrangian shock plot showing the results of combining N+1 Tuning and shock augmented ignition. 98
5.17 An example pulseshape, with a number of profiles with different errors applied 101
5.18 Yield(MJ) for different magnitudes of random error applied to the pulseshape 102
5.19 ρR for different magnitudes of random error applied to the pulseshape 102

6.1 A Lagrangian shock plot showing the 'flat-top' design, which is designed to make imaging the N+1 shock easier. The radiation temperature profile is also shown, and key features are labelled. 106
6.2 A Lagrangian shock plot showing the 'big-dip' design, which is designed to make imaging the N+1 shock easier. The radiation temperature profile is also shown, and key features are labelled. 107

6.3 Example synthetic XPCI lineout of the big-dip design at t=7.31ns into the implosion, the density used is also plotted for comparison. The ice-gas and ice-ablator interfaces are labeled, as is the N+1 shock. An absorption only lineout is also shown. The two shaded regions represent 3 standard deviations from their respective lineouts using a poission error. The spike and dip used in contrast calculations are identified with vertical green dotted lines. 109

6.4 A sequence of synthetic XPCI lineouts over time for the big-dip implosion(top left), paired with the shockplot of the big-dip implosion. The trajectories of the Lagrangian cells in the shockplot are used to interpolate the intermediate values of the synthetic XPCI lineouts, to generate the bottom panel. 112

6.5 A sequence of synthetic XPCI lineouts over time for the flat-top implosion(top left), paired with the shockplot of the flat-top implosion. The trajectories of the Lagrangian cells in the shockplot are used to interpolate the intermediate values of the synthetic XPCI lineouts, to generate the bottom panel. 113

6.6 Sequences of synthetic XPCI lineouts plotted over time for the big-dip design (top panels) and the flat-top design (bottom panels). Shown with a temporal blur which simulates the effect of a detector integration time, the left panels show a 50ps blur while the right panels show a 100ps blur. 114

6.7 Contrast-to-noise-ratio vs time for the selected time range on the big-dip design, showing the effect of a 50ps and 100ps blur. 116

6.8 Contrast-to-noise-ratio vs time for the selected time range on the flat-top design, showing the effect of a 50ps and 100ps blur. 116

B.1 Two lagrangian shockplots, showing a baseline (first panel) and modified (second panel) implosion design, with a lineout of density, electron temperature and pressure taken at the point marked in the plots. The lineouts are shown in the third panel, and taken at 4.38ns. 139

B.2 Two lagrangian shockplots, showing a baseline (first panel) and modified (second panel) implosion design, with a lineout of density, electron temperature and pressure taken at the point marked in the plots. The lineouts are shown in the third panel, and taken at 4.58ns. 140

B.3 Two lagrangian shockplots, showing a baseline (first panel) and modified (second panel) implosion design, with a lineout of density, electron temperature and pressure taken at the point marked in the plots. The lineouts are shown in the third panel, and taken at 4.62ns. 141

B.4 Two lagrangian shockplots, showing a baseline (first panel) and modified (second panel) implosion design, with a lineout of density, electron temperature and pressure taken at the point marked in the plots. The lineouts are shown in the third panel, and taken at 5.00ns. 142

B.5 Two lagrangian shockplots, showing a baseline (first panel) and modified (second panel) implosion design, with a lineout of density, electron temperature and pressure taken at the point marked in the plots. The lineouts are shown in the third panel, and taken at 5.10ns. 143

C.1 Density ramps of two different gradients (left panel), along with the same two density ramps each sampled at different spatial resolutions (right panel). . . . 145

C.2 XPCI lineouts for a density ramp of $0.05\frac{g/cm^3}{\mu m}$, with a range of spatial resolutions. Faint lines with the same colour as the lineout they correspond to show the density profiles used. The regions used for mean deviation calculations are indicated at the bottom of the plot 146

C.3 Mean percentage deviation at different spatial resolutions (step widths) for a density ramp of 0.05 $\frac{g/cm^3}{\mu m}$. 147

C.4 Convergence graph for 0.05 \$g/cm$^3\frac{}{\{\mu}$m\}, mean percentage deviation is plotted as a funciton of spatial resolution 148

C.5 Example synthetic XPCI lineout of the big-dip design at t=7.31ns into the implosion, the density used is also plotted at the resolution used in chapter 6, with different colours for different materials. An absorption only lineout is also shown. 149

C.6 Example synthetic XPCI lineout of the big-dip design at t=7.31ns into the implosion, the density used is also plotted at double the resolution used in chapter 6, with different colours for different materials. An absorption only lineout is also shown. 149

C.7 Convergence plot for the full range of density ramp gradients and spatial resolutions. The left panel has all the points for all the density profiles used in chapter 6 plotted at their respective density gradient and step width, the right panel excludes the doped ablator points as they are far from the region of interest. 150

Acknowedgments

I must, of course, thank my supervisors Robbie Scott and Nigel Woolsey, from whom I've learned a great deal. I would also like to thank the many academics in York, the CLF and further afield who have helped me along the way, in particular Luca Antonelli, Guy Burdiak, Kate Lancaster, Chris Murphy Rob Paddock, and Tony Bell.

For the rest of my acknowledgements: it's not often that you get to write a love letter to your friends, so I'll use this excuse to wax poetic.

Starting with my parents, who have held me through tears of joy and sorrow, through dreams and disappointments, who live out the hymn 'let love be real'. I owe them everything. The rest of my family have also been solid foundations, and Canada often feels like a haven from the world.

I'd like to thank various flatmates who brought me out of my shell: Matt, Luke, Christina and Ryan. I'm also grateful for the friend group that formed around the Fusion flat: you gave my life's gossip more attention than it deserved (but I loved it really). Everyone at the "Many Facets of Calum", "Barbie and Boardgames" or similar parties were wonderful and brought joy to my time in York. From my research group I'd particularly like to thank Phil, Matt, Adam and Arun; their wise (Phil and Matt) unwise (Adam) and downright foolish (Arun) advice was always appreciated.

I have been fortunate to know, and be known by, many communities, each filled with beautiful souls, people who have a sense of embodied wisdom and grace, as well as firebrands who will change the world for the better.

In Edinburgh I'm particularly grateful for Edith, Alistair and the Sawkins. In York I'd like to explicitly thank Victoria, Naomi, the CMC parents, Tina, Will, Naomi, Rebecca, Liz and Phoebe. I'd also like to thank Becca, who convinced me to move to Oxford while at Greenbelt. The WesleyMem Methodists and the Young Adult Quakers have felt like coming home; here I've found moments of intense joy, deep community, and people whose living room I could cry in. At WesleyMem I'd like to thank: Mary & John, Paul & Wendy, Freda & John, Heather, Kaden, and Mark; and among the Quakers: SJ (& Joe), Sian, Matt, Kit and Katie. All of these people made me feel welcome and understood in places where I was new. I'd also like to thank the Quakers up and down the country who make the Young Friends General Meeting (now YFiB) happen. There is hope, joy and love there; people are who they are and

it's glorious.

I'm also grateful to the many friends who knew me before the PhD began and continued to tend my soul throughout. Sean, my oldest friend, and my friends from undergrad: those who I've been to Hel and back with, GEAS, croc'n'roll, physics mystics, etc. There are too many to name, but I will mention Judith, Rosie, Marek, Patrick, Caelyb, Aisling, Laura, Will, Mary, Emma and Simeon.

Everyone mentioned, whether explicitly or as part of a group, has shown me some of the light in the core of every human heart. For that I'm grateful.

Declaration

I declare that this thesis is a presentation of original work and I am the sole author. This work has not previously been presented for an award at this, or any other, University. All sources are acknowledged as References.

Chapter 1

Introduction

1.1 Motivation

Any motivation for a course of action must begin in philosophy or morals, and so shall the motivation for this thesis. This section will start as close to first principles and moral justifications as the author dares, then will work its way down towards the justification for the concrete work presented in this thesis. At the end of this section the structure of the thesis will be outlined and the impact of the work will be emphasised.

The reasons to say human beings matter vary depending on your philosophical perspective: Methodists would say all people are made in the image of the divine [1], the Quakers would reference the inner light [2], while the reader may have their own philosophy. Scientists may be particularly drawn to the idea that humanity is the only known case of the universe really looking at itself and asking deep questions. Humans pursue knowledge and wisdom in a way that nothing else does, and there's beauty worth protecting in that. Furthermore, humans matter whether they are near or far, poor or rich, good or bad [3]; and individual human lives matter, even if something will only harm a few humans it is worthwhile to prevent the harm. This may seem like an unnecessary set of axioms to state, however within the discussion around climate change and energy security there are those who disagree with some of the above paragraph, so it is stated for completeness.

Climate change is a severe threat to humankind as a whole [4], and to particular individual lives [5], this was known as early as 1970s [6]. Leaving aside geoengineering and carbon removal (which come with considerable risks and are often cost prohibitive [7]), the main approach to resolving climate change lies in the decarbonisation of all human activity. There

are many aspects of human life which require decarbonisation, however, the vast majority of carbon emissions (over 70%) come from energy usage [8, 9], which will be the focus of this motivation.

Simply reducing global energy usage is not an option: as it stands, 600 million people in the world do not have access to any electricity [10], and many have extremely limited access [11]. This makes refrigeration and transport of food and medicine difficult, limits access to clean water, and damages health directly through polluting cooking fuels. So to preserve and improve human lives additional supplies of clean energy must be found. In particular, sources of electricity are important due to how versatile it is and how clean it is at the point of use.

The main sources of electricity currently available to humans are: fossil fuels, geothermal energy, hydroelectricity, renewable energy (solar and wind) and nuclear[4]. Fossil fuels are clearly a problem as they are a significant source of carbon emissions, and carbon capture has limited potential to reduce emissions (see figure 4.4 in the IPCC 2023 report [4]). Geothermal energy and hydroelectricity are expensive, especially outside of the few geologically or geographically favourable locations [4]. In recent years the cost of solar and wind have reduced significantly [12], the cost of batteries for short term storage has also reduced significantly over time [12]; renewable energy is a significant promising solution to climate change and will likely make up the majority of cheap reductions in carbon emissions [4]. However, a grid which fully relies on renewable energy creates challenges due to the intermittent nature of solar and wind, particularly over the longer term. While batteries support short term fluctuations in energy supply, building enough capacity to supply countries through an especially dark and windless winter is prohibitively expensive. Having even a small amount of constant and consistent energy production reduces the difficulty of maintaining electricity supplies significantly, and currently the UK uses dispatchable methane gas power plants to manage high demand. One option for consistent energy production with low carbon emissions is nuclear energy. However, the fuel required for conventional nuclear power plants poses a nuclear weapons proliferation risk, and there are significant concerns around waste management. In addition, while the total deaths and deaths per kilowatthour of electricity from nuclear energy are lower than any other energy source (except solar) [12], the perception of meltdown risks poses significant challenges to building power plants, in part due to the ‘low probability, high consequence’ nature of the risk.

Fusion has the benefits of conventional nuclear energy, without many of the problems. The radioactive waste is the largest concern, and with careful reactor design the waste can be low level within 100 years [13]. Fusion is a low carbon energy source, the fuel is readily available and the raw inputs (deuterium/lithium) are not a proliferation concern. The power produced by fusion is also reliable and consistent, and could even be adjustable with intelligent inertial fusion energy reactor design (eg by adjusting the repetition rate); allowing it to complement intermittent renewable sources.

Even in the case where fusion energy is not developed in time to significantly mitigate climate change, it's development is still worthwhile. In particular, having a fuel source which can be extracted from simple seawater could reduce the motivation for wars over energy supplies. And looking to the longer term future of humanity, the high density nature of fusion fuel would be well suited to long haul space voyages. So there are a number of non-climate change related benefits that come from working fusion reactors.

Within Fusion, there are two main branches: Inertial Fusion Energy (IFE) and Magnetic Confinement Fusion (MCF), a brief overview of these are given in section 1.2, and more detail on IFE follows in section 1.3. MCF began much earlier than IFE, and is broadly more mature, however, IFE has produced ignition [14] while MCF has yet to do so, suggesting IFE may be the more promising route.

To advance IFE further, gain must be increased significantly (see appendix A), and more work is needed on detailed reactor design. Both of those goals require diagnostics to understand the implosion experiments in more depth. One key diagnostic is X-ray Phase Contrast Imaging (XPCI)[15], details are given in section 2.3, which has the potential to allow for in-flight imaging of capsules as they implode, and to allow for the imaging of shocks as they travel through the shell [16]. XPCI could allow for significant improvements in the understanding of implosion dynamics, potentially refining equation of state models, and allowing for improved experimental design which could lead to higher yield implosions. To use diagnostics effectively, synthetic diagnostics are often required. This is particularly true for XPCI, where results can be difficult to interpret without a model of what would be expected from a given experiment. Synthetic diagnostics are also useful for designing experiments, they can help optimise a measurement to obtain as much useful information from an experiment as possible.

In chapter 4 of this thesis, a synthetic diagnostic tool will be presented, which can generate

synthetic X-ray Phase Contrast Images in the high energy density (HED) regieme. This work has applications outside of IFE, and can be used to study HED shocks in many contexts, but the code was built for IFE applications.

In chapter 5, this thesis will propose a way to increase the yield of an IFE implosion by tuning the N+1 shock, whilst also improving a range of other fusion metrics. This chapter focuses on an indirect drive implosion, but with results which should be also be applicable for direct drive.

Finally, in chapter 6, the synthetic diagnostic tool will be used to design an experiment to image the N+1 shock, allowing understanding of this shock and it's impact on implosions to be understood further, and supporting the work in chapter 5.

1.2 Fusion Overview

Fusion energy research began as an outworking of the nuclear physics developments through the 1950's, atomic weapons had shown that fusion reactions could produce energy and many physicists were looking for peaceful applications of this science [17]. Initially, the most promising method of fusion appeared to be magnetic confinement fusion, which works by confining a plasma within magnetic fields. Many different confinement schemes were tried, but the tokamak emerged as the most promising, and research focused on that direction for the following years [18].

After the invention of lasers in 1960, a small amount of classified research was done into inertial confinement fusion, primarily for weapons purposes. The Halite/Centurion program in particular demonstrated that the principles of inertial confinement fusion were sound[19], however, it wasn't until 1972 that any of this research became public [20]. At this point, inertial fusion energy (IFE) research began in earnest, decades later than MCF. In 1983 the Joint European Tourus (JET), a tokamak, began operating [21]; a number of experiments ran over the 14 years until in 1997 it had completed an experiment which produced fusion power equal to 60-80% of the input energy [21]. National Ignition Facility (NIF) was finished in 2009[17], and 13 years later achieved the first experiment to produce more fusion energy than was used to drive the reaction [14]. Both NIF and JET use more electricity than the energy used to drive the reactions, however, these results are both steps on the path towards reactors which produce net electricity. Based on the potential generated by the culmination of decades

of publicly funded research, recent years have seen a number of fusion companies and startups attempting to design and build demonstration reactors, with the goal of commercialising the industry. This thesis will focus on ICF designs for a NIF like facility.

The principle behind both IFE and MCF, as well as fission energy, is embodied by the nuclear binding energy curve, shown in figure 1.1, where the binding energy per nucleon is shown for a range of isotopes of different elements. If nuclei are combined or split to form new nuclei, and the new nuclei have higher total binding energy, then there will be excess energy released by the act of binding the nucleons in this more energetically favourable state, this excess energy becomes the kinetic energy of the nuclei and has the potential to be extracted.

In the case of nuclear fission, large isotopes are broken into smaller, more energetically favourable ones, while in the case of fusion, small isotopes are combined into large more energetically favourable ones. In both cases the progression is towards iron, which is the most energetically favourable configuration. A particular reaction of note is the Deuterium-Tritium reaction:

$$D + T \rightarrow He(3.5MeV) + n(14.1MeV)$$

[22]

This reaction produces a large amount of energy due to how tightly bound Helium-4 (an alpha particle) is.

For a fusion reaction to happen between two nuclei, those nuclei must overcome the Coulomb barrier, since they're both positively charged. This means they can only fuse within collisions of sufficient energy, although quantum tunnelling allows for a small probability that nuclei can penetrate this barrier without having sufficient energy to classically overcome it. This is critical in that it allows for fusion reactions to occur in dense material, where there are large numbers of collisions, at lower temperatures than would be required without quantum tunnelling. The temperature of the nuclei affect the cross section of a reaction, and this can be plotted for many different potential fusion reactions. This is one way to determine which reaction is most useful for a reactor, if the maximum cross section is too low, or the temperature required to reach a high cross section is too high, the reaction is likely a poor choice for a reactor. The left panel in figure 1.2 shows the cross sections against temperature for a number of potential reactions, which makes it clear that Deuterium-Tritium is a good candidate.

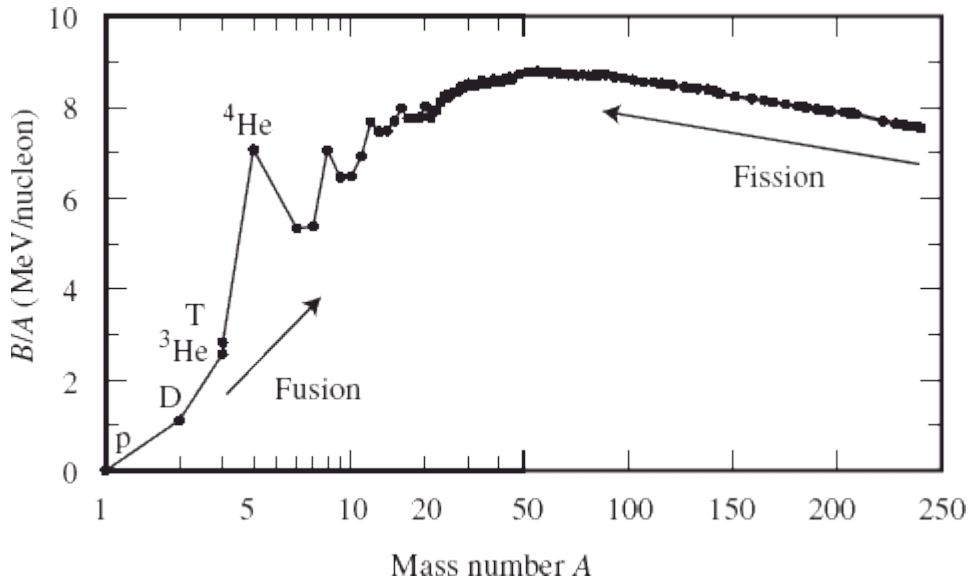


Figure 1.1: A graph of binding energy per nucleon against number of nucleons in nucleus, with key isotopes labeled, note the scale is logarithmic up to mass number 50, and linear from then on [23].

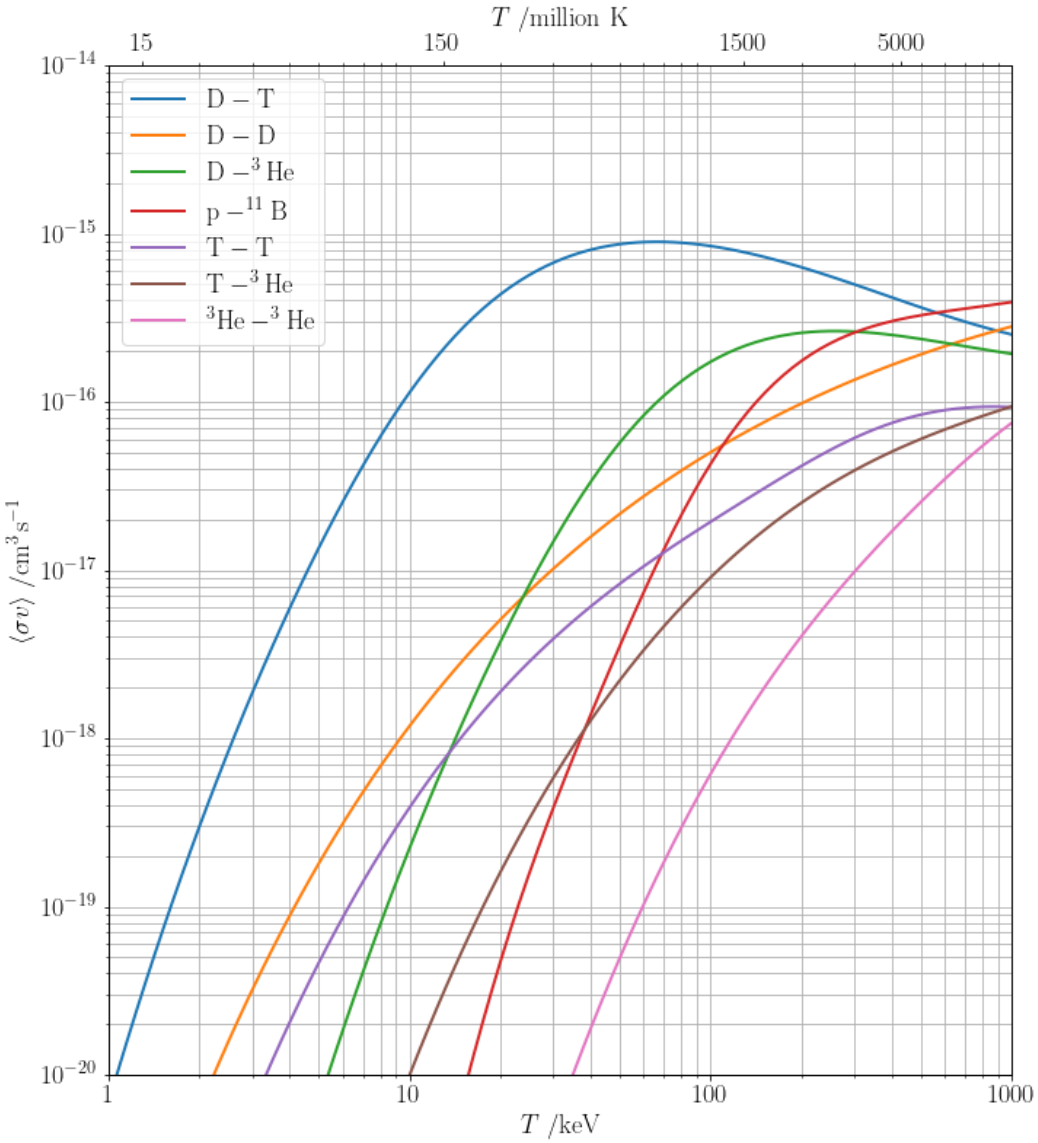


Figure 1.2: A graphs showing reaction cross section averaged over a Maxwell-Boltzmann distribution vs temperature for a few key fusion relevant reactions [24].

There are other factors involved in choosing a fuel, the availability of the fuel matters, and Deuterium can be extracted from heavy water, which can be filtered from sea water. Tritium is more difficult to obtain, but could be obtained by bombarding lithium with neutrons generated from a fusion reactor [25]. The neutron generation is potentially a problem, since it causes damage to wall materials which leads to radioactive waste. Other fuels which don't

produce neutrons may prove to be better in the long term, but often have lower reaction rates or require higher temperatures, which makes them more challenging to use. For now the field has largely settled on DT fusion.

1.3 Inertial Fusion Energy

The principle behind Inertial Fusion Energy (IFE) is to compress some fusion fuel to high enough energy densities that a fusion burn can begin, and then confine it there with the inertia of an imploding spherical shell for long enough to burn a significant portion of the fuel, therefore producing energy [26].

This is conventionally done using a hollow sphere (a capsule), the outer layer of the shell acts as an ablator, then the majority of the fuel forms an inner DT ice layer making up the rest of the shell, with a small amount of low density DT gas filling the hollow cavity within the shell. Radiation impinges on the capsule and ablates the surface away, as the material of the ablator is launched off the surface, this creates a 'rocket effect'. Focusing on one particular point on the shell, mass is launched radially outward as it ablates away, by conservation of momentum, the remaining shell must be pushed radially inward; as this happens at every point on the shell, the whole shell compresses and implodes inwards.

There are two main types of IFE: direct drive, where a laser is used to directly ablate the surface of the capsule [27]; and indirect drive where the laser is used to drive a hohlraum which generates a thermal X-ray bath that ablates the surface of the capsule [28]. These schemes have different constraints and concerns; this thesis will focus on indirect drive, in part because the only ignition scale facility currently operating is the NIF, which is indirect drive. For energy production, direct drive is likely to be more useful, but lessons from indirect drive often transfer well, particularly the imaging and shock tuning discussed in this thesis.

In either case, as the shell collapses inwards, the gas in the core is heated and compressed to higher energy density. From there there are a few main schemes, covered in more detail in 2.2. The simplest to describe is the central hotspot design, in this design the gas gets hot and dense enough that fusion reactions begin to occur, this generates energetic alpha particles some of which deposit energy in the gas and heat the hotspot further, as well as passing through the low density gas towards the wall. Once the heating caused by these alpha particles depositing their energy exceeds the energy deposited in the shell by the collapsing

shell, the implosion is considered to be a burning plasma. However, this is not sufficient to be considered ignition, as the hotspot loses energy to thermal conduction through the fuel/shell, as well as through radiation losses; once the implosion starts to disassemble, there are further losses due to the work done by the expansion. The implosion is considered to have ignited once there is enough heating from the fusion reactions to overcome all heat losses from the hotspot [29].

When the alpha particles encounter the ice shell surrounding the gas, the high density means that many more alpha particles are absorbed, so quickly deposit their energy in a thin layer of the DT ice. A propagating burn wave can then occur if there is enough deposited energy to heat the first thin layer of the ice to the point that it can begin to fuse and generate more alpha particles to heat the next layer. If this propagating burn wave burns enough of the fuel to produce more energy than the laser energy that was used to compress the implosion, then the implosion has reached scientific breakeven [14].

As well as the central hotspot ignition scheme, there are a number of alternative schemes, including some exotic ones [30], but most research on alternatives to central hotspot ignition has focused on shock ignition and fast ignition. These still compress the fuel to high energy density, but not enough to cause central hotspot ignition, therefore reducing the energy requirements for compression. Once at high energy density another trigger mechanism is used to start the propagating burn, a shock in the case of shock ignition, and fast electrons or fast protons in the case of fast ignition. By using a different trigger mechanism, these schemes reduce to total compression required, but at the cost of using some energy for the trigger mechanisms [31].

An alternative scheme, discussed in more detail in section 2.2.8 is shock augmented ignition, which aims to combine the benefits of shock ignition and central hotspot ignition [32].

This thesis will focus on indirect drive central hotspot ignition designs, as well as shock augmented ignition designs.

Chapter 2

Background

2.1 Hydrodynamics and Shocks

This section will give the necessary hydrodynamics background and shock physics to understand the motivations for some of the constraints on implosion designs, these constraints and the designs they lead to will be explored further in section 2.2.

For the purpose of clear illustrations, this section will start with the Euler equations, these are considered to be the simplest set of equations that are still useful for describing high energy density systems [33]. They are a simplification of the Navier-Stokes equations in the case of no viscosity or thermal conductivity, are a single fluid and do not include electromagnetic effects. Simple equations are chosen for ease of illustration, and more complex treatments can be found in more detail in Drake [33] or Colvin & Larsen [34].

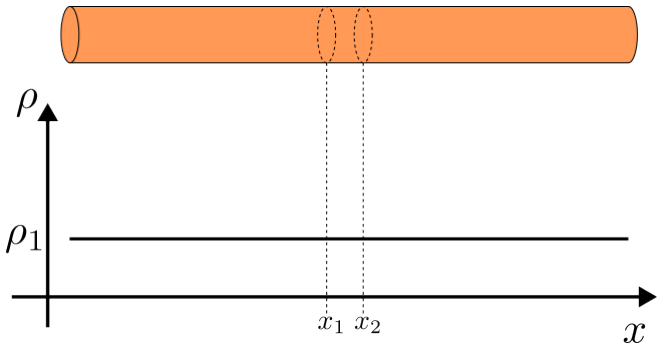


Figure 2.1: An unperturbed 1D column of fluid (shown in orange), a plot of density against position is shown below the column of fluid, with a thin slice from x_1 to x_2 identified.

Often throughout this section, a 1D treatment will be used, this can be considered to be an ideallised column of fluid, as shown in figure 2.1, initially with constant density throughout.

The column is shown as a column with a radius for ease of illustration, but should be though of as purely 1D. A thin layer is identified in the column, which will be used in the discussions around shocks. The 1D column of fluid can be considered an approximation of a small section of a spherical implosion, although spherical systems behave slightly differently as they have convergent geometry, differences will be mentioned when relevant. In general: imagining the sphere as a series of layered shells of constant radius, as radius increases, the surface area of the shells increase, which spreads out effects that move outward; as radius decreases, the surface area of the shells decrease, which concentrates effects that move inwards.

In this simplified 1D column of fluid, the Euler equations for a non-relativistic and non-viscous fluid are as follows (a good derivation can be found in Colvin and Larsen [34]):

conservation of mass:

$$\frac{\partial \rho}{\partial t} = -\frac{\partial \rho u}{\partial x}$$

conservation of momentum:

$$\frac{\partial \rho u}{\partial t} = -\frac{\partial (P + \rho u^2)}{\partial x}$$

conservation of energy:

$$\frac{\partial (\frac{\rho u^2}{2} + \rho \epsilon)}{\partial t} = -\frac{\partial (\rho u (\epsilon + \frac{u^2}{2} + \frac{p}{\rho}))}{\partial x}$$

Where ρ is the mass density, u is the fluid velocity, P is the fluid pressure, and ϵ is the specific internal energy of the fluid.

If a small perturbation is applied to this column of fluid, say by artificially increasing or decreasing the density and pressure at some point, that perturbation will travel outward at the adiabatic sound speed of the fluid. If a perturbation is introduced by pushing a piston into one side of the column, this will cause a compression wave to propagate inwards at the adiabatic sound speed of the fluid. In the context of inertial confinement fusion, this is analogous to the pressure waved caused by an ablator ablating away and exerting pressure on the outer surface of the sphere. In spherical geometry pressure waves moving outward from the core are weakened, while inward moving pressure waves are strengthened.

2.1.1 Shock formation

In the column of fluid example, with a piston pushing in from the right hand side, if the compression wave is strong enough a shock can begin to form. This is best explained by

imagining the compression wave as a sequence of smaller compression waves following each other, as shown in figure 2.2. Here the start of the wave (at x_1 is entering unperturbed material, and travels at the adiabatic sound speed of that material (c_{s1}), while the later part of the wave (at x_2) is going into compressed material, which has a higher sound speed (c_{s2}), so as the wave progresses it steepens up into discontinuity (in real fluids particle collisions smooth the discontinuity over a few mean free paths, which can be modelled in fluids as viscosity). In spherical systems an inward moving shock is strengthened by convergent effects, and an inward moving pressure wave can strengthen into a shock, whereas outward moving shocks weaken.

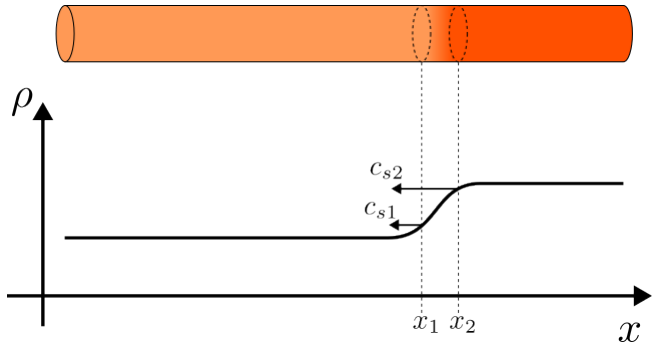


Figure 2.2: A 1D-like column of fluid with a pressure wave travelling through it from right to left, with a plot of the density against position shown below.

2.1.2 Jump Conditions

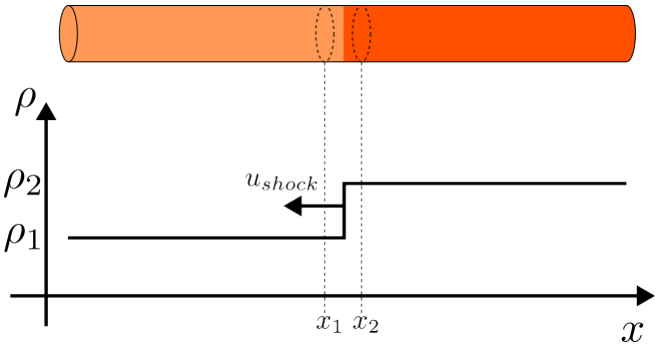


Figure 2.3: A 1D-like column of fluid with a shock travelling through it from right to left, with a plot of the density against position shown below.

Figure 2.3 shows the column of fluid from before, but now with a shock travelling through it with velocity u_{shock}. Two test points are shown, one before the shock hits (x_1), and one in the already shocked material (x_2). Using this thin layer either side of the shock, three jump

conditions can be derived from the Euler equations shown earlier. Following the derivations in Colvin and Larsen [34] and Zeldovitch and Razier [35] and starting with the mass conservation equation. The integral across the thin layer can be found:

$$\int_{x_1}^{x_2} \frac{\partial \rho}{\partial t} dx = -\int_{x_1}^{x_2} \frac{\partial \rho u}{\partial x} dx$$

If the layer is thin enough that the density does not build up or accumulate in the layer over time, then the left hand side goes to zero, giving:

$$0 = \rho(x_2)u(x_2) - \rho(x_1)u(x_1)$$

Rearranging, gives the first of the jump condition equations

$$\rho_1 u_1 = \rho_2 u_2$$

The same derivation can be done for the momentum and energy conservation equations to get the other jump conditions

$$\rho_1 u_1^2 + p_1 = \rho_2 u_2^2 + p_2$$

$$\rho_1 u_1 \left(\epsilon_1 + \frac{u_1^2}{2}\right) + p_1 u_1 = \rho_2 u_2 \left(\epsilon_2 + \frac{u_2^2}{2}\right) + p_2 u_2$$

Note that this describes the state either side of a transition, and does not depend on the thickness of the transition.

Assuming the density and pressure before the shock are known, and that the internal energy can be found from pressure and density using an equation of state (which is material dependent), this means there are three equations for four unknowns (ρ_2, p_2, u_2, u_1), so with some measure of the shock strength, the end state can be fully determined.

2.1.3 Shock Hugoniots

The jump conditions can be rearranged to find the Rankine-Hugoniot equations, which describe the set of states a material can be in after experiencing a shock, based on the conditions of the material before the shock, the equation of state and the shock strength.

Starting with the conservation of momentum equation, the conservation of mass equation can be substituted in to remove either the velocity before or after the shock. Then by

rearranging, and substituting in the specific volume $V = 1/\rho$, the following pair of equations can be found:

$$u_1^2 = \frac{1}{\rho_1}\frac{P_2 - P_1}{1/\rho_1 - 1/\rho_2} = V_1\frac{P_2 - P_1}{V_1 - V_2}$$

$$u_2^2 = \frac{1}{\rho_2}\frac{P_2 - P_1}{1/\rho_1 - 1/\rho_2} = V_2\frac{P_2 - P_1}{V_1 - V_2}$$

Substituting these equations into the conservation of energy equation, and using a difference of two squares to remove the denominator ($V_1^2 - V_2^2 = (V_1 - V_2)(V_1 + V_2)$), a form of the shock Hugoniot can be found:

$$\epsilon_2 - \epsilon_1 = \frac{1}{2}(P_2 - P_1)(V_1 + V_2) \tag{2.1}$$

This can be rearranged to find the usual shock Hugoniot:

$$P_2 = H(V_2, P_1, V_1) = 2\frac{\epsilon_2 - \epsilon_1}{V_1 + V_2} + P_1$$

Here the shock strength is characterised by the change in specific volume across the shock (equivalent to the density jump), from which the final pressure can be found, alternately the final density could be found from the shock pressure. Although again, the equation of state of the material is required to determine the internal energy from the pressure and density.

2.1.4 Equation of State and Maximum Single Shock Compression

The jump conditions and Hugoniot equations both leave the internal energy as an unknown, which must be found from the density and pressure, the equation of state provides a closure which gives this value.

The simplest meaningful equation of state is the ideal gas law [33]:

$$PV = Nk_bT$$

Where P and V are Pressure and Volume as before, N is number of particles, k_b is boltzmans constant and T is Temperature. This can be expressed in terms of number density (n):

$$P = nk_bT$$

The number density can be obtained from the ionisation state (Z), mass of ions (A), proton mass (m_p) and density. So: $n_{atoms} = \rho/(Am_p)$, and $n_{particles} = n_{atoms}(1 + Z) = \rho(1 + Z)/(Am_p)$.

This model is a useful approximation but, for real plasmas with different ionisation levels, more accurate equation of state models are often needed. For the high energy density context, these are based on first principle calculations [36] constrained to experimental measurements [37]; although in some cases there are simple equations which are appropriate approximations, see secitons 3.1.2 and 3.1.3 in Drake [33].

An ideal gas is still a reasonable approximation, and taking a form of the ideal gas equation of state derived in Colvin and Larsen [34]:

$$\epsilon = C_V T = \frac{1}{\gamma - 1} PV$$

where the adiabatic index is defined as the ratio of specific heats $\gamma = C_p/C_V$. Substituting this equation of state into equation 2.1, rearranging, and taking the case where $P_2 >> P_1$, gives the maximum density ratio across a shock:

$$\frac{\rho_2}{\rho_1} = \frac{\gamma + 1}{\gamma - 1}$$

So in the limit of very strong shocks, the maximum compression ratio is asymptotically limited to a value set by the adiabatic index (although this assumes the adiabatic index does not change across the shock). For a monatomic ideal gas, this limits the maximum compression from a single shock to a factor of 4. So to get the compression needed in ICF, multiple shocks must be used (although, depending on the degrees of freedom available within the molecules making up the material, the maximum compression can go up to a factor of 13 for some types of ideal gasses [35]).

2.1.5 Strong Shocks Increase Entropy

Again using a simple model, for an ideal gas with constant heat capacity, the change in entropy across a shock is given by:

$$\Delta S = c_v ln \left(\frac{P_2 V_2^\gamma}{P_1 V_1^\gamma} \right)$$

[35] For weak shocks, the ratio of specific volumes and the pressure ratio are close to unity

$$\Delta S \approx c_v ln(1) = 0$$

so weak shocks introduce little additional entropy to the system. For strong shocks, the maximum compression is limited to a factor of ~4 in strong shocks, so the ratio of specific volumes can be treated as constant, meaning that for the entropy simplifies to:

$$\Delta S \approx c_v ln\left(\frac{P_2}{P_1}\right)$$

So for strong shocks the increase in entropy is proportional to the log of the shock strength.

This can be conceptually understood by thinking about a shock going through some unshocked material from the perspective of the shock. Thinking in terms of particles: in the frame of the shock, all the unshocked particles are moving towards the shock front. The unshocked material moves in an ordered manner, with all its velocity pointing towards the shock front. When the shock passes through the material, collisions occur between particles. This takes the ordered velocity from the unshocked material and thermalises it into disordered motion, thereby increasing the entropy of the now shocked material. A stronger, and therefore higher velocity shock, has more kinetic energy to convert into thermal energy, so increases entropy more than a weaker shock would.

2.1.6 Rarefactions

Taking the column of fluid with a piston at one end, the piston can be withdrawn out of the fluid instead of pushing into the fluid, this reduces the pressure at the end of the column with the piston. The reduction in pressure propagates into the column as a rarefaction or release wave. Unlike compression waves, release waves cannot form shocks (without phase transitions), since waves travel more slowly in the rarefied material, so the later parts of the rarefaction will not catch up with the earlier parts of the rarefaction.

2.1.7 Shocks and Rarefactions at Interfaces

As a shock crosses an interface from a high density region to a lower density region or vacuum, the shocked material is free to quickly flow into the lower density region, this lowers the pressure in the region of high density material which the shock previously traversed. This

"unloading" acts similar to the piston being pulled out of the column of fluid as described in section 2.1.6, and creates a rarefaction which propagates back in the opposite direction from the shock.

Another interaction of interest is when a rarefaction encounters a region of fixed pressure, which will be relevant in later sections discussing rarefactions encountering the ablation front in an ICF implosion. At the ablation front the ablation pressure is determined by the intensity of the radiation which is impinging on the surface and driving the ablation process. When a rarefaction approaches the ablation front it represents a drop in pressure relative to the pre-rarefaction material, and since the ablation pressure is effectively fixed by the radiation intensity, the interaction between a rarefaction and the ablation front increases the gradient in pressure between the ablation front and the rarefied material. This gradient in pressure propagates in the opposite direction from the rarefaction wave, back in towards the core of the implosion. If this gradient is strong enough, the pressure wave may form a shock, or steepen up into a shock as it propagates. In ICF, this feature was noted by Betti et. al. [38] and is especially relevant in ICF implosions as the convergent geometry of the implosion aids the shock strengthening process for inward propagating shocks.

2.1.8 Mutli-Dimensional Instabilities

Thus far a 1D model has been sufficient for illustrative purposes, however, there are some features which are inherently 2D and 3D. In particular, hydrodynamic instabilities require a multidimensional treatment. When a region of low density accelerates a higher density region, the Rayleigh-Taylor instability (RTI or RT instability) forms. This instability is seeded by non-uniformaties, which grow into spikes and bubbles of the low and high density regions interpenetrating each other. If the acceleration takes the form of a shock, the Richtmyer-Meshkov instability (RMI) can be seeded by surface non-uniformities, which also results in interpenetrating spikes and bubbles between the two fluids. Along the surfaces of these bubbles and spikes shear flow occurs, which allow the Kevin-Helmholtz instability to form. The combination of these instabilities mixes the two fluids together [33].

Non-uniformities throughout the layers in a capsule can contribute to RTI and RMI seeding at other layers: outer layer non-uniformities lead to non-uniform shocks or pressure waves which seed RTI and RMI at inner layers, and inner layer non-uniformities lead to non-uniform rarefactions propagating outwards after shocks or pressure waves pass through

the inner layers which seeds RTI and RMI instabilities at the outer layers. Non-uniformities of different sizes seed different modes of instabilities, with large non-uniformities seeding low mode instabilities, while small non-uniformites seed high mode instabilities. High mode instabilites saturate quickly, while low mode instabilities grow more slowly, so intermediate mode instabilities are the largest concern for ICF.

In the ICF context, the ablation plasma is much lower density than the bulk shell material, and the ablation process drives an acceleration inwards. This allows the RT instability to grow, and as a result some of the ablation plasma can mix into the bulk shell of the implosion. This is partially mitigated by ablative stabilisation, since as the ablation front removes material, it preferentially ablates the peaks of non-uniformities, creating a modulation in ablation pressure that dampens instability growth. The RT instability also occurs later in the implosion, during the deceleration phase, as the low density hotspot slows the higher density shell. As fusion occurs in the hotspot, alpha particles and heat conduction from the hotspot can ablate the inner surface of the ice, acting as another form of ablative stabilisation [39]. The implosion process is described in more detail in section 2.2, where the implications of the RT instability on implosion design will be explored.

2.2 Implosion Design

This section will include the key background relevant to designing implosions, starting with the ablation processes involved in direct and indirect drive implosions, before outlining the central hotspot ignition scheme. Hydrodynamic instabilities and their relevance for key fusion metrics will be described, providing motivation for some of the design choices made by scientists at the NIF, particularly focusing on shocks and shock timing. Finally, advanced ignition schemes will be described, with a focus on shock augmented ignition.

2.2.1 Ablation, Direct Drive and Indirect Drive

A core idea of an ICF implosion is to use radiation to ablate the surface of a shell of material, causing it to collapse inwards and compress the fuel contained in the capsule to the high energy density conditions required for fusion to occur. This process will be described in more detail in following sections, for now ablation and driver for it will be the focus.

The ablation process is an application of newtons third law, which states that two bodies

exerting forces on each other will have equal magnitude and opposite directions [40]. This gives rise to the "rocket effect", where if some material is launched in one direction (eg the exhaust of a rocket), then the object launching the material (ie the rocket) must move in the opposite direction to conserve momentum. In the ICF context, radiation is impinging on the surface of a spherical capsule, the radiation heats the surface until material from the surface of the capsule is launched off the surface. By conservation of momentum, this exerts a force on the surface of the capsule, and since the whole surface of the sphere is being ablated, each point on the sphere is being accelerated towards the centre of the sphere. The result is compression of the spherical capsule.

Ablation in ICF is driven by radiation, either by direct irradiation of the surface of the capsule by lasers (direct drive), or by an X-ray bath produced by a hohlraum which is in turn heated by laser irradiation (indirect drive). The ablation pressure produced by the radiation on the surface of the capsule is determined by the radiation intensity and ablator material and can therefore be adjusted by adjusting either of these. This allows shocks to be launched and their strength to be adjusted by changing the radiation intensity on the ablator surface over time.

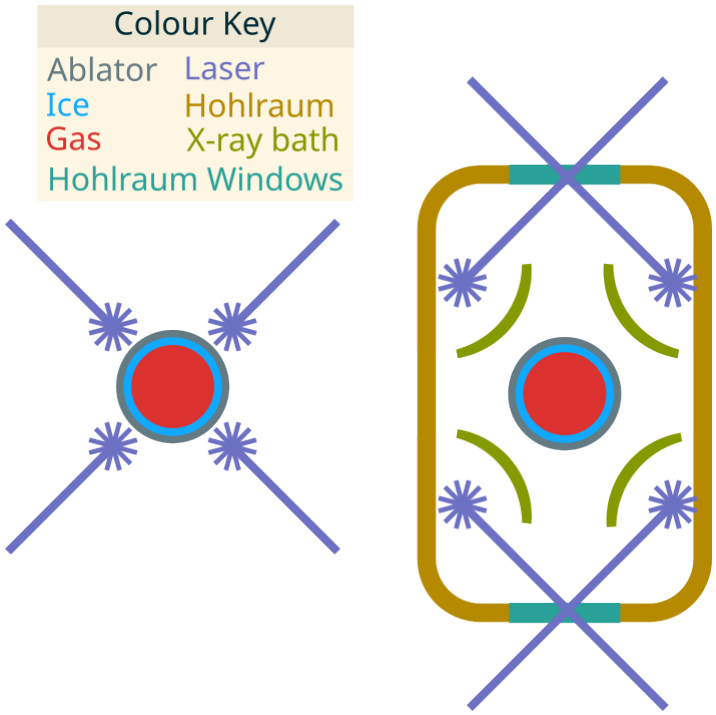


Figure 2.4: Schematic showing a direct drive implosion (left) and an indirect drive implosion (right), with key features colour coded for clarity. In both only show a few laser beams, while in reality many more beams are required for even illumination of the capsule.

Direct drive (shown on the left of figure 2.4) is often considered the favoured route towards inertial fusion energy, due to its efficient coupling between laser energy and radiation on target [41], as well as simple target design relative to indirect drive. Laser spots have a structure to them called a speckle pattern, where some parts of the laser spot have higher or lower intensity, and unfortunately as direct drive ablates the capsule surface with the laser itself, this pattern is imprinted on the capsule surface. As the implosion progresses the imprinted speckle pattern seeds density fluctuations which grow via the Rayleigh-Taylor instability, the problems these cause are discussed in a later section. This problem can be mitigated by techniques such as smoothed spectral dispersion or induced spatial incoherence, which smooth the structure of the laser spot to reduce the speckle pattern.

Indirect drive (shown on the right of figure 2.4) works by using lasers to heat the inner walls of a hollow gold cylinder (the hohlraum), the walls then re-emit light as X-rays which bathe a spherical capsule at the centre of the hohlraum. This has the key advantage that it smooths the laser drive as thermal conduction and the incoherence of the X-rays evens out the speckle pattern of the laser, which reduces the seeding of RT instabilities. The downside of this process is that the conversion of laser energy to X-ray energy, and X-rays in the hohlraum to X-rays on the capsule, both have losses associated with them, and the coupling of laser energy to radiation on target suffers as a result. Indirect drive also requires a more complex target: a hohlraum must be manufactured in addition to the capsule, and the capsule must be suspended within the hohlraum.

Another key difference between indirect and direct drive is the capsule material, plastic and high density carbon readily absorb high intensity optical laser light, so can be used as ablator materials. Both also absorb X-rays to an extent, however, the high energy x-rays produced by the M-band in gold hohlraums could penetrate deep enough to deposit energy in the core of capsules. This results in preheating of the fuel, which makes it harder to compress. To mitigate pre-heating risks, indirect drive capsules include dopeants in the shell to prevent high energy x-rays from penetrating to deeply into the capsule. Indirect drive must also contend with gold bubbles which form in the hohlraum around the locations of laser spots, these bubbles block laser light once developed and disrupt the implosion drive. Currently indirect drive designs use high density carbon as the ablator, even although it is less efficient than plastic, because the high density allows for thinner shells which shocks can traverse more quickly, which reduces the time available for gold bubbles to form.

Currently, the largest direct drive facility is Omega in Rochester, New York state. Omega has 30kJ of laser energy on target, and is a sub-ignition scale facility. The only ignition scale facility currently operating is the National Ignition Facility (NIF) in Livermore, California; which is an indirect drive facility. As of writing, the NIF has 2.2MJ laser energy on target, although upgrades are being considered to increase this. In December 2022, NIF completed an experiment in which a target produced 3.15MJ of energy, more than the 2.05MJ of laser energy used to drive the implosion [14], and in April 2025 NIF completed an experiment with the same laser energy but producing 8.6MJ of fusion energy. This is the key reason indirect drive was chosen as the focus of this thesis, it is the only scheme for which ignition scale experiments can be executed as well as designed.

2.2.2 Central Hotspot Ignition

The method for reaching ignition which has proven the most effective so far is central hotspot ignition [14, 29]. This involves heating a small region of gas, the 'hotspot', at the core of the spherical implosion using the work done on the gas by the capsule shell as the capsule compresses and as shocks traverse the gas. This hotspot must be heated sufficiently for fusion reactions to occur, typically a temperature of at least 5keV is required. The mass of the remaining ablator and fuel ice confines the hot gas long enough for significant numbers of fusion reactions to occur, these reactions produce alpha particles which heat the hotspot further, allowing more reactions to take place. For this self-heating to occur, a minimum hotspot areal density (ρR) must be reached [33]:

$$(\rho R)_{hotspot} > 0.3 g/cm^3$$

Once the reactions have began in earnest, the ice layer which is acting as the confining mechanism is also exposed to these alpha particles, and since the alpha particles are stopped a short distance into the high density ice layer, all of the energy going into the ice is used to heat the thin layer closest to the hotspot. In an igniting implosion this layer must be heated enough that it also begins to fuse, at which point the next innermost thin layer of ice begins to be heated. To reach ignition, enough fuel must burn rapidly enough to overcome the loss mechanisms that cool the implosion (eg disassembly, conduction, radiation) and therefore to sustain the reaction. To ensure enough of the fuel ice burns and ignition occurs there is an

additional areal density requirement for the total fuel, in indirect drive implosions this 'cliff of ignition' is typically:

$$(\rho R)_{shell} > 1.2 g/cm^3$$

[42]

The fuel burnup is related to the areal density:

$$f = \frac{\rho R[g/cm^2]}{\rho R[g/cm^2] + 7}$$

[23]

So higher areal density leads to higher fuel burnup and therefore higher yield.

2.2.3 Hydrodynamic Instabilities

The Rayleigh-Taylor and Kevin-Helmholtz instabilities were mentioned in section 2.1.8, and act to mix two fluids together when a low density fluid is accelerating a higher density fluid. This occurs at two points during an implosion, the first is as the ablation occurs, this creates a diffuse ablation plasma which is pushing on the high density shell of ablator and DT ice. During this phase of the implosion, depending on the relative densities of the ablator and fuel ice, the low density ablator plasma can be mixed into the fuel by the Rayleigh-Taylor instability. The second time this occurs is as the implosion approaches stagnation, when the low density hotspot is decelerating the high density shell, which mixes shell material into the hotspot.

Both cases are potentially problematic as they can result in ablator material with high atomic number (high-Z) atoms mixing into the fusion fuel. Mixing can also be caused by the presence of the fill tube which is used during capsule manufacture to fill the capsule with DT, the tent which acts as a support structure to hold the capsule in place, and any other impurities or non-uniformities in the shell. These structures and non-uniformities act as seeds for hydrodynamic instabilities, and can cause non-uniformitites in the radiation drive, which can also mix high-Z material into the fuel. High-Z materials are much more effective at emitting bremstrahlung radiation, so if they are in the core of an implosion they will allow much more radiation to escape, cooling the hotspot and making ignition more difficult. Mixing at the ice-gas interface can also be a problem as it mixes cold ice into the hotpsot,

cooling it and making ignition harder to achieve.

2.2.4 Fusion Metrics

Alongside areal density, there are a number of fusion metrics which are relevant for implosion design. A common metric used across fusion, including magnetic confinement fusion, is the Lawson criterion, sometimes expressed as a triple product:

$$nT\tau > 3 \times 10^{15} keVs/cm^3$$

where n is the number density of deuterium and tritium ions, T is the temperature of the fuel, and τ is the confinement time of the fuel. In practice there are a few other ignition criteria which are more useful in ICF, but they are roughly equivalent to the Lawson criterion. These are described in a report by Tipton [43]. These metrics were particularly useful before the first ignition shots [29, 14] as ways to measure how close to ignition an implosion was.

Yield and Gain are of course key metrics, and maximising gain is critical for the success of IFE, however, these can be difficult metrics to use in practice as yield becomes highly non-linear close to ignition. This means that small changes in implosion design can result in large changes in yield, and different shots with nominally identical designs can have significant variation in yield. As a result, when designing implosions other metrics must be used to check how robust a high yield design actually is.

In-flight aspect ratio (IFAR) is a key metric for measuring hydrodynamic instability. IFAR is the ratio of the capsule radius to the thickness of the shell, so implosions with high maximum IFARs are particularly thin relative to their radius. This is relevant because as Rayleigh-Taylor instabilities grow causing mix, the thicker the shell is, the harder it is for ablator material to mix into the fuel. So low IFAR designs are more hydrodynamically stable.

The fall line is another metric that can be used to measure the potential for ablator material to mix into the fuel. This measures the worst case scenario for mixing: assume that the ablator shell is not decelerated after peak velocity is reached, and project how far into the fuel the ablator material would reach at bangtime if it continued at peak velocity. In this thesis the fall line behaviour will be given as the radius reached by the ablator-ice boundary if it continued at peak velocity, a smaller radius therefore means more potential for mix.

Implosion velocity is a key metric relating to Rayleigh-Taylor growth, implosion velocity is

the velocity of the shell as it collapses inwards, and normally refers to the maximum velocity the shell reaches while imploding. Higher implosion velocities increase the growth rate of RT instabilities, so are less hydrodynamically stable. There are benifits to higher velocity, as higher implosion velocity means the shell has more kinetic energy available to be converted into heating the hotspot over the course of the implosion. For a NIF scale central hotspot ignition, 330km/s is considered the minimum implosion velocity needed to get ignition [28], but advanced schemes can lower this requirement.

Adiabat is another key metric, and acts as an easier to measure proxy for entropy. The adiabat (α) is given by the pressure over the Fermi degenerate pressure P_f:

$$\alpha = P/P_f \propto P/\rho^{5/3}$$

Where the Fermi degenerate pressure is the pressure that would be exerted by the fuel if it had zero temperature and Fermi-degenerate electrons. Materials with higher adiabat are harder to compress, and so in implosions, a higher shell adiabat will make it difficult to obtain high areal density and high yield [42]. Unfortunately, lower adiabat implosions tend to be more susceptible to Rayleigh-Taylor instabilities, so there is a balance to be struck.

2.2.5 Shock Design

The main factors that set the adiabat of an implosion are the strengths of the various shocks which go through the capsule, particularly the earliest shocks. Stronger shocks increase the adiabat more than weaker shocks, as discussed in seciton 2.1.5. When two shocks merge, they form a stronger combined shock, which increases the adiabat more than the two weaker shocks would separately. When generating a given compression ratio, a larger number of weaker shocks will result in a far lower adiabat implosion than smaller number of stronger shocks [22]. The standard NIF design currently uses three shocks, and these shocks are timed to merge just after crossing the ice-gas interface, thereby minimising the amount of ice which the strong shock passes through and therefore minimising the adiabat [44, 45].

However, while shocks merging too early will increase the adiabat, allowing the shocks to pass too far into the gas without merging will also increase the adiabat. This is because after the first shock crosses the ice-gas interface and breaks out, the ice begins to decompress and a rarefaction begins to work outwards away from the centre of the implosion. If a second shock

travels through this decompressed material, it is travelling through material with reduced density and pressure, so the pressure ratio and therefore effective shock strength increases.

2.2.6 Fast Ignition

Fast ignition is an advanced fusion scheme, which still requires significant compression, but uses fast electrons or fast ions instead of a central hotspot to ignite the implosion. For electron fast ignition: after the fuel is compressed and ready to ignite, a high intensity laser impinges on the capsule, boring a hole and generating fast electrons which can penetrate to the core of the capsule and trigger ignition. For cone guided fast ignition, a gold cone is embedded in the capsule to guide the laser pulse so that the fast electrons are generated near the core of the implosion. For proton or ion fast ignition, an ion or proton beam is used instead, typically by focusing the high intensity laser on a foil outside the capsule [31].

2.2.7 Shock Ignition

Shock ignition is another advanced fusion scheme, which separates the compression and heating phases of the implosion. Shock ignition implosions are compressed to high areal density in a similar way to central hotspot ignition implosions, but with a slower implosion velocity and lower adiabat to minimise heating. By having lower implosion velocity, this scheme is less susceptible to Rayleigh-Taylor growth [46]. Additionally, by having lower velocity, more mass can be assembled with a given kinetic energy, allowing for higher fuel burnup fractions and higher gain.

The assembled implosion is then ignited with a late time spherical shock [47], which is launched by rapidly increasing the radiation intensity on the capsule. This strong shock heats the fuel it passes through, and is strengthened by spherical convergence. Ideally the shock is timed to meet the rebound shock from the compression phase immediately before entering the gas core, the shock collision amplifies the shock and this strong shock heats the gas core into an igniting hotspot. To launch this shock, requires that the laser driving the compression is able to increase in intensity rapidly and go to high intensity: 10^{16} W/cm^2 as opposed to 10^{15}W/cm^2 for central hotspot ignition designs. This risks damaging optics in the laser chain. Higher intensity lasers also have increased laser plasma instabilities, which reduce the energy delivered to the shell and generate hot electrons which preheat the fuel making it harder to compress [46].

2.2.8 Shock Augmented Ignition

Shock augmented ignition (SAI) aims to achieve the benefits of shock ignition, without the need for high intensity lasers. This is achieved by assembling the fuel in a similar way to central hotspot ignition, but part way through the implosion the radiation intensity is reduced, then increased back up. The drop in radiation intensity allows the capsule to partially decompress and preconditions the plasma so that when the radiation intensity is increased again, a shock can is launched. As a shock is used to ignite the implosion, a lower implosion velocity can be used, which reduces the susceptibility to Rayleigh-Taylor instabilities. In addition, since the capsule is allowed to decompress during the dip, the peak IFAR is lower, which also reduced the susceptibility to RT instabilities. The dip allows a shock to be launched without requiring much higher intensity than central hotspot ignition designs, and therefore reduces concerns around laser plasma instabilities.

Experiments with warm Deuteriam-Deuterium targets in a direct drive configuration have shown that with appropriate timing of the drop in laser intensity, significant improvements in areal density (50%) and hotspot pressure (112%) can be achieved [48].

Indirect drive SAI experiments have also been performed on the National Ignition Facility. These also showed increased areal density relevant to equivalent non-SAI implosions, and demonstrated that hohlraums can achieve the required radiation temperature profiles [49].

The implosion design explored in chapter 5 is complementary to shock augmented ignition, and will explore a scheme which combines the N+1 Tuning design and the shock augmented ignition design.

2.3 Xray Imaging

X-ray imaging has been used extensively during the over 100 years since the discovery of X-rays [50], and across a range of fields (particularly medical imaging [51]). X-ray imaging is useful because X-rays are more weakly attenuated by matter than optical light, so X-ray imaging allows scientists see the internal structure of objects which are opaque to optical wavelengths.

In the context of ICF, X-rays have been used for some time to discern the structure of capsules and implosions, which is particularly useful for mid-implosion or in-flight measurements when the plasma corona is ionised and optical light has a critical surface at which it

is fully absorbed and therefore cannot penetrate through the capsule [52]. Recently there has been interest in developing more advanced imaging techniques like X-ray phase contrast imaging for use in ICF [15]. Capsules also emit their own X-rays as materials in implosions can have high enough temperatures for their thermal emission to enter the X-ray regime.

This section will describe the theory and state of the art for absorption X-ray imaging, as well as phase contrast imaging; then it will go on to briefly discuss self-emission and the X-ray sources available to act as backlighters for absorption and phase contrast imaging.

2.3.1 Complex Refractive Index

One key equation to help understand the differences between absorption and phase contrast imaging with X-rays is the complex refractive index:

$$n = 1 - \delta + i\beta \tag{2.2}$$

which characterises an X-rays interaction with matter. The $i\beta$ component represents attenuation and is used in absorption imaging, while the $1-\delta$ component represents the phase shift a wave experiences as it goes through some material and is used in phase contrast imaging. The phase shift manifests either as the deflection of X-rays as they go through an object, or as diffraction, depending on the imaging regime.

For most materials in the X-ray regime δ is small and positive, making the real component of the refractive index less than 1 (see [53] for a discussion on this), and for X-rays above 5keV δ is much larger than β [54]. The values taken in cold materials are proportional to the mass density, and the value of δ is proportional to the electron density in plasmas when the X-ray energy is far away from the absorption edges [55]. The value of β in plasmas with X-ray energy far away from absorption edges is proportional to the electron density squared [56]. As a result both the absorption and phase components give information about the density fluctuations within the object they image. However, the effect of δ is more difficult to extract than β, so absorption imaging is easier to use and was developed first. The imaging setups which are designed to extract the phase effects do so as either the gradient or laplacian of the phase shift, as a result these setups are useful for imaging structures with sharply varying densities (eg shocks or interfaces between materials).

2.3.2 Absorption Imaging - Theory

Traditional radiography measures the reduction in intensity caused by a material's absorption, this intensity drop is given by:

$$I = I_0 e^{-\mu}$$

with the absorption coefficient as:

$$\mu = \frac{4\pi}{\lambda} \int \beta(s) ds$$

where I_0 is the initial intensity, s as the path of the ray, λ as the wavelength, and $\beta(s)$ as the real component of the refractive index at a particular location in the object [57]. In the case of an object with constant β, this collapses to the Beer-Lambert law

$$I = I_0 e^{-\frac{4\pi}{\lambda}\beta s}$$

A diagram of a typical point-projection based absorption imaging setup is shown in figure 2.5, here X-rays (shown in green) are emitted from a source on the left, the X-rays go through the hollow capsule which is being imaged (shown as a cross section of a blue shell), until they hit the detector (shown as a vertical purple line). As the X-rays go through the object they are partially absorbed, and above the object a synthetic absorption image is shown, the outline of the shell is clearly visible, and by looking at the lineout (shown on the right of the detector image), the drop in intensity is clear. On the lineout, the edges of the shell are picked out. The source to object distance (R_1) and object to detector distance (R_2) are also labeled, absorption imaging is also sometimes called "contact imaging" because of the very short distance between the object and the detector (R_2).

Absorption imaging allows differences in the line integrated density to be detected, which can be combined with other information about the target (or multiple images from different lines of sight) to determine information about the target geometry and density.

2.3.3 Absorption Imaging - Applications in ICF

Absorption radiography is used extensively in inertial confinement fusion, and has been used to determine information about in-flight capsule dynamics [58].

A number of different methods are employed, and Kozioziemski et al.[52] provide a good

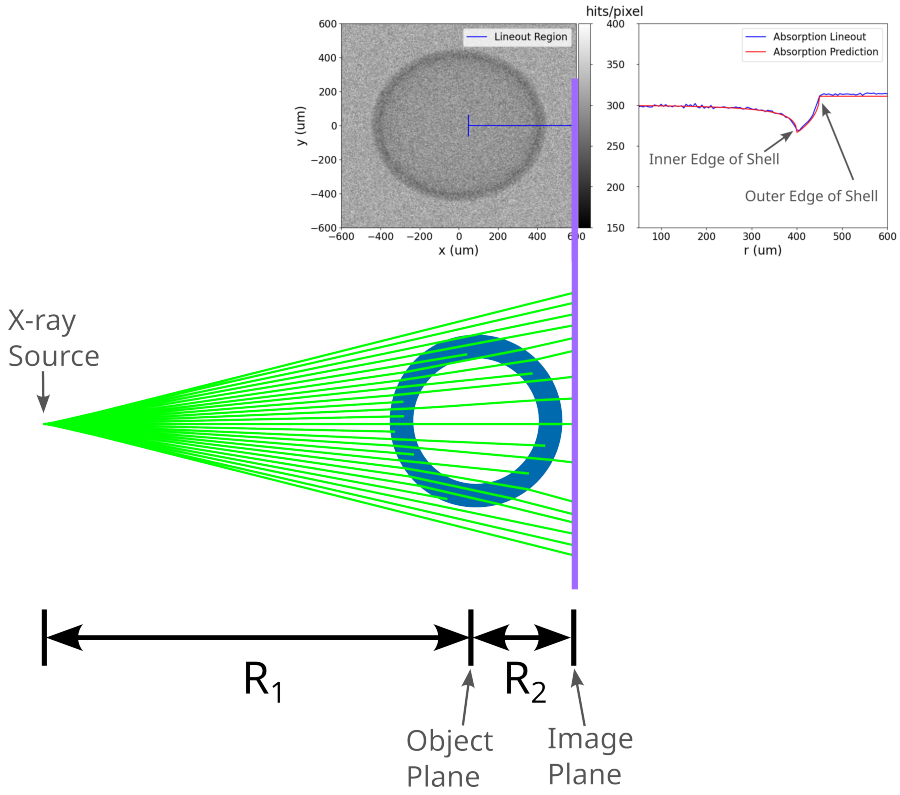


Figure 2.5: Diagram showing absorption imaging

overview, but point projection imaging will be the focus for this section of the thesis. A typical point projection setup is shown in figure 2.5, it is a useful tool in part because the experimental setup is relatively simple, but it is the focus here because the setup is very similar to the propagation based imaging setup that will be discussed in the X-ray phase contrast imaging section.

An ideal imaging diagnostic would have high spatial and temporal resolution, which in are coupled by the velocity of the object. In the ICF context objects are often moving at 100's of km/s (or 100's of μm/ns in more relevant units), as a result objects will spatially blur if the temporal resolution of a diagnostic is too low. An ideal imaging diagnostic would also have high good contrast across the features of interest and high signal to noise ratio. In practical experiments there are a number of trade offs, the rest of this section will illustrate them and give an idea for the current state of the art of absorption imaging.

Starting with resolution, spatial resolution is limited by penumbral blurring from the effective source size and the pixel size on the detector. In point projection imaging, for a finite source size, the light from the top of the source and the bottom of the source take

different paths through the object to reach the detector, this blurs the object by a distance (d) given by this equation (for the distance on the object plane):

$$d = s \times \frac{(M-1)}{M}$$

[52] where s is the source size and M is the magnification (given by: $M = \frac{R_1+R_2}{R_1}$[57], see figure 2.5 for R_1 and R_2). This means smaller source sizes are desirable for higher resolution, and has motivated the use of laser driven wires (described in more detail in section 2.3.7). Alternatively larger physical sources, or sources which expand over the imaging time, can be used by placing a pinhole immediately after the source of X-rays. This allows a wider range of sources, some of which are physically large, while still limiting the effective source size [59].

The penumbral blurring equation also shows that larger magnification increases the penumbral blur and therefore reduces resolution, however, this is a trade off with the other key factor affecting spatial resolution: detector pixel size. For a fixed detector pixel size (typically many microns), increasing the magnification allows for higher resolution images. Ideally a balance point is found where magnification is optimised so that penumbral blurring and the detector resolution are comparable [52].

Typically, in the ICF context resolution better than 10μm is difficult to achieve with point projection alone [52]. Pinholes on NIF are not typically less than 10μm [16], in part because small pinholes limit the flux of X-rays.

Some advanced imaging techniques can work around this trade off using X-ray optics of various kinds, one such example is a fresnel zone plate. Fresnel zone plates focus X-rays with diffraction, and have been shown to have as good as 3μm spatial resolution [60]. However, in practice Do et al [60] found the resolution was limited to 18μm by motion blur on a 600ps gated image plate, and by low signal to noise ratio on a framing camera which had a shorter time window.

This sets up another trade off relating to the flux and signal to noise ratio, reducing the motion blur requires a short time window to gather data, but this reduces the total photon flux on the detector. In addition, reducing spatial blur typically involves reducing the effective source size, which also reduces the total photon flux on the detector.

One way to work around this trade off is to have brighter X-ray sources, developments in laser wakefield accelerators are particularly promising, and X-ray sources are discussed in

more detail in section 2.3.7. Another way to improve signal to noise ratio is to reduce the noise from the implosion, most of the noise comes from self-emission. Self emission is discussed in more detail in section 2.3.6, for the purpose of imaging it is a thermal source of X-rays. The noise from self emission can be avoided by using high energy X-rays as a backlighter and filtering the softer X-rays, or by using a narrow band X-ray source and filtering detectors to just that band [61].

The final limitation of absorption imaging is of particular importance, as it forms the motivation for using X-ray phase contrast imaging. This is that imaging contrast is limited by the density contrast of the actual object. In the case of a shock going through some material, the maximum jump in density across the shock is typically a factor of 4 [35], since absorption imaging is sensitive to the density, this means the difference in transmission between shocked and unshocked material is limited, which makes it difficult to experimentally observe shocks with absorption radiography.

2.3.4 Phase Contrast Imaging - Theory

While absorption imaging is extremely useful, as discussed, it also has some limitations, particularly on contrast when trying to understand shocks and view the interfaces within a capsule. An advanced technique which improves upon this is to use phase contrast effects to highlight interfaces, which works particularly well in contexts like ICF where many materials are constructed from weakly absorbing materials. Montgomery [15] provides a good summary of XPCI as applied to ICF and high energy density science.

Returning to equation 2.2, as well as absorption, the phase of light is also shifted as it goes through a material, the total phase shift resulting from the δ component of the refractive index is defined as:

$$\phi = \frac{-2\pi}{\lambda} \int \delta(z) dz$$

[57] This phase shift cannot be directly measured, however, experimental setups can be designed to infer information about the phase shift, generally as the gradient or laplacian of the phase. There are a number of different X-ray phase contrast imaging techniques, which extract phase information in different ways, for a brief description and comparison of the three main categories of techniques, see Diemoz et. al.[62], for a more in depth discussion see chapter 49 of Russo [63]. This thesis will focus on propagation based imaging (PBI), which

uses a very similar setup to the point projection absorption radiography shown in 2.5. Figure 2.6 shows the PBI setup, with the key difference from the absorption setup being that the detector is placed further away, allowing phase effects to become significant.

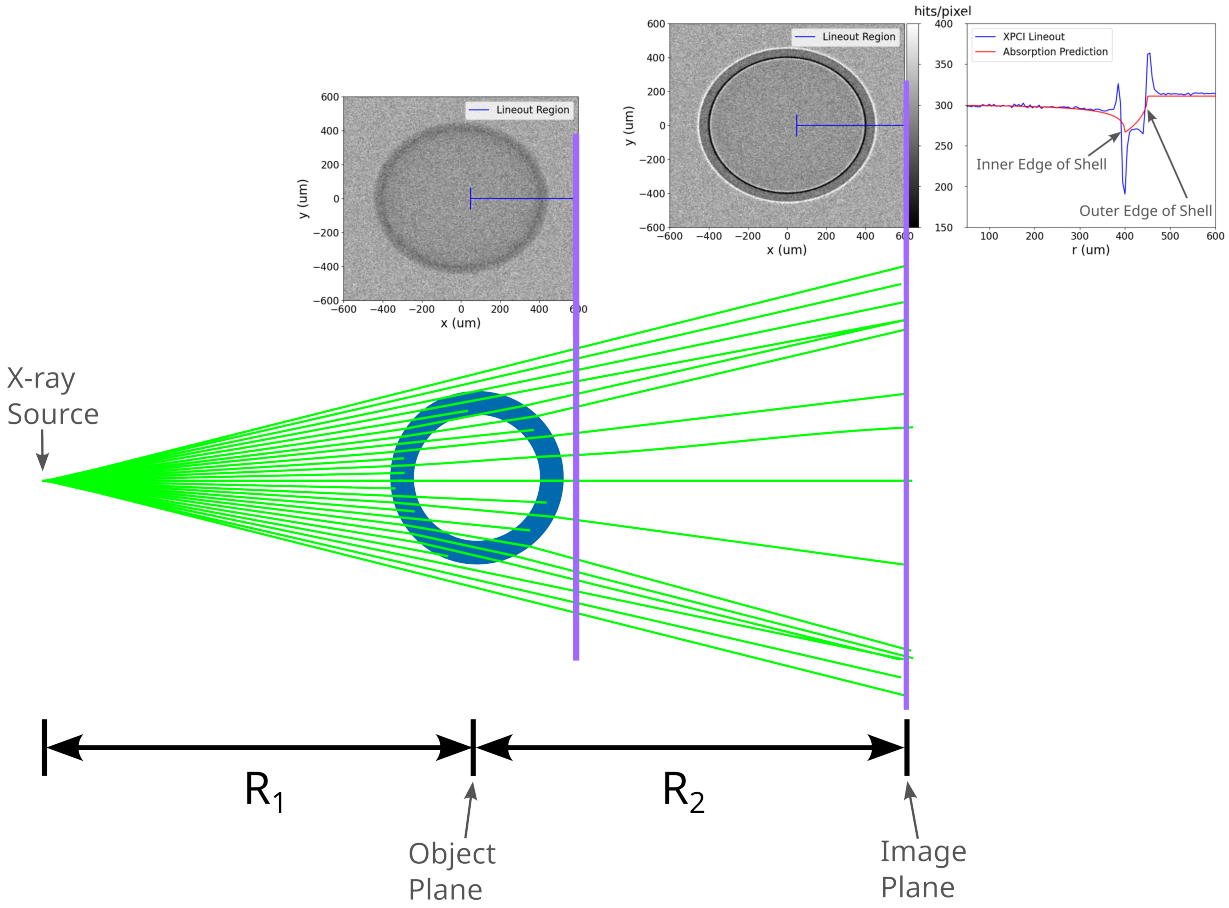


Figure 2.6: Diagram showing refraction generating phase contrast effects

For a full treatment from kirchhoff diffraction theory, see Pogany et. al. [64] and Cowley [65]. However, there are a number of approximations which mean that in the ICF context, the geometric optics limit can be used [66]. These approximations are:

- in the paraxial limit
- the object is optically thin (ie total $\phi + \mu$ is small)
- in the diffractive near field (ie for large fresnel number: $1 << N_f = \frac{a^2}{\lambda R_2}$, where a is the typical size of the feature of interest)

With these approximations the intensity from a backlit pure phase object in a propagation based imaging setup is [15]:

$$I \approx I_0 \left(1 + \frac{\lambda z}{2\pi} \nabla^2 \phi\right) \tag{2.3}$$

The physical interpretation of this in terms of geometric optics is that in areas of sharp δ changes: light is refracted from high δ regions to low δ regions, so the increased distance to the detector allows the refracted light to highlight sharp changes in δ (and therefore density). An illustration shows this effect in figure 2.6, where X-rays come in from the left, go through some target, are refracted, and hit the detector strip on the right.

This can be modelled as simple refracting using Snell's law, where the real component of the refractive index ($n = 1 - \delta$) is the part used:

$$n_1 sin(\theta_1) = n_2 sin(\theta_2)$$

However, it is worth noting that Snell's law is based on a discontinuity in refractive index across a boundary. To refract light through materials with smoothly varying refractive index, either Snell's law must be applied repeatedly over small steps (like a numerical integration scheme), or a fully continuous equation of motion must be applied (as in Pediredla et. al. [67]). To apply Snell's law over small steps in a continuously varying material, the gradient of the refractive index can be used as the surface normal for calculating θ_1 and θ_2, the practical implementation of this as part of the GREENER code is discussed in section 4.2.

2.3.5 Phase Contrast Imaging - Applications in ICF

XPCI was first suggested for use in ICF as a way to characterise capsules [68], and this has allowed voids, fine features and imperfections within the ablator and ice layers to be detected [69]. Later, it was noted that since X-rays easily penetrate ICF plasmas, XPCI could be used for in-flight imaging [70].

From equation 2.3 it is clear that the phase effects in propagation based imaging are based on the laplacian of the phase, and that therefore XPCI is particularly sensitive to sharp gradients in refractive index. This is what enables features like interfaces, voids and imperfections to be observed, the sharp changes in density are highlighted, even although the absolute change in density may not be large. This suggests XPCI would be useful for imaging shocks, as the absolute change in density across a shock is limited to roughly a factor of 4, so absorption contrast is limited, but the stronger the shock, the sharper the jump in

density will be, so the laplacian of the phase can be much larger. A number of studies have been completed examining in-flight shocks in high energy density conditions [71, 72, 73].

As part of this effort, platforms for X-ray phase contrast imaging were developed at OMEGA [72] and NIF [74], with resolution as good as 1μm on OMEGA [75] and 5.8μm on NIF[16]. The platform on NIF was used to image early time shocks in-flight, with a temporal resolution of 100ps, and spatial resolution of 5.8μm. Late time shocks are more challenging to image due to the motion blur caused by the implosion collapsing, and the limited time window a shock takes to cross the shell due to the shell being compressed.

There have also been proposals to take this further by imaging the evolution of voids during an implosion [76].

Two of the challenges which apply to all diagnostics, are designing optimal experimental setups and interpreting results. These are particularly acute for XPCI, because the refraction can deflect the location of interfaces on detectors, making it harder to understand how an image relates to a target. Synthetic diagnostics are an effective way to resolve these challenges, as they allow many experimental setups or targets to be simulated, then either used to design an experiment or to be compared with results.

For pure absorption or pure phase objects, analytical solutions can be found in many cases, but real experiments have both, as well potentially complex target geometry. Even in the geometric optics limit, this requires computational simulations for synthetic diagnostics, this is explored further in section 3.1.

2.3.6 Self-emission Imaging

ICF implosions are often hot enough to emit a significant flux of X-rays, this is a useful diagnostic in its own right [52], and is regularly used to determine information about hotspot shape, volume and temperature, with time resolution [77]. X-ray self emission can also be used to learn about mixing, since any ablator mixed into the hotspot will increase the bremsstrahlung self emission [78]. Deliberate doping of layers in a capsule can be used to gain further information about hydrodynamic instabilities [79]. Advanced techniques like multi-monochromatic Imaging (MMI) can provide even more diagnostic power, MMI allows simultaneous spectral, 2D spatial and temporal resolution of an implosion at OMEGA[80, 81, 82], although this is yet to be implemented on NIF. With multiple MMI diagnostics, tomography can be performed to gain 3D spatial information. In addition, for indirect drive

the hohlraum emits X-rays, which can be used to determine the temperature of the hohlraum over time [83].

However, for the purposes of this thesis self emission from the hotspot is a source of noise which makes backlighter based X-ray phase contrast imaging more difficult. Particularly because the X-rays emitted by the hotspot are often in a similar range to X-ray backlighters. This detrimental effect can be limited by using monoenergetic backlighters and filtering detectors, and by choosing backlighters with energies which are far away from the peak of the expected self emission spectrum.

2.3.7 X-ray Sources

The ideal X-ray source for imaging would be a synchrotron or XFEL, unfortunately these are not co-located with ICF experiments, so alternatives must be found. ICF facilities do have high power lasers, which makes laser driven sources particularly attractive, and a range of sources have been developed [59].

The simplest laser driven sources are wires or foils, here a high power laser impinges on the wire or foil and forms a plasma from the outer layers of material. The plasma itself acts as a thermal source, emitting bremsstrahlung X-rays. In addition, hot electrons are accelerated in the plasma and impact the cold core of material, on impact they knock inner shell electons out of atoms. An electron at a higher energy level in the atom will drop in to replace the missing electron, releasing a photon in the process; these are called characteristic X-rays.

Since brehmstrahlung radiation is a thermal source, it's spectrum is that of a black or grey body, depending on how transmissive the source is. The maximum brightness is therefore dependent on the temperature and the source size only, and since small source sizes are desirable and X-ray energy choice matters, this is a significant limitation and drives the desire for non-thermal sources.

Characteristic X-rays have a spectrum based on the energy levels of the material being used, with different energy transitions producing different mono-energetic emission lines. The kα and kβ lines are the transitions from the L and M shells into the K shell, and the kα line is commonly used as an X-ray source [84]. Characeristic X-rays are non-thermal and can therefore be brighter than brehmstrahlung, they are also mono-energetic, which means that filters can be used to select only the characteristic X-ray, which allows background noise to be significantly reduced.

More advanced X-ray sources are also avaiable, of particular interest are laser wakefield accelerators (LWFA), which are non-thermal sources that offer synchrotron-like radiation from a tabletop source. These have been used for X-ray phase contrast imaging [85], and are continuously improving in level of control and automation [86].

2.4 Shock Diagnostics: VISAR

Section 2.3.5 described how XPCI can be used as a diagnostic to image early time shocks in implosions, and in chapter 6 of this thesis, an experimental design to examine the N+1 shock will be described. However, there are a few other diagnostics which can be used to understand shocks, in particular the Velocity Interferometer System for Any Reflector (VISAR) [87].

VISAR works by measuring the phase shift that results from optical light reflecting off a moving surface using interferometry, and it can be used to measure shocks travelling through the shell of an implosion capsule. For a more in depth description of VISAR, see Dolan [88].

To use this technique to observe shocks in an implosion, the diagnostic must be used on the inner surface of the capsule, since the ablation process forms an opaque plasma on the outer surface. This can be achieved using a gold cone embedded into the capsule shell, such that there is an optical path from the diagnostic to the inner surface of the capsule. This works well to image the early time shocks as the ice layer is transmissive, although D2 ice is used, so shock timing implosions are not perfect surrogates for DT implosions. However, there is a more significant problem in that after the main shock breakout occurs, the ice becomes a plasma and the diagnostic is no longer effective [89]. As a result the VISAR technique cannot be used to measure late time reflected shocks, such as the N+1 shock.

Chapter 3

Methods

3.1 Imaging Simulation Codes

To design experiments using X-ray phase contrast imaging (XPCI) as a diagnostic, and to interpret the results of XPCI experiments, it is useful to have codes which will produce synthetic images of the experiments. There are two main methods of generating synthetic images which will be described here: wavesolvers and raytracers

Wavesolvers treat light as a wave, and solve Fresnel-Kirchoff integral to determine how it propagates through a target. This method is not restricted to the geometric optics limit, and is more general than a raytracer, but does not parallelise well and requires large amounts of ram. Wavesolvers need to store the entire wavefront in memory all at once, then solve the full integral before propagating the wave, this makes it difficult to split the problem across multiple cpu's or nodes in a cluster.

Raytracers treat light as a particle, and generally only work in the geometric optics limit, they work by tracking a photon from the source to the detector as it travels through the target. This allows each ray or photon to be treated separately, so the problem can easily split between different cpu's or nodes in a cluster. In addition, if an object is particularly large or complex, it can be split into sections, with each section being raytraced on a different node, meaning memory usage need not be a limiting factor. Finally, raytracers allow for a whole object to be raytraced with coarse details, and then raytrace a feature of interest in much finer detail, allowing for reduced computational cost.

The GREENER code, discussed in chapter 4 is based in GEANT4, which is a monte-carlo based particle tracker, and is used as a raytracer. Chapter 4 contains much more detail on

this code. One of the difficulties with a raytracing approach is how to handle refraction in a continuous medium. This can be done in two main ways, either refracting at interfaces, or via an integration scheme based on the gradient of the refractive index [90]. These two methods are better suited different situations [91], however, as the primary concern of this project is picking out interfaces and sharp changes in density, surface based refraction was thought to be a better choice. Additionally, there was already a module written to handle surface based refraction in geant4[92, 93], making the choice even clearer. Nevertheless, importing and handling complex geometry from simulations was a non trivial task, and is covered in more detail in section 4.2.

There are a number of codes already written to produce synthetic X-ray images or to model X-ray optics. PhaseX [55] is a wavesolver written to produce synthetic X-ray phase contrast images for ICF experiments, it appears extremely useful, unfortunately it is closed source and was not available for use. SHADOW [94] is a raytracing code which is often used to model synchrotron optics, so is equiped to handle X-rays, and OASYS [95] is a toolkit which allows for easy use of SHADOW. These codes are very useful for modelling X-rays, and could likely have been used, however, are intended for optics with clearly defined surfaces and would require similar solutions to the ones described in section 4.2. Finally, Fluka [96] is another monte-carlo based particle tracker, and is similar to Geant4, although it has the advantage that a module has been written to handle diffractive effects using a scheme based on Huygens principle [97]. The choice to use Geant4 over Fluka was based on the extensive documentation available for Geant4, and the fact that the geometric optics approximation works well for the synthetic images of interest to GREENER.

3.2 Experiment Simulation Codes

To design experiments and analyise results, experiment simulation codes are needed, and these provide the inputs for the imaging simulation codes discussed in section 3.1 and chapter 4. Many codes exist for use in ICF to simulate different processes, using different approaches and approximations to obtain their results.

For full ICF implosion scale simulations, the type of codes which are used the most frequently are Radiation-Hydrodynamics codes. These codes treat the simulation as a hydrodynamics problem, using the Navier-Stokes equations, with radiation transport models

included. Within Radiation-Hydrodynamics codes, there are three main types: Eulerian, Lagrangian and Arbitrary-Lagrangian-Eulerian(ALE). Eulerian codes have a fixed grid, and the values of density, pressure, temperature, etc are defined at the points on the grid, then the values at each grid point are evolved with time using the relevant equations of motion. Lagrangian codes have a moving grid, where each cell defines a packet of material, then the grid deforms and moves according to the relevant equations of motion as the simulation evolves. ALE codes decouple the motion of the mesh from the material motion, and can allow the mesh to move arbitrarily, anywhere between being fixed in place like an Eulerian code, or moving with the material like a Lagrangian code.

Eulerian codes are inherently diffusive, as material crosses a cell boundary it is instantly considered to be spread across the whole cell; this makes it difficult to study shocks, which require strong gradients, using these codes [33]. The artificial diffusion in Eulerian codes can be mitigated by having many small cells in areas close to high gradients, and the code FLASH [98] does this by using adaptive mesh refinement to dynamically adjust the grid as the simulation progresses. Lagrangian codes do not have this problem, but instead struggle to accurately model mixing and in 2D or 3D simulations the grid can become tangled (see section 4.2.1 for more details). Tangling is a major problem for Lagrangian codes, which ALE codes seek to circumvent. ALE codes, such as Odin, often use a Lagrangian mesh for as long as possible, then once the mesh becomes tangled or too much deformation occurs, they remap the grid onto a new mesh and continue as a lagrangian code[99]. However, ALE codes are more general and can allow the mesh to move independently of the material; in areas of the mesh where it moves with the material, ALE codes act like Lagrangian codes, and in areas where the mesh moves differently from the material, ALE codes act more like Eulerian codes [100].

Radiation-Hydrodynamics codes struggle with one feature in particular: electron thermal conduction, often referred to as transport. There are a few different transport models, and the Spitzer-Harm model [101] is commonly used in radiation-hydrodynamics codes, it states that the thermal transport is proportional to the gradient of the temperature. This model is a diffusion approach, built on the assumption that the electron mean free path is much smaller than the temperature length scales. This means it works well for small temperature gradients, but if temperature gradients are large enough to make the electron temperature length scale smaller than the electron mean free path, the model breaks down. In the ICF

context, particularly around the ablation front, the temperature gradients can be very large, limiting the validity of the Spitzer-Harm model [102]. In computational terms, large gradients in the Spitzer-Harm model lead to unphysical transport: gradients can get so large that more energy is transported between one cell and the next than could be transported if every electron simultaneously moved cell. Every electron simultaneously moving from one cell to the next is equivalent to the free streaming limit, and this limit can be used to mitigate the shortcomings of the Spitzer-Harm model with the help of a flux limiter. The flux limiter sets the maximum allowed thermal transport to some fraction of the free streaming limit, a fraction of 0.1 is typically appropriate for electrons in the ICF context [103].

A more robust treatment of the transport problem would use either a Particle in Cell (PIC) or Vlasov-Fokker-Planck (VFP) Code. PIC codes are useful for simulating some laser plasma interactions, one prominent example is EPOCH [104]. These codes work by tracking the motion of particles moving through a grid of cells, then using the distribution of particle locations to update the electromagnetic fields on the grid, before moving the particles according to the newly updated fields. Vlasov-Fokker-Planck (VFP) codes, such as FASTVFP [105], work by solving the VFP equation, and accurately describe non-local energy transport, which occurs in the large temperature gradients present in ICF conditions [105]. Both PIC and VFP codes are computationally expensive, and while useful for modelling parts of ICF implosions, they are too computationally expensive to model an entire implosion. They can be used to benchmark approximations like the flux limiter, and have been used to determine that an electron flux limiter of 0.1 is a reasonable approximation.

3.2.1 HYADES

HYADES [106, 107] is a 1D Lagrangian Radiation-Hydrodynamics code, and was chosen as the code for use in the work in this thesis. A 1D code was chosen for its computational efficiency, the results in chapters 5 and 6 involved thousands of simulations, which was enabled by the use of a 1D code. While a 1D code cannot model 2D effects like Rayleigh-Taylor growth, their likely impact on an experiment can be estimated using measurements like the in-flight aspect ratio (IFAR), as discussed in section 2.2.4. A Lagrangian code was chosen as it models shocks more accurately and shocks were a key component of the research. And finally, HYADES was a particularly good choice as it is designed specifically for simulating high energy density experiments, and can simulate planar, cylindrical or spherical geometry.

The code models three fluids: elecrons, ions and radiation. Ions and electrons are assumed to be in local thermodynamic equilibrium, and thermal transport is modelled using the Spitzer-Harm model with a flux limiter. In HYADES the flux limiter can either use the classical method where the transport is taken as the minimum of the Spitzer-Harm model and the flux limited free streaming transport. In HYADES, for electrons the flux limited free streaming thermal transport is taken as:

$$q_e^{fs} = f_e^{fs} n_e k_b T_e v_e$$

where f_e^{fs} is the flux limiter, n_e is the electron density, k_b is boltzmans constant T_e is the electron temperature and v_e is the electron thermal velocity. More details are in documents [108] provided by Cascade Applied Sciences inc.

HYADES underpredicts yield when compared to experiments, comparisons with other codes (eg LILAC) which don't underpredict yield suggests that the is caused by the hotspot in HYADES simulations ablating too much DT ice into the hotspot [109]. This results in a cooler denser hotspot, lowering yield. To compensate for these issues, the electron flux limiter can be lowered at late time in an implosion, to reduce the transport and therefore reduce the ablation of the DT ice, bringing the yield of HYADES simulations closer to experiments. In the simulations presented in this thesis, the electron flux limiter was chosen to be 0.1 for most of the simulation time, then lowered to 0.05 close to bang time. These values were chosen to best match the "baseline simulation" discussed in section 5.5 to the experimental yield of NIF's N210808 shot. The ion flux limiter was left as the HYADES default value of 0.4.

There are two models available to handle the radiation field in HYADES, the simplest is the 'grey' model, which assumes the radiation follows a planckian spectrum and therefore simplifies the radiation transport. However, the model used for the simulations in this thesis is the 'multigroup' model, which treats the spectrum as a number of weighted bins, and transports them separately with their own opacities/emissivities. This allows non planckian spectra to be used, which allows an M-band fraction to be included. The M-band fraction is the portion of the radiation impinging on the capsule which is emitted from the M-band part of the spectrum of the gold hohlraum, and experimental values from the N210808 shot are used.

For alpha particle transport, HYADES uses a diffusive model, this is a good approximation in the laser fusion context for plasma temperatures below 30keV [110].

As discussed in section 2.1.4, hydrodynamics equations require an equation of state to close the system of equations. HYADES can use precomputed tables or inline calculations (eg the ideal gas law) as the equation of state. In this thesis DT was modelled using qeos [111], while the high density carbon and helium gas fill were both modelled using precomputed tables from the LEOS and SESAME datasets respectively, which were provided by the HYADES developer.

3.3 Betatron Spectrum

For the synthetic images in chapter 4, the spectrum of a betatron source was required. This was done using a code provided by Robbie Scott, which numerically solves the equations for a relativistic electron in instantaneously circular motion provided in Jackson [112].

Jackson equation 14.76

$$\eta(\theta) = \frac{\rho\omega}{3c}(\frac{1}{\gamma^2} + \theta^2)^{1.5}$$

and Jackson equation 14.79

$$\frac{d^2I}{d\omega d\Omega} = \frac{e^2}{3\pi^2 c}\left(\frac{\omega\rho}{c}\right)^2\left(\frac{1}{\gamma^2}+\theta^2\right)^2 K_{2/3}(\eta(\theta))^2 + \frac{\theta^2}{1/\gamma^2+\theta^2}K_{1/3}(\eta(\theta))^2$$

where ρ is density, γ is the lorentz factor for the electrons, ω is frequency, θ is the angle away from the laser wakefield accelerator axis, c is the speed of light, e is the electron charge and K are bessel functions.

The lorentz factor was calculated using electron energies from Fourmaux et. al. [113] and the density was calculated by rearranging the equation for the critical frequency (Jackson 14.81):

$$\omega_c = \frac{3}{2}\gamma^3\left(\frac{c}{\rho}\right)$$

The on axis values (when $\theta = 0$) were calculated to give a spectrum of relative intensity against photon energy. The spectrum was then calibrated against experimental data when absolute photon counts were required. This is noted when relevant.

Chapter 4

GREENER

4.1 Chapter Overview

This chapter will describe the details of the GREENER code, which was written for this thesis by the author and can be used to generate synthetic X-ray phase contrast images of fusion relevant targets. The code can also handle arbitrary complex 3D meshes, although there are some outstanding issues with this, which are detailed later. The structure of the code is shown in figure 4.1, most of the GREENER code is built within GEANT4 [93], a robust and well tested monte-carlo particle tracker and physics simulation package, which can be used as a physics based raytracer; this has the advantage that there are many physics packages and extensions which can be used to include new physics as required. To handle X-ray refraction: GREENER uses an extension developed by researchers in Trieste as part of the PEPI project [114], which was in turn based on previous work by Wang et. al. [115] this extension was intended for edge-illumination XPCI, but also works for propagation based XPCI, which is how it is used in GREENER.

The opacity values are taken from GEANT4's livermore physics package, and the real component of the refractive index is taken from CXRO, which is based on [116]. In this chapter, the geometry construction which was written for this thesis by the author is explained, and the code is applied to a number of test problems, demonstrating how the code can be used to design and analyse experiments.

The GREENER code allows for varying levels of complexity in target geometries, however, wherever possible the simplest geometry is used:

- For 1D Lagrangian simulations of implosions, the geometry can be modelled as layered

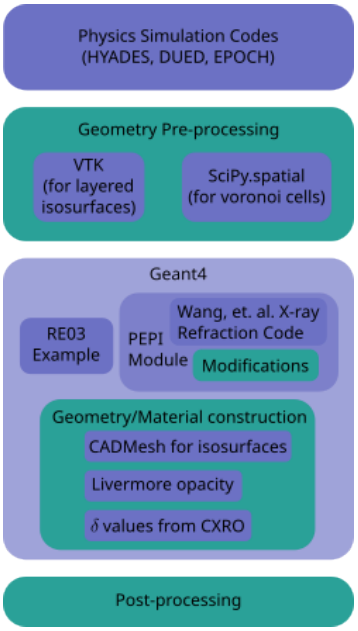


Figure 4.1: A diagram showing the different components of the GREENER code, starting with the output of a physics simulation, preprocessing it, using the geometry in the GREENER code, then post-processing the results. The authors contributions are shown in green, while modules written by others are shown in purple.

concentric spherical shells. Examples of this method being applied are given in section 4.2, and an example convergence test is provided in appendix C.

- For 2D Lagrangian meshes in cylindrical coordinates, the geometry can be modelled as rotated polycones, this is described in more detail in section 4.2.1, with an example application provided in section 4.4.

- For fully 3D Eulerian meshes, the CADMesh plugin [117] is used to import arbitrary layered surfaces, this is described in detail in in section 4.2.2, with an example application provided in section 4.5.

4.2 Geometry Construction Methods

The extension GREENER uses to handle X-ray refraction operates by applying snell's law as photons cross boundaries and surfaces between objects defined in Geant4. Geant4 builds objects in the simulations out of logical blocks (eg spheres, cubes, rings), and because snell's law is applied at boundaries, this works well for simple objects constructed as solid logical blocks. An example of an image of a simple sphere is shown in figure 4.2, this sphere was

constructed form GEANT4 builtin geometries, and backlit by 4.5keV X-rays. The parameters were chosen for illustrative purposes: the source was 1 metre behind the target, the detector was 1 metre ahead, and the sphere is a 1mm diameter sphere of artificially low density CH (ρ=0.5g/cm^3). This construction can be extended further to handle layered capsules, so long as spherical symmetry can be assumed, by constructing them from concentric spherical shells.

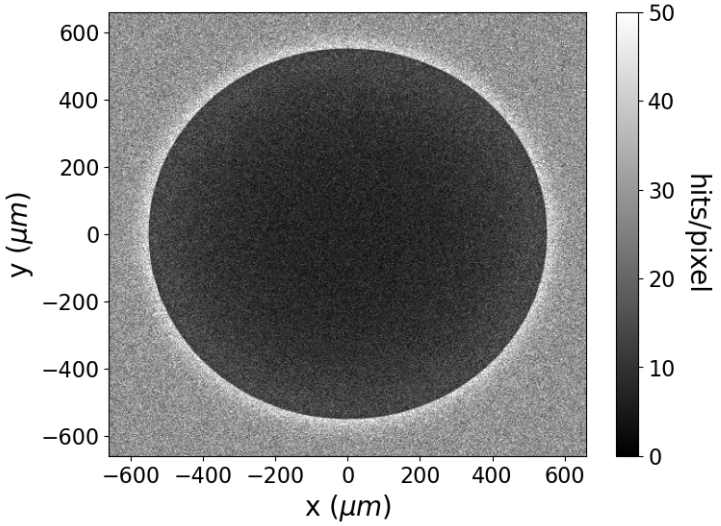


Figure 4.2: Synthetic image of a sphere (radius 0.5mm) constructed as a simple GEANT4 solid, backlit with fully columnated 4.5keV X-rays, simulated with 10 million rays.

Complex geometries which cannot easily be constructed from simple objects are more difficult, which is particularly relevant if the code is to be used to generate images from the output of multi-dimensional simulation codes. In particular, care must be taken to ensure that the surfaces of the objects constructed are aligned with the true surfaces (or gradients of refractive index) within the simulation, so that when geant4 refracts a ray across the surface, it does so at the correct angle. The naive approach to imaging the output of multi-dimensional simulation codes would be to simply build the complex objects from voxels (essentially small boxes, akin to 3D pixels), figure 4.3 shows the phase contrast image that results from a "sphere" constructed with this method. The left panel shows a low resolution grid, and the right shows a higher resolution grid. The voxel like structures are clearly visible, and even with higher resolution constructions, artificial structure is always introduced by the voxels. These grid artefacts are structured in nature, so cannot simply be averaged or smoothed away (eg by taking a radial integral).

Figure 4.4 clarifies the problem, on the far left of the figure a ray is shown travelling

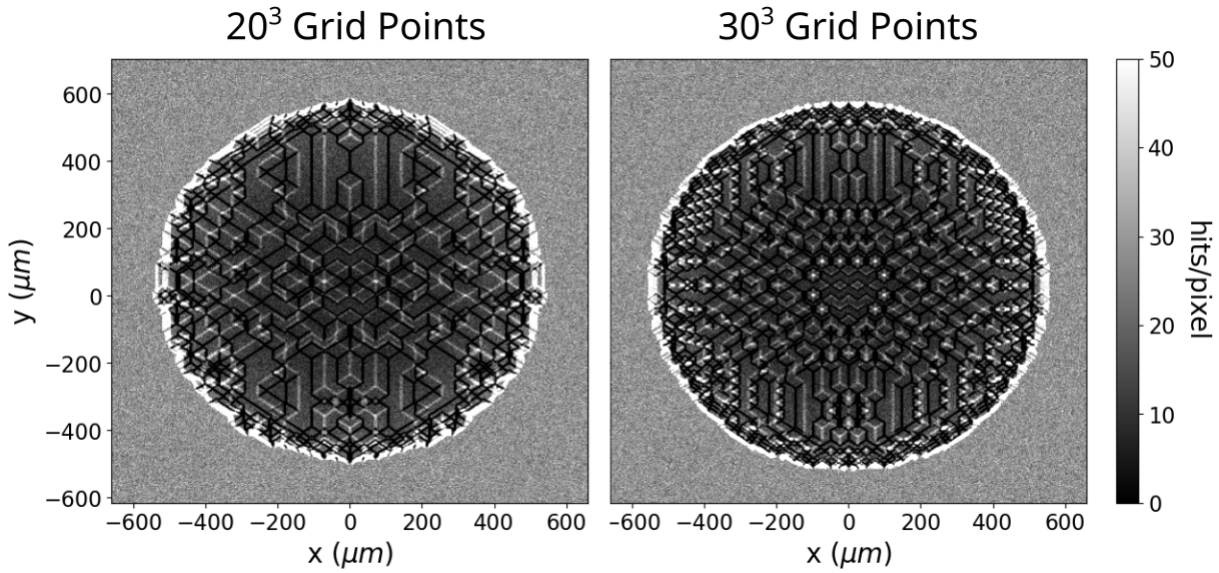


Figure 4.3: Synthetic image of two spheres (radius 0.5mm) constructed from voxels, backlit with fully columnated 4.5keV X-rays, simulated with 10 million rays. Left has 20^3 grid points, right has 30^3 grid ponits

through a material of changing refractive index, as it is parallel to the gradient it does not refract. The three diagrams on the right show what happens when the smooth gradient is discretised onto grids with different angles. The two rightmost diagrams show how artificial refraction is introduced. This effect requires that surface angles are preserved when constructing the target.

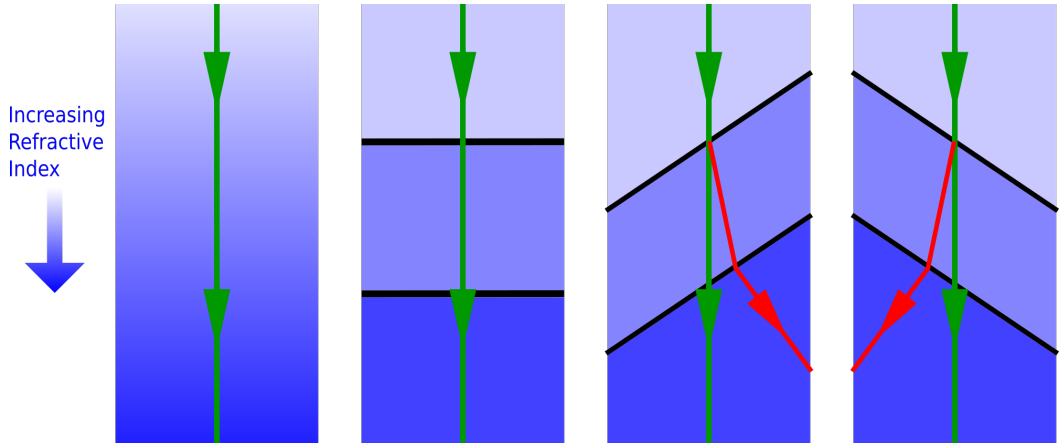


Figure 4.4: Diagram of a ray travelling through a refractive index gradient, showing how different discretisation grids give different refracted rays.

There are a number of solutions to this problem, which depend on the nature of the particular geometry being modelled. For a 1D Lagrangian code of a simulation, each Lagrangian

cell represents a packet of material which can be treated as a spherical shell, this sort of geometry can be constructed directly in geant4 as many concentric spherical shells, examples are provided in section 4.3. A 2D Lagrangian code of a cylindrically symmetric geometry is more complex, but can still be built with geant4 built in tools, this is done using the voronoi method described in section 4.2.1, with an example application is shown in section 4.4. Alternatively, for fully 3D objects, or for Eulerian meshes, a method based on finding surfaces of constant refractive index can be used, this is described in more detail in section 4.2.2, and an example application is shown in section 4.5.

4.2.1 Voronoi Method

For Lagrangian simulations, each point on the grid represents a packet of material which moves as the simulation progresses, this means the geometry of the grid points correspond to the physical geometry much better than an Eulerian simulation where the grid points are fixed and potentially arbitrary. As a result, in Lagrangian simulations, the boundaries of the zone represented by the packet around a grid point are good proxies for the interfaces which rays would be expected to refract off in the physical object which the simulation represents. These boundaries can be found from a set of points in a simulation using Voronoi diagrams, and then used to construct the geometry in GEANT4. In the case of 2D cylindrical systems, this is done with polycones.

A Lagrangian grid can be defined such that the points representing packets of material make up the nodes in a mesh, or are at the cell centres of a grid of cells. Since the cell centres can be used as the nodes to define a new mesh, these constructions can be interchanged for the purposes of defining the boundaries of packets of materials. For an untangled mesh, in either case, the boundaries of a packet of material can be found using the points on the mesh (either the connected nodes, or by connecting the cell centres). Unfortunately any mesh tangling will cause problems for this simple construction; examples of mesh tangling mechanisms are shown in figure 4.5, where the bowtie and overlap tangling are shown. These disrupt the geometry of the mesh and make it impossible to construct sensible zones around points using only the mesh.

If the mesh is tangled, untangling it is non-trivial, this is a problem for hydrocodes, but since in this case the mesh doesn't need to be able to be used in hydrocodes afterwards, the problem is simpler. The mesh doesn't need to be a regular grid with sensible connectedness,

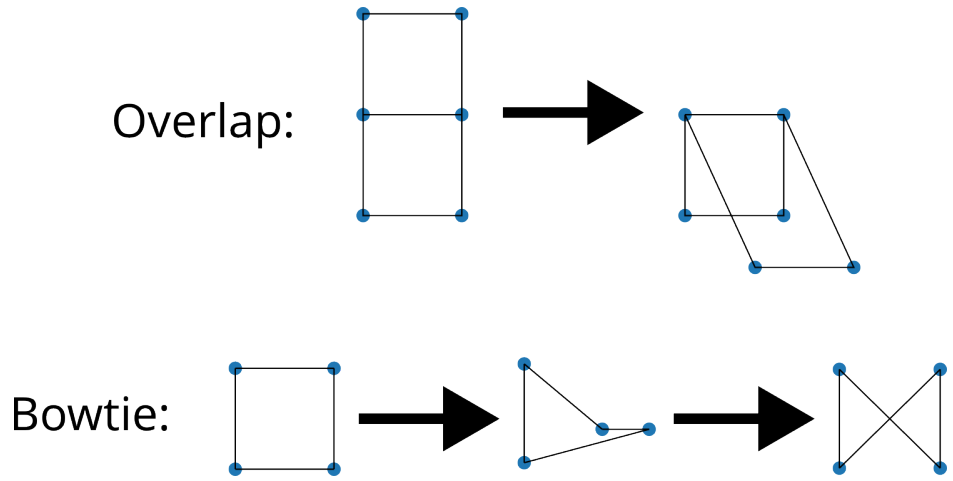


Figure 4.5: Two tangling mechansims in a Lagrangian mesh: Overlap, where a zone overlaps another zone, and bowtie, where the edges of a zone cross each other.

and the existing structure of connections does not need to be preserved, the new mesh just needs to have a zone around a point, and that zone can be any arbitrary shape. As a result of these looser requirements, rather than detangling the mesh, a fully new mesh can be constructed in the form of a voronoi diagram. A voronoi diagram defines a zone around each grid point, as the set of points in space which are closer to the grid point than they are to any other grid point. This is a simple and repeatable way to find boundaries for the cells, it works whether the grid points are defined at the nodes or cell centres, and there are many pre-written codes to do this. In GREENER: `scipy.spatial`[118] was used.

To demonstrate this method, a test sphere was constructed as if it were a 2D cylindrically symmetric Lagrangian simulation, figure 4.6 shows the grid and voronoi diagram for a low resolution test sphere. Here the Lagrangian points are shown as blue points, in this mesh the material properties are defined at these nodes, so these blue points are taken to be the centre of the Lagrangian packets of material. Voronoi cells can then be constructed around each packet of material, the boundaries of these cells are shown as black lines, with dotted lines for the outermost layer of cells which cannot be closed with this method. The orange points show the corners of the voronoi cell boundaries, these are used to define the corners of the cells when importing them in geant4.

For a 2D simulation of a cylindrically symetric object, the voronoi cells can be rotated around the central axis to construct the volume associated with each cell. This can be done using the GEANT4 Polycone constructer, which takes any arbitrary set of points which make up the boundaries of a 2D cell and builds the corresponding 3D rotated volume in GEANT4.

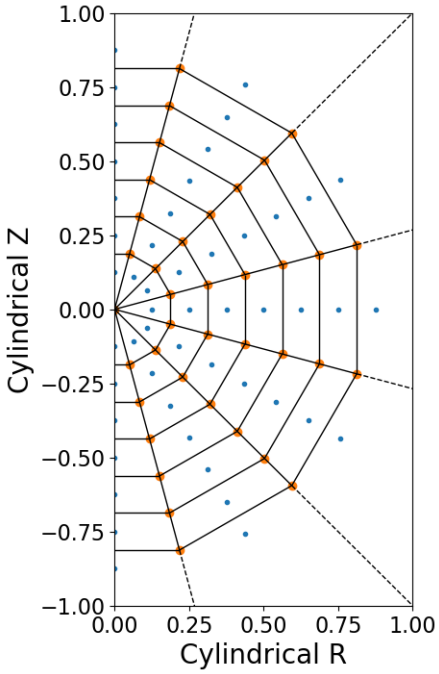


Figure 4.6: Voronoi diagram for a sphere, constructed in cylindrical coordinates. Blue points show the grid points used, black lines show the edges of each cell, and orange points show the grid points of the new mesh, which can be used to construct the geometry in GEANT4. Black dotted lines show edges of cells which are unbounded, these must either be closed or discarded.

Figure 4.7 shows two spheres constructed with the voronoi method, it is clear that increased resolution in the grid gives better results, unlike the voxelised spheres. Though there are still some artifacts, this is because discretising a sphere into flat surfaces is always inherently an approximation and means there are parts of surfaces which are shifted or angled incorrectly, and there are sharp corners which cause sudden shifts in the direction of refraction.

There are also problems when applying this to real simulations in that the surfaces generated from a grid of points may not be correctly aligned with the true surfaces, and there are sometimes back angled surfaces that go the wrong way. Figure 4.8 shows an example of this, where a grid of points is being shocked from below with a curved shock. The points have been arranged to force a back angled surface to occur at the middle pair of points in the shock.

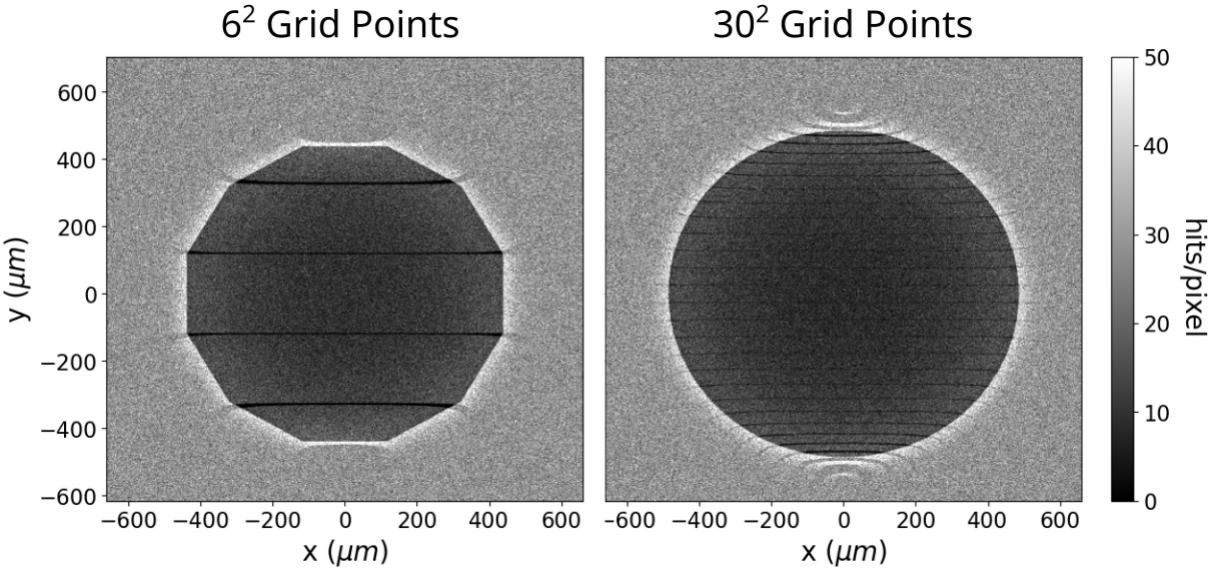


Figure 4.7: Synthetic images of two spheres (radius 0.5mm) each constructed from a voronoi diagram generated from a Lagrangian grid, backlit with fully columnated 4.5keV X-rays, simulated with 10 million rays. Left has 6^2 grid points in the underlying 2D mesh, while right has 30^2.

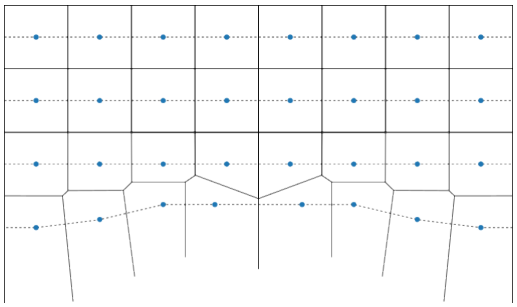


Figure 4.8: A Voronoi diagram showing a grid of points experiencing a curved shocked from below. The blue dots represent the points on the mesh, with dashed lines showing the true curvature of surfaces and solid lines representing the voronoi surfaces. The central points of the shock have been deliberately spread in such a way that the voronoi surfaces become angled against the direction of curvature.

4.2.2 Isosurface Method

The most general method for handling complex geometry in geant4 is to build it from layered "isosurfaces" (surfaces of constant refractive index). This is similar to the 1D layered sphere method, except that instead of spheres, the layers can take any shape. This involves finding many layered surfaces, where each surface is constructed from the set of points with equal refractive index, and contains all the points with higher refractive index, which means a surface could take the form of multiple isolated blobs, a single solid, a hollow solid, or any arbitrary set of shapes. These surfaces are by definition perpindicular to the gradient of the refractive index, and therefore the correct angle will be used in snells law calculations as a ray goes through the surface. This method is particularly useful for Eulerian grids, where the grid cells are often not related to the physical geometry of the object, and finding the correct surfaces to refract off is non-trivial, or 3D Lagrangian meshes where there are no simple geant4 geometries to construct cells from.

To find the required isosurfaces, the software package VTK [119] was used, this package uses the marching cubes algorithm to extract surface meshes from arbitrary grids. The grids used in rad-hydro codes are constructed in the VTK framework, and then the surfaces are extracted as `.stl` files, which can be imported into GEANT4 using the CADMESH package [117]. Example images of spheres constructed from stl surfaces are shown in figure 4.9. Here an eularian grid was constructed, with density decreasing with distance from a central point, an isosurface contour was found at radius=0.5mm, and exported as an stl file. The stl file was then used to construct a solid sphere with density 0.5g/cm^3. Two spheres are shown, one with low grid resolution, the other with high grid resolution, comparable to the voxelised spheres shown in figure 4.3. It is clear that in the isosurface method, increasing the grid resolution gives resutls which more closely approximate the pure sphere shown in figure 4.2, so this is an improvement over the voxelisation method. The isosurface method also gives good results at lower resolution, the low resolution grid in figure 4.9 is much lower resolution than the low resolution grid in figure 4.3, yet it preserves the qualitative properties of the pure sphere more closely.

The isosurface method does have some downsides, some artefacts are still present even in the high resolution image in figure 4.9, for example: there are extra bright rings which come from surfaces which aren't quite aligned correctly. There are also problems that result from poor mesh quality, in complex simulations the surfaces produced might not be properly closed

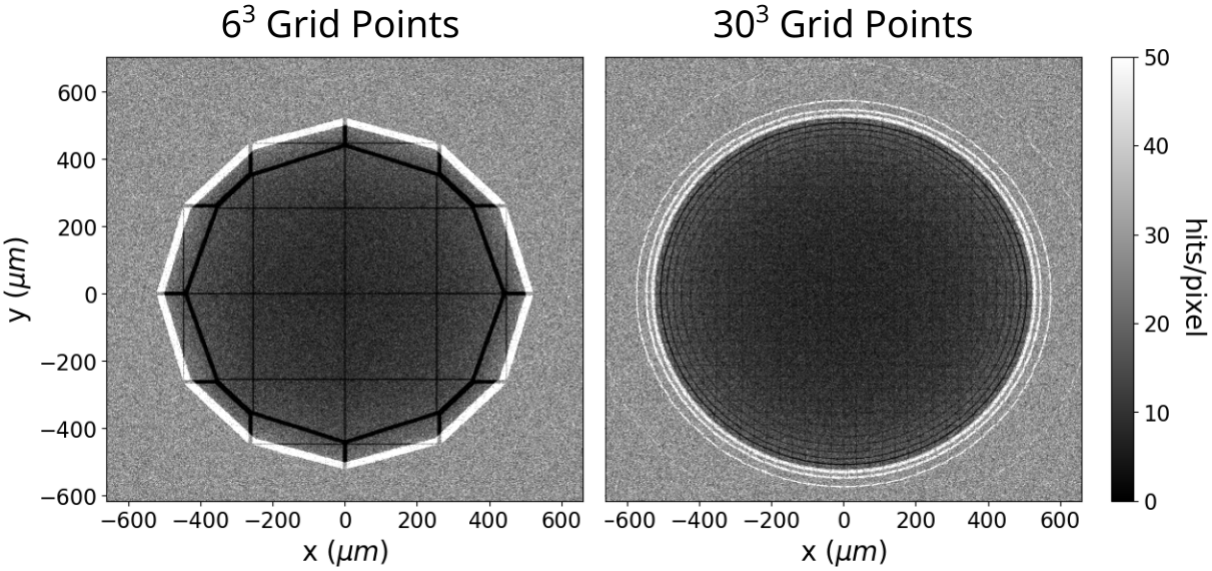


Figure 4.9: Synthetic images of two spheres (radius 0.5mm) each constructed from an isosurface generated from an Eulerian grid, backlit with fully columnated 4.5keV X-rays, simulated with 10 million rays. Left has 6^3 grid points, right has 30^3 grid ponits

(ie not "watertight"). In this case, a particle can lose track of which object or layer it is in, and as a result the particle will stop resolving the object and, instead, take a straight line path to the detector, giving an image where the object appears to have holes in it. Fortunately VTK provides algorithms for fixing many types of holes in surfaces and checking they are watertight, although VTK does not fix all meshing defects, so care should be taken to ensure there are no geometry issues remaining after fixing meshes.

4.3 1D Problems: Spherical Capsules

In this section, the principles behind imaging the results of 1D Lagrangian codes are demonstrated and benchmarked against experiment. This is the simplest geometry to image that will be shown, it is simply constructed from layered concentric spheres, and geant4 has built in methods to construct spheres. To test the code, a static capsule with a monochromatic backlighter offers the cleanest comparison. This is because the geometry of a static capsule can be known in much more detail than any target which is undergoing an implosion.

In work by Dewald et al. [74] a static capsule was imaged with a 5μm 7.8keV monoenergetic backlighter. The authors confirmed that the backlighter was 15mm back from the capsule, with the detector 670mm forward from the capsule, contrary to the paper [120].

The capsule consists of a hollow 66μm thick shell of 3.5g/cm^3 high density carbon, with inner radius 844μm, coated in a 50μm thick layer of 1.2g/cm^3 CH plastic, with inner radius 910μm. The pixel size was 9μmx9μm, with an additional 25% magnification through the streak camera, so to simulate this a pixel size of 7.2μm was used in GREENER. Resolution on the object plane was limited by the backlighter to 5μm. Figure 4.10 shows a comparison of the experimental lineout data to the results of a GREENER simulation. Here we see good agreement between simulation and experiment, except in the region where a tungsten wire was present in the experiment, since the wire was not included in the simulation. The simulation used 200 million rays, which took 10 minutes on 64 cores.

This setup is particularly useful as it is the same setup as was used later for measurements of in-flight shocks [16], so in addition to being a direct comparison of code to experiment, it strongly supports the use of GREENER in the experimental design shown in chapter 6, which proposes using the same in-flight shock measuring setup to image the N+1 shock.

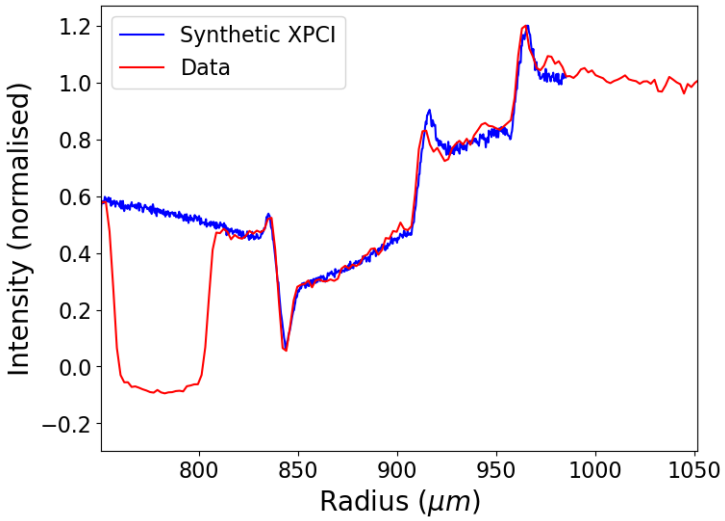


Figure 4.10: A lineout of the capsule imaged in Dewald et al. [74], compared to a synthetic XPCI lineout produced by the GREENER code. Note that the synthetic code does not include the tungsten wire at the left edge of the experimental data.

Another test case of a simple static capsule can be found in a paper by Fourmaux et al.[121], here a capsule from LMJ was imaged using a betatron source. This is more difficult to simulate since the exact spectrum of the source is not known, although an approximate spectrum can be generated based on the data in the companion paper [113]. The source was treated as a betatron source with critical energy of 15.1keV and an electron energy of 340 MeV, the spectrum was calculated using the method in section 3.3, and is shown in figure

4.11. Using this spectrum, a simulation was run with the source of size 1.7μm placed 75cm behind the capsule, backlighting it onto a detector placed 280cm ahead of the capsule, with 20μm pixels on the detector. The parameters were chosen to match the experiment, and in both the experiment and simulation, resolution was limited by the detector pixel size, limiting the resolution to 4.2μm in the object plane.

The capsule details are shown in table 1 of Fourmaux et al. [121], capsule 1 was used, which has a total outer diameter of 2049.3μm; the capsule consists of an innermost 11.4μm layer of CH, followed by 47.6μm of Germanium doped(0.44%) CH, followed by 10.8μm of Germanium doped(0.29%) CH, and an outermost layer of 11.9μm of CH. A comparison between the experimental image from the paper and the synthetic X-ray phase contrast image generated with GREENER is shown on the left side of figure 4.12, while a comparison with an absorption only synthetic image (equivalent to contact imaging) is shown on the right side of the figure. From these comparisons it is clear that the synthetic XPCI is much closer to the experiment than the absorption only image. A comparison of the lineouts is shown in figure 4.13, here it is still clear that refractive effects are necessary to explain the image, but the agreement isn't perfect, this is likely because the exact spectrum of the betatron source is not known, and a more accurate spectrum would likely give better agreement. Since the fringes are less extreme in the synthetic diagnostic than they are in the experiment, this suggests the spectrum used was harder than the experiment. The true spectrum could be found in a more systematic way by using the GREENER code to simulate many mono-chromatic images and superimposing them with different weights, this would allow the weights most closely matching the experimental data to be found, and therefore the experimental spectrum.

The experimental image used in both panels of figure 4.12 and the lineout used in figure 4.13 were both extracted from the paper [121], and reproduced with permission.

4.4 2D Problem: Shocked Cylinder

This section will apply the voronoi method to simulations of an experiment performed on the Target Area West chamber at the UK's Central Laser Facility. The experiment featured a cylinder of CH foam, which was shocked from one end of the cylinder by a laser induced shock. The cylinder also had layer of tungsten on the end of the cylinder to prevent electrons accelerated by the laser solid interaction from penetrating into the foam. A schematic of

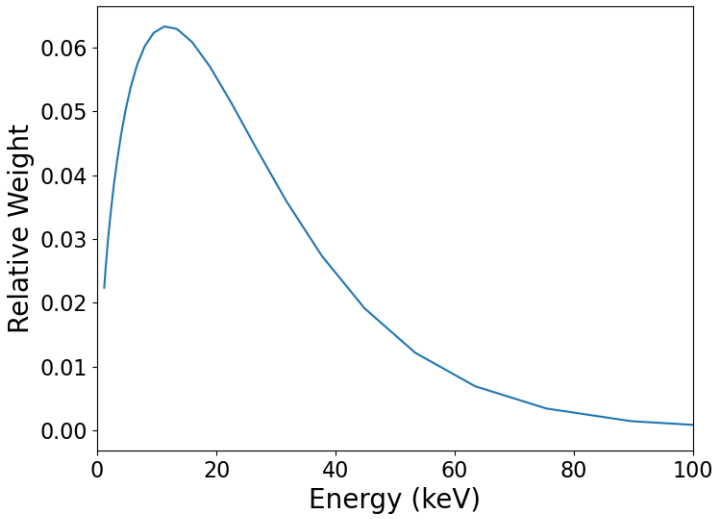


Figure 4.11: The weighted spectrum used to approximate the betatron source in Fourmaux et al. [121] using the method outlined in section 3.3.

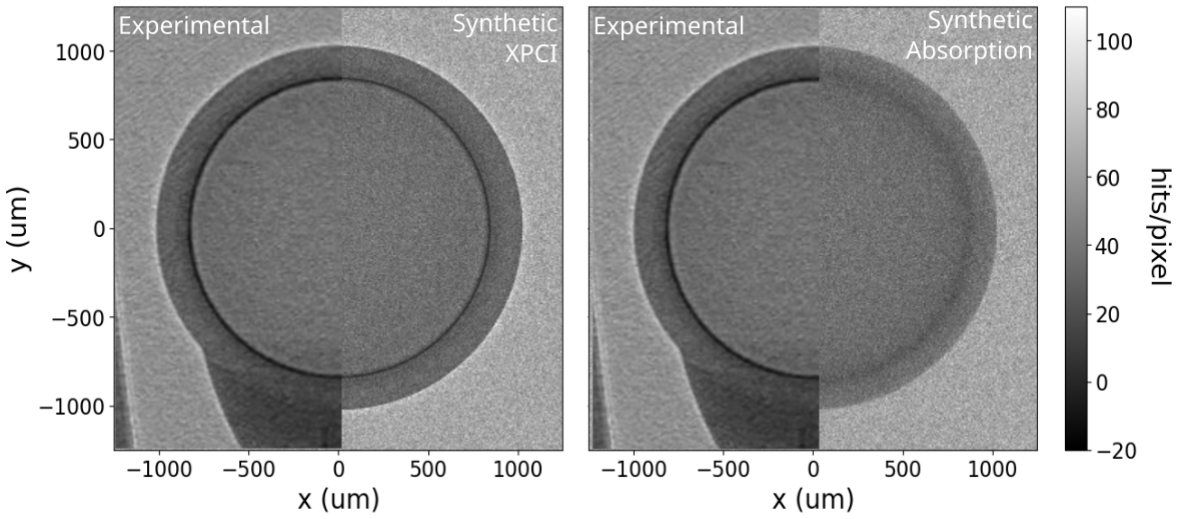


Figure 4.12: A comparison of the experimental image for capsule 1 in Fourmaux et al.[121], to synthetic images produced by GREENER. The left synthetic image has refractive effects enabled, while the right synthetic image is an absorption only image.

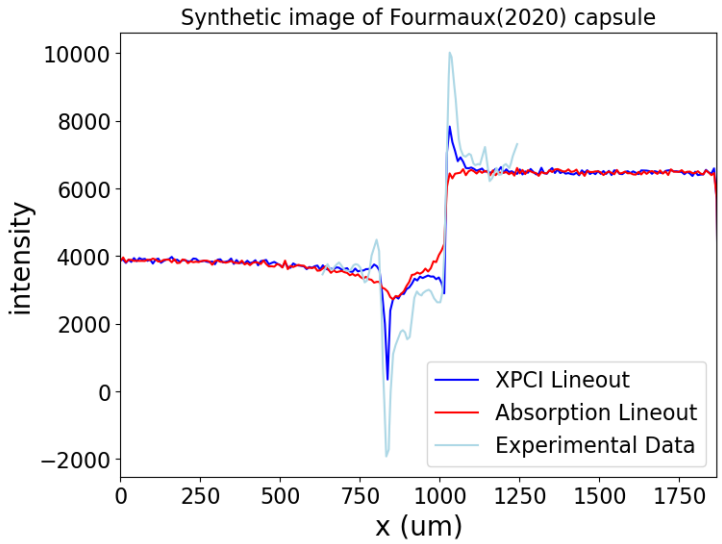


Figure 4.13: A comparison of the lineout for capsule 1 in Fourmaux et al.[121], to synthetic lineouts produced by GREENER.

the target is shown in figure 4.14. The shock was then imaged by a laser driven copper backlighter, producing a 5μm source at copper's kα line (8keV), the source was placed 5cm behind from the target, with a detector placed 70cm ahead of the target. This configuration was chosen to allow for phase contrast effects to occur.

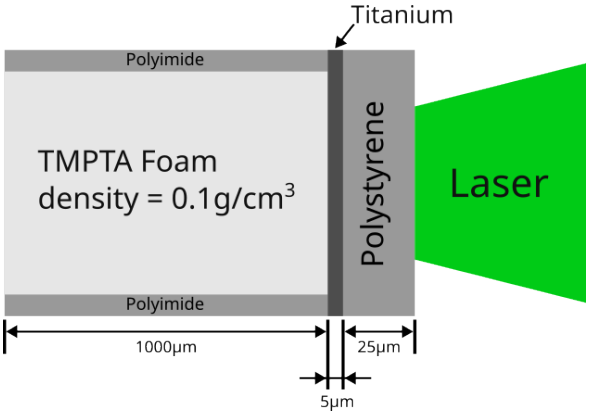


Figure 4.14: Schematic of a cross section of the cylindrical target used, showing the direction of laser incidence.

A number of shots were performed, and the laser parameters used to launch the shock were intended to be 550J, in a 400μm spot for 2 ns. Unfortunately due to misalignments in the laser chain, less energy was delivered, in addition, on shot 69 the phase plate was removed to increase the energy meaning the focal spot size is also unknown. This shot was the shot with the highest energy and strongest shock, and therefore has the clearest X-ray

phase contrast effect. DUED simulations of this shot were performed by Luca Antonelli with a range of estimates of these values, the closest fitting simulation had a laser energy of 240J over 2ns and a focal spot size of 400μm, so that simulation will be used in this section. On the shot being analysed, the backlighter was fired 25ns after the shock launch, however, due to the uncertainties in the laser parameters, the simulation timestep used does not match this time, instead the timestep chosen was chosen to be the step which most closely matched the experimental image.

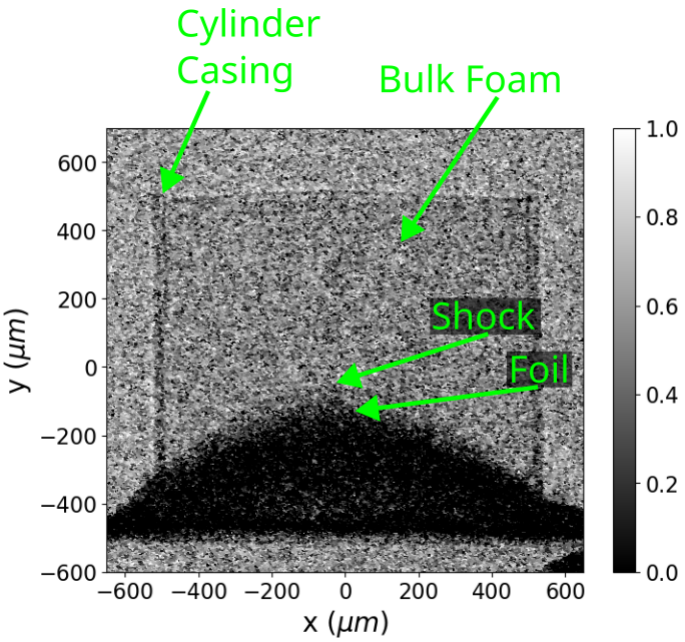


Figure 4.15: The experimental image produced on shot 69, with key features labeled.

The experimental data from shot 69 is shown in figure 4.15, where the colorbar was normalised and optimised for contrast. The cylinder and foam are clearly visible, and the foil can be seen to have been deformed by the laser driven shock; in addition, there is a shock travelling through the foam ahead of the metal foil, which is labelled for ease of identification.

Figure 4.16 shows the synthetic image along with the density profile from the simulation used to generate it. In figure 4.17, the synthetic image is compared to the experimental data, showing qualitative similarities: in particular, the contrast on the shock appears similar, and the transmission of the metal foil appears similar.

Although quantitative comparisons are limited by the difficulty in matching hydro simulations to the experiment, the synthetic image provides a good demonstration of the capabilities of the GREENER code, showing that 2D Lagrangian density profiles can be used to generate

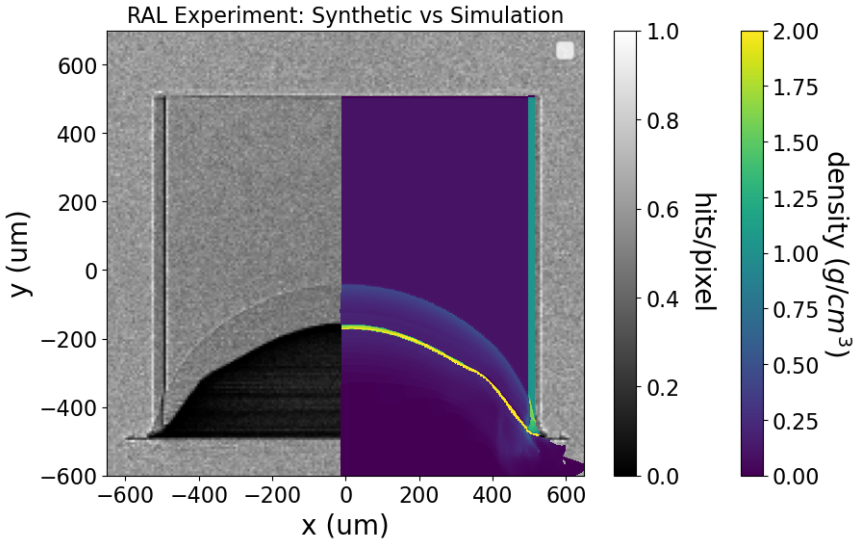


Figure 4.16: A synthetic image of a simulation of the experiment (left), and the 2D DUED density profile used to produce the synthetic image (right).

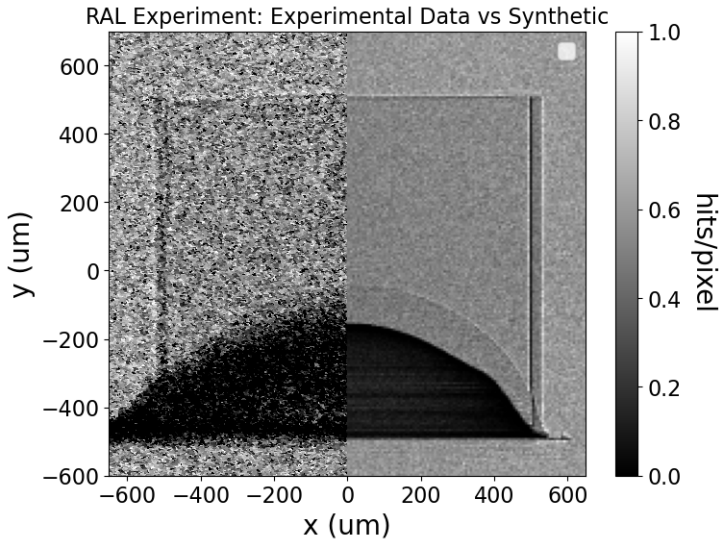


Figure 4.17: The experimental image produced on shot 69 (left) compared to the synthetic image produced from a simulation (right).

synthetic images.

4.5 3D Problem: Relativistic Transparency Experiment

In this section the isosurface method will be applied to a 3D Eulerain mesh, which was needed as part of work done to design an experiment for a proposal to study relativistic transparency.

Relativistic transparency occurs when a sufficiently intense laser impinges on a plasma, such that the electrons become relativistic, and their mass increases. If the masses of the electrons are increased enough, the plasma frequency in an otherwise overdense plasma can be lowered below the laser frequency, at which point a previously opaque plasma becomes transparent [122]. This effect allows an intense laser to tunnel through a plasma more effectively than would otherwise be expected, and the experiment proposed was intended to observe the tunnelling created by this effect. A discussion could be had about whether the physical mechanism for the tunnelling is relativistic transparency or hole boring, but for the purpose of this thesis the relevant fact is that some tunnelling mechanism occurs.

EPOCH simulations were performed by Kavin Tangtartharakul, the simulations are shown in figure 4.18 and were set up to match the planned experiment, with a laser impinging on a block of foam from the left of the figure. Three simulations are shown, each with a different foam density, and each after the structures are well developed. The laser wavelength was 415nm, meaning the critical density for CH plastic was 20mg/cm^3, the foam densities were chosen to be 5, 10 and 20 times the critical density (n_{cr}). It is clear that the tunnelling effect is much more pronounced in the 5n_{cr} case than the others, however, this was discarded as it was incorrectly believed that a foam density this low would be difficult to manufacture. The 20n_{cr} case shows little to no tunnelling, so would not be suitable. This leaves 10n_{cr} as the clear choice for the experiment. Using the data dump shown in figure 4.18, for the 10n_{cr} simulation, synthetic X-ray phase contrast images can then be used to determine what detector and source setup is appropriate, how achievable the experiment is and what quality of results would be expected.

Since the simulations are fully 3D structures on an Eulerian grid, the layered isosurfaces method outlined in section 4.2.2 must be used to import the structure into Geant4. Once in GREENER, these structures were then backlit with the proposed X-ray source, the full details of this source are given in [123]. The source is a betatron source, from a laser wakefield

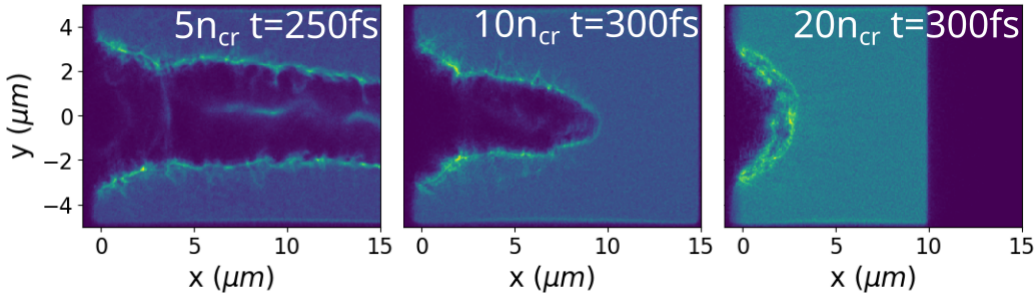


Figure 4.18: Simulations of proposed relativistic transparency experiment, for 3 possible densities, at timesteps once structures are well developed.

accelerator, with a critical photon energy of approximately 10keV, the spectrum used is shown in figure 4.19. The source size was taken as 1μm, with a divergence of 20mrad. Within the 20mrad cone, the total flux was 7e5 photons per eV at 5keV, this was used to calculate the total number of photons in the detector field of view for each configuration tested.

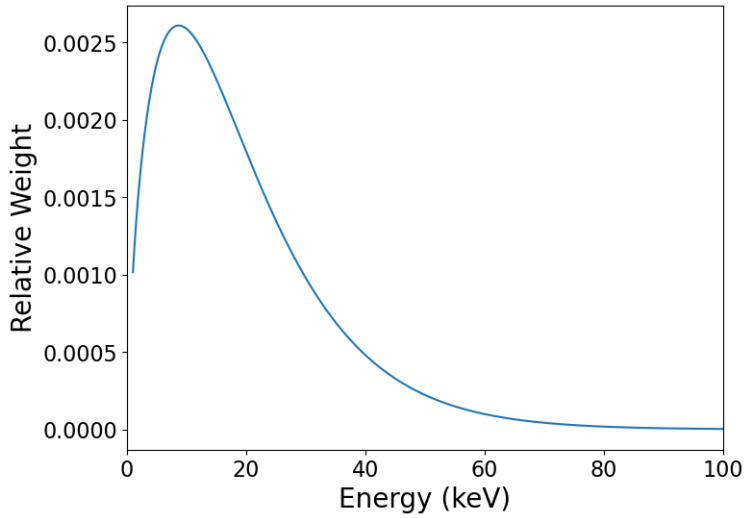


Figure 4.19: The weighted spectrum used to approximate the betatron source expected for the relativistic transparency experiment

The detector was placed 3 meters back from the target, and the source was initially placed 10cm before the target.

An example absorption only image, equivalent to contact imaging, of the 10n$_{cr}$ sample at t=300fs is shown in figure 4.20. Here the structures are not visible at all, despite artificially high photon flux, indicating that XPCI is necessary to image the structures in these foams, introducing refractive effects makes the structures much clearer.

The imaging configurations shown throughout this subsection all have low enough photon

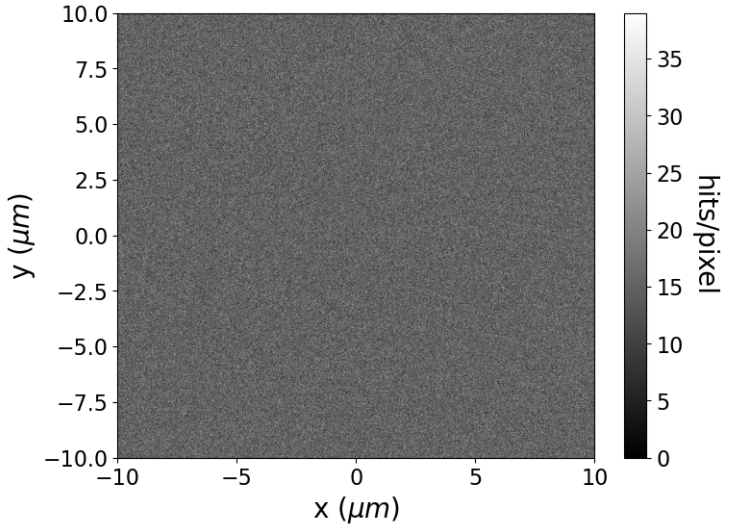


Figure 4.20: Absorption only synthetic X-ray image of an epoch simulation of the relativistic transparency experiment, no features can be seen without XPCI.

flux, that every photon could be simulated by GREENER, meaning the shot noise in the synthetic images is the same as would be expected in the actual experiment. The scales on the colorbars were chosen to maximise contrast on the structures being imaged.

As the source is a LWFA, in addition to the X-rays, there is also an electron beam. This electron beam would saturate and cover the image if it was allowed to hit the detector, so a magnet must be used to bend the electrons away from the imaging axis before they hit the detector. Initially the plan was to do this before the target, which required the source to be placed at least 10cm away from the target. A synthetic X-ray phase contrast image with this configuration is shown in the top left panel of figure 4.21, here it is clear that the photon count is too low to easily pick out the structures within the foam, so the source needed to be placed closer to the target to increase photon flux. Moving the source to 5cm or 2.5cm away from the target, produces the figures shown in the other two panels of 4.21, these show clear improvements in contrast, suggesting that this source needs to be placed closer to get clear images. This experimental setup can be achieved by placing the bending magnets after the target, which should be acceptable as the X-rays will get through the target before the electron beam can disrupt the target.

The reason images at 5cm and 2.5cm have a clearer signal vs noise, is entirely as a result of decreased shot noise due to higher photon flux. The signal from the phase contrast effects is actually better at 10cm, this is made clear with figure 4.22, where the source was placed

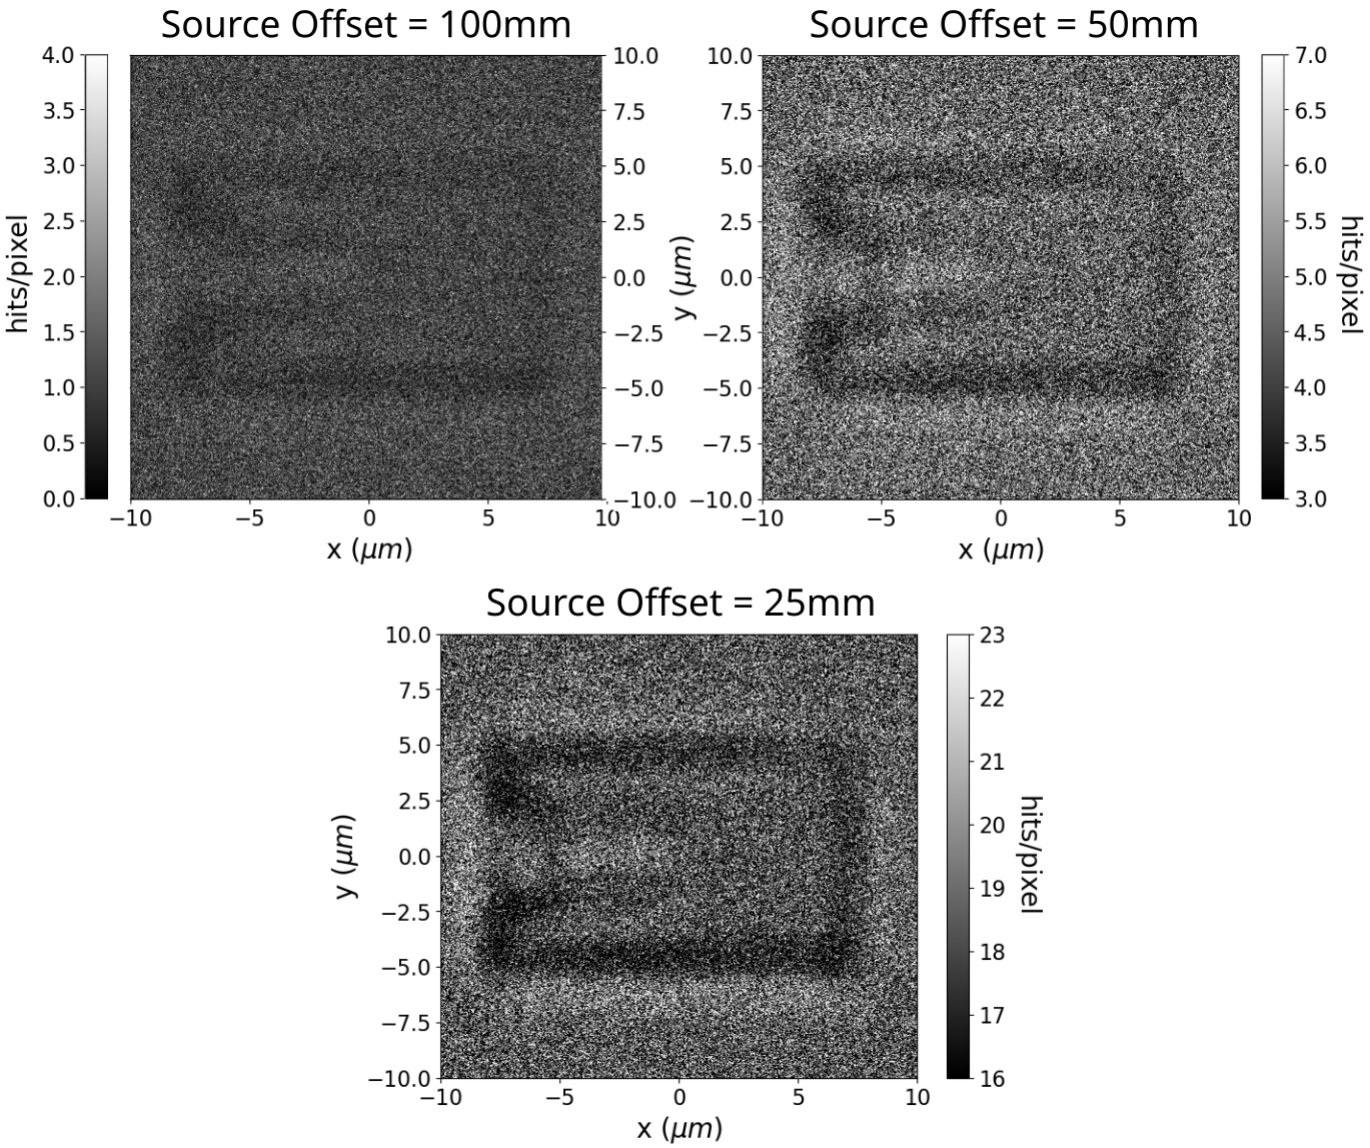


Figure 4.21: Synthetic X-ray phase contrast images of the relativistic transparency experiment, showing different possible source offset distances.

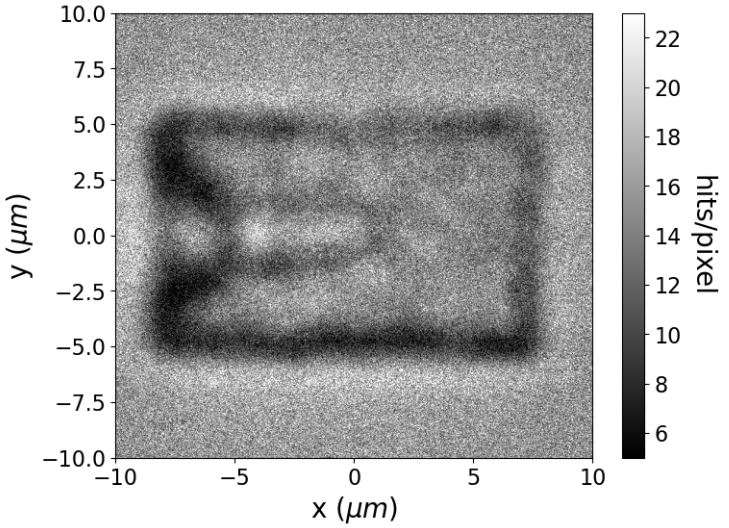


Figure 4.22: Synthetic X-ray phase contrast image of the relativistic transparency experiment, with a source offset of 1m, and artificially high source brightness.

a full 1 meter back from the target, but with increased flux so that it has the same number of photons on the detector as the 2.5cm. This image is the clearest of all the images, due to the strong phase contrast effects, highlighting the trade off between increased contrast vs increased photon flux.

4.6 Chapter Summary

This chapter has described the details of the GREENER code, which was written by the author for this thesis, within the GEANT4 framework. GREENER is used to generate synthetic X-ray phase contrast images of fusion relevant targets. Writing it involved developing multiple methods to import different complex geometry in such a way that surface interfaces were preserved. Synthetic images were shown for 1D, 2D and 3D problems, and comparison was made to experiment in the 1D and 2D cases. The code matched the experiments well, which supports its use as a synthetic diagnostic to design and interpret experiments.

Chapter 5

N+1 Tuning Design

5.1 Chapter Overview

Recent developments in indirect drive inertial confinement fusion have crossed the threshold of ignition, and developing ways to further increase yield has become increasingly important.

This chapter presents HYADES [106] simulations of implosions, in which radiation temperature profiles have been modified to mitigate the effects of the N+1 shock, resulting in improved implosion performance. Later in the chapter, the N+1 tuning method is combined with shock augmented ignition [32]. This combination provides a 10x increase in 1D yield over the baseline simulation, as well as improved ρR, hotspot pressure, ion temperature, shell velocity and IFAR. Suggesting a promising way forward to increase yield. The baseline implosion chosen was the NIF's N210808 shot, later implosions[14] at the NIF are more robust, so further work could explore whether the increase in implosion performance still holds in these regimes. This chapter includes simulations with alpha heating turned off that still show improved implosion metrics, suggesting the improvements in yield are not simply due to being close to the ignition cliff.

The N+1 shock forms in indirect drive implosions when reflected shocks merge in the ice layer. These reflected shocks emerge from the dynamics of the first three shocks in a typical high adiabat NIF implosion interacting with interfaces within the capsule shell. The formation and propagation of the N+1 shock introduces additional entropy into the ice layer, and therefore limits performance. This chapter starts by describing the details of this shock formation, the impact of the N+1 shock and the motivation for tuning it. The chapter then goes on to discuss the modifications that can be made to a radiation temperature profile to

mitigate the detrimental effects of the N+1 shock. As well as showing a path to improved yield, the modifications shown in this chapter aid in the understanding of the N+1 shock and its effects. Finally, this chapter will show that the increased yield resulting from these proposed modifications is robust against the small pulse shaping errors typical of the NIF, supporting the case for fielding this design as an experiment.

5.2 Origin of the N+1 Shock: Rarefactions and Reflected Pressure Waves

Starting with an IFE capsule, as the implosion experiment begins, the X-ray radiation flux on the surface of the capsule increases. This radiation drives the process of ablation, which increases the pressure at the ablation surface, a gentle increase would propagate inwards as a pressure wave, but in a typical NIF design the increase in pressure is rapid enough to form a shock which travels inwards towards the centre of the capsule [22].

When the shock crosses an interface into a lower density region (eg the ablater-ice interface or ice-gas interface), the shock continues towards the centre, while a release wave or rarefaction forms and propagates back outwards away from the centre [35]. As the rarefaction passes through material, the rarefied material is at a lower pressure than it was before the rarefaction takes place, this means that when the rarefaction reaches the ablation front, it lowers the pressure of the material which is right up against the ablation front. The ablation pressure remains the same, so this creates a stronger pressure gradient which needs to be rectified. The situation resolves because the pressure difference drives a pressure wave back in towards the centre of the capsule [38].

The pressure wave which forms is strengthened by the convergent effects of spherical geometry, and by the inherent steepening of pressure waves propagating close to the speed of sound, as a result it may steepen up into a shock. The reflected shock may also merge with other shocks or reflected shocks to form a stronger shock.

5.3 Formation of the N+1 Shock

In current (high adiabat) NIF implosions (such as N210808), an initial three shocks are launched [124], and then there are a number of reflections. Figure 5.1 shows a Lagrangian

frame plot of the density gradients in a 1D HYADES [106] simulation of the N210808 implosion, with stronger density gradients shown as stronger colours. The Lagrangian frame plot is generated by taking each cell at each timestep, and instead of plotting at the position it occupies at the timestep in question, instead it is plotted at the position it occupied at t=0. This means the Lagrangian cells are shown as if they were stationary, and shocks can be tracked more easily as they travel through the different layers in the capsule. Inward travelling shocks and outward travelling rarefactions appear on the plot as green lines, while the ablation front appears in purple as its density gradient goes in the opposite direction. The shocks are labelled, and for a given shock "Shock X", it's reflection from an interface is labelled "Shock Xr".

As the implosion progresses, the radiation drive is increased to launch the first shock (Shock 1). Shock 1 creates an outward propagating rarefaction as it crosses the ice-ablator interface, this rarefaction then reflects off the ablation front to form a new shock (Shock 1r) which moves back in towards the centre. The radiation drive is increased again to launch shocks 2 and 3, which work their way in towards the centre along with shock 1r. These shocks each create their own outward propagating rarefactions as they cross the ice-ablator interface, the rarefactions then reflect off the ablation front to form yet more new reflected shocks (Shock 2r, 3r and 1rr).

Meanwhile, the shocks 1, 2, 3 and 1r all continue through the ice layer towards the ice-gas interface, and in the implosion shown in figure 5.1, they all merge to form Shock M just before the ice-gas interface. In a well designed implosion, the merger should happen just before the ice-gas interface, but with enough of a buffer that small pulse shaping errors won't push the merger into the gas layer. This is because strong shocks increase the adiabat more than weak shocks, so the target averaged adiabat of the implosion is minimised by minimising the volume of ice which experiences the strong merged shock. However, if the merger is allowed to happen after the shocks cross into the gas layer, then rarefactions will form as the shocks cross the interface, which lowers the density of the material which subsequent shocks go through, which effectively strengthens them, and increases the adiabat as a result.

As long as the merger happens before the ice-gas interface, when shock M crosses the ice-gas interface it creates an outward propagating rarefaction (Rarefaction M), which works its way out towards the ablation front. This rarefaction then reflects off the ablation front to form Shock Mr, which begins moving inwards towards the centre.

Now (between roughly 5 and 6 ns into the implosion) there are 4 reflected shocks (Shocks 2r, 3r, 1rr, Mr), all moving towards the centre. They merge part way through the ice layer to form the "N+1 Shock", this is problematic as the N+1 shock is stronger than the 4 other shocks, and significantly increases the adiabat in the material it passes through. The increase in adiabat can be observed as a horizontal green feature that occurs at each of the N+1 merger points. Material above the merger has experienced multiple weak shocks, while material below the merger has experienced a single strong shock, so the material below the merger is at a higher temperature and adiabat, which results in a lower density at the same pressure. This density gradient appears as a green horizontal line in the lagrangian frame shock plot.

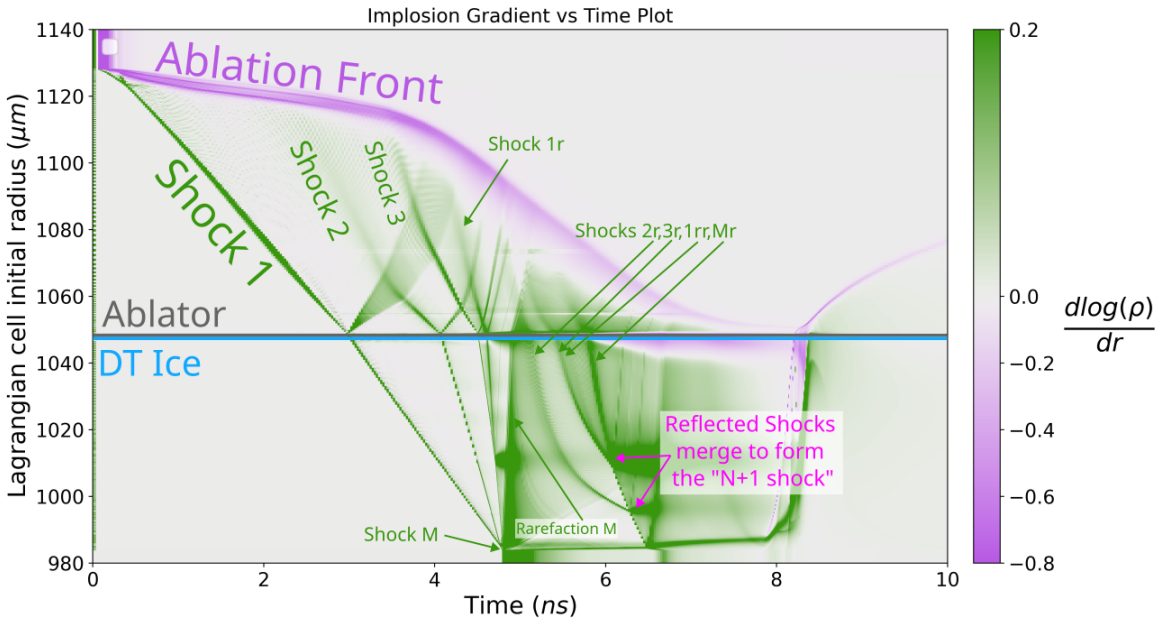


Figure 5.1: A Lagrangian shock plot with the main features, shocks and rarefactions labeled. This shows the first three shocks as they propagate through the ablator and ice into the gas, as well as the rarefactions they create, then the reflected shocks which form from those rarefactions, until finally the reflected shocks merge to form the N+1 shock.

5.4 Justification for Tuning the N+1 Shock

The "N+1 Shock" is much stronger than the pre-merger shocks(shocks 2r, 3r, 1rr and Mr), so as a result the N+1 shock increases the adiabat much more than the unmerged shocks would have on their own. This reduces the fuel compressibility, which in turn reduces the maximum ρR, therefore reducing the yield obtained.

Figure 5.2 shows a lab frame plot of the N+1 shock as it travels through the ice layer; the jump in adiabat is clearly visible. A Lagrangian frame plot is shown in figure 5.3 which shows the effect on adiabat across the whole implosion. This also provides the explanation for the horizontal green lines that appear at shock mergers (eg when shock 3 and shock 1r merge around 4.75ns or when shocks 3r, 1rr and Mr merge around 6.25ns in figure 5.1). The material above the horozontal line has experienced a few weak shocks, while the material below the line has experienced one strong shock, so is at a higher adiabat. The higher adiabat means the fuel is less compressable, so is at a lower density for the same pressure, creating a density gradient that appears horizontal due to plot being in the lagrangian frame.

In figure 5.3, the initial hotspot edge is shown, which is the beginning of the gas cells, if the merger can be delayed until the shocks are close to or inside the gas layer, then the corresponding adiabat increase will affect less of the fuel, and therefore it's detrimental effect on implosion performance will be less reduced.

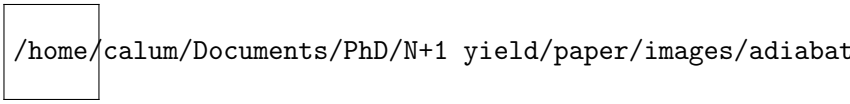


Figure 5.2: A plot of density, pressure and adiabat against radius in the lab frame, for the N210808 shot at a time chosen to show the N+1 shock travelling through the ice layer (t=6.4ns), where the increase in local adiabat due to the N+1 shock is clear.

Tuning the N+1 shock consists of altering the laser pulse shape, and therefore radiation temperature profile, to delay the N+1 shock merger by accelerating the earlier reflected shocks (eg shock 2r), to prevent the later reflected shocks (eg shock Mr) from catching up, thereby delaying the merger until after the shocks cross into the DT gas. The N+1 shock can also be mitigated by slowing the later reflected shocks to delay the mergers; in either case the ice layer is kept on a lower adiabat, so the fuel remains compressible. The details of how different radiation temperature profiles can affect the strength of reflected shocks are discussed in appendix B, using the designs from sections 5.5 and 5.6 as illustrative examples.

The next few sections will show that this technique works in simulations.

5.5 Baseline Simulation

To test the idea of N+1 tuning, simulations were run using the 1D Lagrangian code: HYADES [106]. The simulations were based on the NIF's N210808 shot [29, 125], and used the capsule

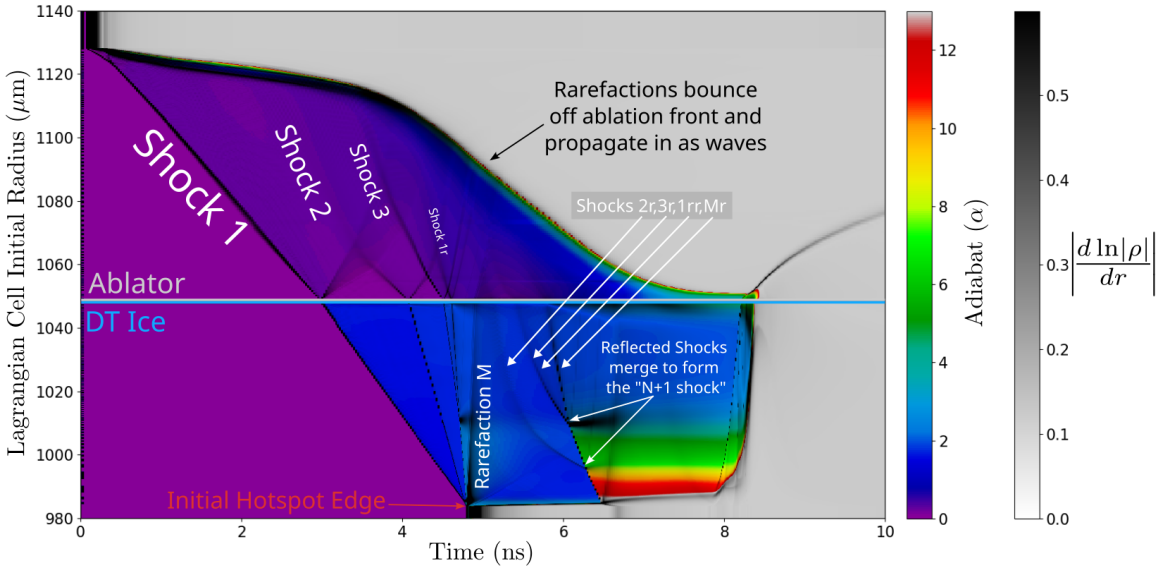


Figure 5.3: A Lagrangian shock plot, with shocks shown in black, combined with a Lagrangian frame plot of the cell adiabat, shown as spectral colours. This shows the significant increase in adiabat precipitated by the N+1 shock merger

layers and materials from that experiment.

Dante data from N210808 was used for the radiation temperature profile, with shocks 1, 2, and 3 retimed so that they merged close to the ice-gas interface. This is necessary because there are a number of uncertainties, in particular: the exact equation of state is not known, and the radiation temperature experienced by the capsule is not exactly the same as the Dante measurements. This means that the shock timing in HYADES does not resemble the designs intended shock timing until the radiation temperature profile has been modified. Figure 5.5 shows the Dante data along side the radiation temperature profile used in the simulations. The dante data is smoothed, the late time bump which comes from the capsule igniting is removed, and the early shocks are re-timed to get the radiation temperature profile. Retiming simply aimed to time the shocks so they merged as close to the ice-gas interface as possible, making this an imperfect match to the N210808 experiment. However, without any retiming, shocks 1 and 2 merge high in the ice (993μm instead of 985μm after retuning), and shock 3 doesn't merge at all, it crosses the ice-gas interface on its own 0.3ns after shocks 1 and 2 have merged and crossed the ice-gas interface. Table 5.1 shows the tuned and untuned implosion metrics for reference.

For radiation transport a multigroup radiation model was used, the spectrum used was

Table 5.1: Table of implosion metrics, comparing simulations of N210808 before retuning the first 3 shocks, and after. Both simulations have the same flux limiters. ρR, T_{ion} and $P_{hotspot}$ are burn averaged.

Implosion Metric	Unaltered N210808	Retimed N210808
Yield (MJ)	0.53	1.15
ρ R (g/cm^2)	0.595	0.554
Adiabat	4.55	4.17
T_{ion} (keV)	6.65	8.3
$P_{hotspot}$ (Gbar)	212	253
max v_{shell} (km/s)	422	436
IFAR (75%)	15.08	25.3
Fall Line Radius (μm)	16.9	14.8

a thermal spectrum with an additional M-band fraction based on experimental data, the M-band fraction from the experiment and after removing the late time bump is are also shown in figure 5.5. 70 radiation groups were used in the radiation transport model, this was a compromose between accuracy and computational speed. A convergence test is shown in figure 5.4, which shows that below 40 or so radiation groups the hotspot ion temperature at bang time, and maximum areal density both become highly sensitive to the number of radiation groups. Yield is shown for reference, but the simulation is close to the ignition cliff, it is highly sensitive to small changes in the ion temperature and areal density, so it is less useful for testing convergence. Above 40 radiation groups the metrics stabilise into a downward trend, however, with increasing radiation groups comes increasing computational cost. Above 70 radiation groups the computational time increases significantly, and HYADES does not allow more than 150 radiation groups by default, so only points for 84 and 140 are plotted. For the purpose of this thesis the number of radiation groups will remain fixed between different modified radiation temperature profiles, and the actual yield/ion temperature will be tuned by adjusting the flux limiter, so this trend is not a concern.

In the real N210808 experiment, the actual shock timing is unlikely to be as good as the designed timing, so when the flux limiter was tuned to match the yield (as described in section 3.2.1), this is likely an underestimate of how much the flux limiter should be lowered to match the yield. In addition, after tuning the shocks and flux limiter, the yield was still slightly lower than the experiment (1.15MJ rather than 1.37MJ), and the bang time was earlier (8.3ns rather than 8.8ns). So this baseline simulation is not a perfect match to the experiment, which is why the improvements in yield found by modifying the baseline

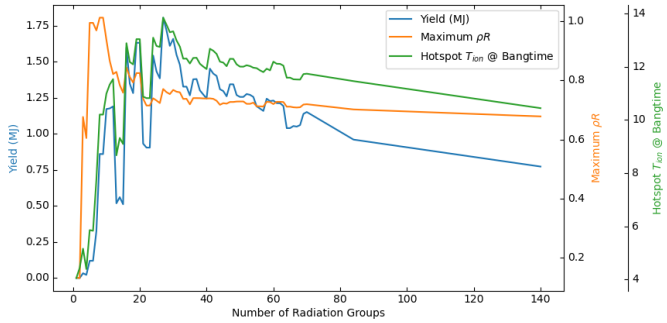


Figure 5.4: A convergence test for the number of radiation groups to be used in hyades simulations, with the hotspot ion temperature at bang time, maximum areal density, and yield plotted for each number of radiation groups up to 70, after 70 radiation groups, only 84 and 140 groups are plotted.

simulation will be measured against the values found in the baseline simulation rather than the experimental values.

Throughout this chapter, the radiation temperature profile is modified in various ways, however, the M-band fraction remains unchanged. There is considerable uncertainty on the M-band fraction as the Dante diagnostic only allocates 2-3 channels to measuring it[126]. While calculating how the M-band fraction would change with a given change to the radiation temperature profile would be a complex and interesting area for future work, it may not be possible without experiments, hence the M-band data was left as it was in the baseline. There is also a noticeable jump in the M-band fraction around 8.5ns, this was the result of a bug in the bump-removal code, however as it is a small jump which is late enough in time that it should not have a significant effect on the results.

5.6 Peak Only Design

There are a number of ways the radiation temperature profile can be altered to tune the N+1 shock, the simplest of which is to introduce an appropriately timed peak in the radiation temperature profile. This modification was refined using a constrained Bayesian optimisation algorithm to find the highest yield design, the SCIKIT-OPTIMISATION module was used, which implements a gaussian process based Bayesian optimisation algorithm. Bayesian optimisation was chosen as it is particularly well suited to optimisation problems where the goal is to find the maximum or minimum output from a computationally expensive function with a few

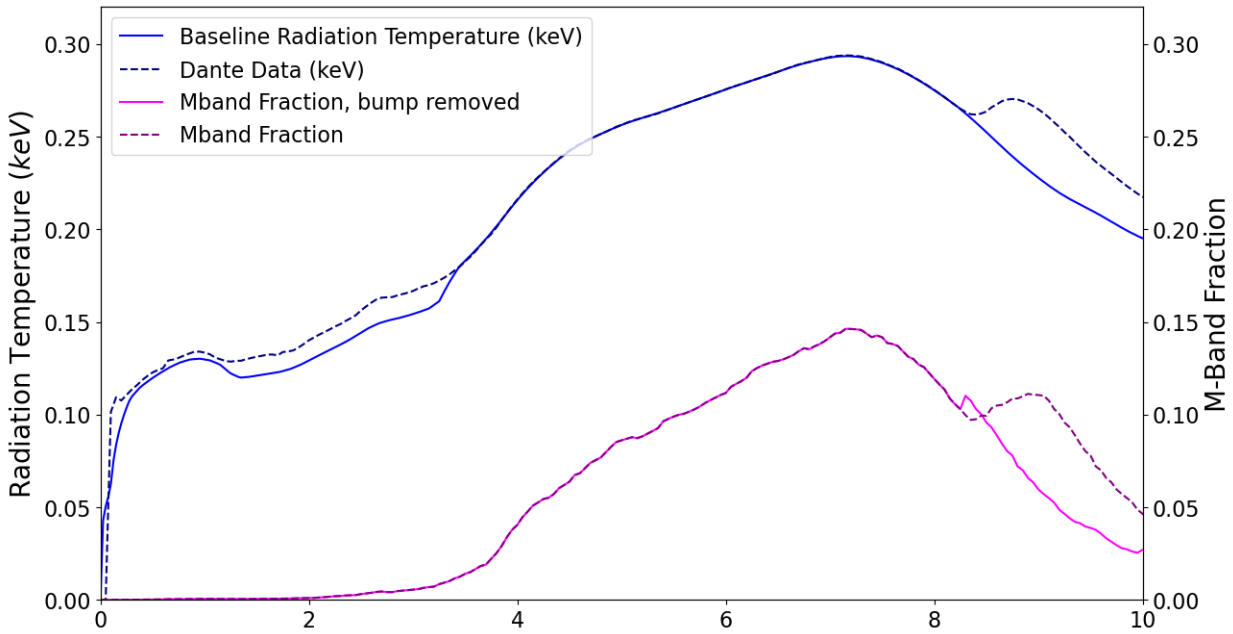


Figure 5.5: The experimental data from N210808 for the radiation temperature and M-band fraction, shown along side the radiation temperature profile and M-band fraction used as the baseline for the simulations.

input parameters. SCIKIT-OPTIMISATION provides a good tutorial [127], and Frazier [128] provides a more in depth description.

The optimisation was constrained using dante data from N210808 and a shock augmented ignition(SAI) shot (N221025, see section 2.2.8 for details on SAI). The dante measurements suggest that the maximum rise that can be achieved in the N210808 hohlraum is 60V/ns, while the fastest drop seen in the SAI dante data is 45eV/ns, so those constraints were applied to the optimisation. More complex hohlraum models [129], or improved hohlraum design could potentially loosen some of these constraints, but for this analysis, these experimental constraints will be used.

The total energy in the radiation temperature profile was also conserved so that any modification requires no additional energy from the laser driving the hohlraum, this was done using a scaling law to convert radiation temperature to laser energy required [130].

Figure 5.6 shows the modified radiation temperature profile in black, with the original one in blue, as was mentioned earlier: both use the same experimentally determined M-band fractions over time. Comparing this with figure 5.1, it is clear that the first few reflected shocks (2r,3r,1rr) have been accelerated, such that the later shock (Mr) does not catch up with them until later than in the N210808 design. This effectively delays the mergers which

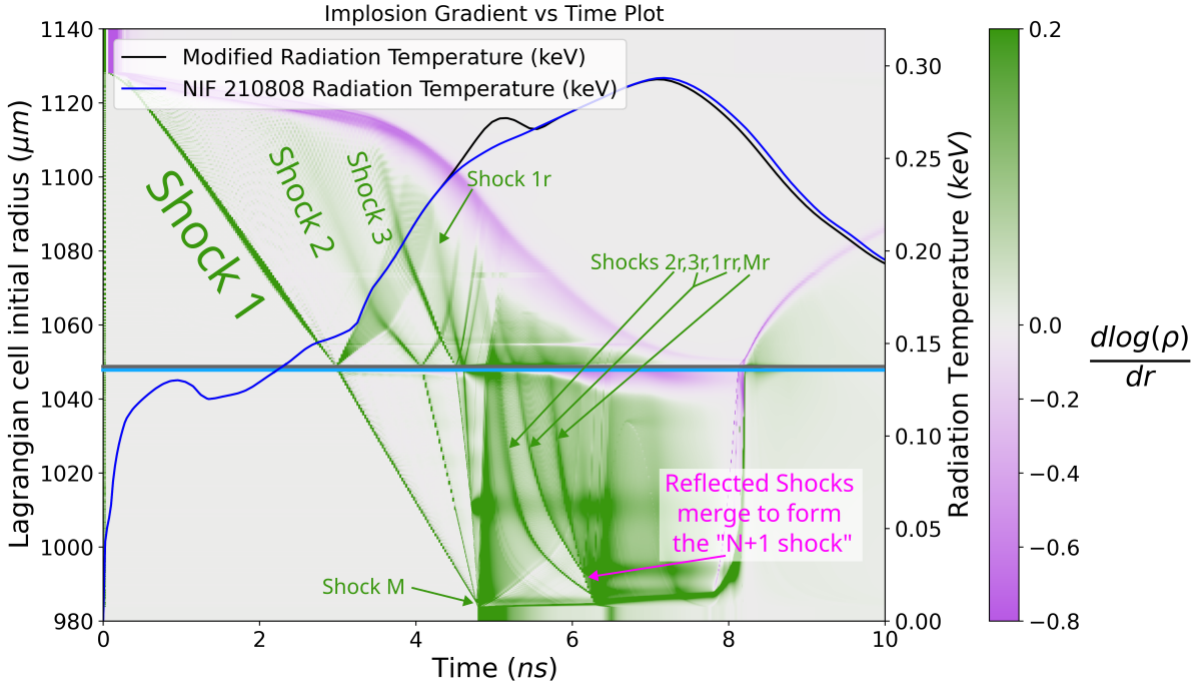


Figure 5.6: A Lagrangian shock plot for the result found by constrained bayesian optimisation, showing N+1 Tuning has been effective, as the N+1 shock mergers have been delayed until closer to the ice-gas iterface than in the baseline case.

form the N+1 shock until much closer to the gas-ice interface. In the baseline case shocks 3r, 1rr and Mr merge at 1011μm in the lagrangian frame, then shock 2r merges with them at 995μm, whereas in this design the mergers occur at 994μm and 988μm respectively. As a result of these mergers occurring lower in the ice, the amount of material which experiences a strong N+1 shock is reduced. In turn this limits the effect of the adiabat increase caused by the strong shock.

The consequences of the N+1 tuning and the corresponding reduction in adiabat are seen across a range of implosion metrics. These are shown in table 5.2, where the comparison is made with a simulation of N210808 so that identical measures can be used. The lower adiabat of this design allows for increased compression, so ρR is higher. This contributes, along with the higher hotspot pressure and ion temperature, to a higher yield. This implosion also has higher implosion velocity, which might contribute to higher yield as there is more kinetic energy available to convert into internal energy. However, higher implosion velocity also suggests increased susceptibility to Rayleigh-Taylor instabilities, as does the high IFAR of this design. Rayleigh-Taylor instabilitites are a 2D effect, which reduce the yield of an implosion through mechanisms like mixing, and are not captured by these 1D simulations.

The fall line radius does suggest that the peak only design is better in the worst case scenario where the ablator freely passes into the ice layer without decelerating, but this metric does not account for the fact that the baseline and SAI only designs are likely to have reduced mix. In particular for the SAI design, in addition to reduced mixing due to lower IFAR, the deceleration is from a reduced peak velocity. The reason the fall line radius is lower in the SAI simulation is that the time from peak velocity to bang time is longer (0.5ns vs 0.46ns in the baseline case and 0.4ns in the peak only case). So the usefulness of this metric is limited for these designs, and IFAR and implosion velocity will be the focus for discussing mix. Fortunately, IFAR and implosion velocity can be brought down by introducing a small dip after the peak, this will be discussed later in this chapter.

Table 5.2: Table of implosion metrics, directly comparing simulations with and without the N+1 Tuning modification. ρR, T_{ion} and $P_{hotspot}$ are burn averaged.

Implosion Metric	Baseline (N210808)	Peak Only	SAI Only
Yield (MJ)	1.15	5.21	6.37
ρ R (g/cm^2)	0.554	0.619	0.699
Adiabat	4.17	2.81	3.03
T_{ion} (keV)	8.3	15.5	15.4
$P_{hotspot}$ (Gbar)	253	627	548
max v_{shell} (km/s)	436	443	333
IFAR (75%)	25.3	34.5	16.7
Fall Line Radius (μm)	14.8	19.8	-9.8

5.7 Timing Scan

To demonstrate that the benefits of this modification primarily result from N+1 Tuning, it is easiest to relax the hohlraum requirements, and allow the radiation temperature to rise or drop arbitrarily rapidly, which allows the peak to be freely shifted to occur at different points along the radiation temperature profile. This is shown in figure 5.7, where yield is plotted as a function of the midpoint of the peak, also shown is the baseline yield and a rough idea of the area where hohlraums cannot achieve the requested rise and fall rate in radiation temperature. The fact that there are a range of timings which all give higher yield suggests that this scheme is robust to timing issues and uncertainties of $\pm$ 100ps.

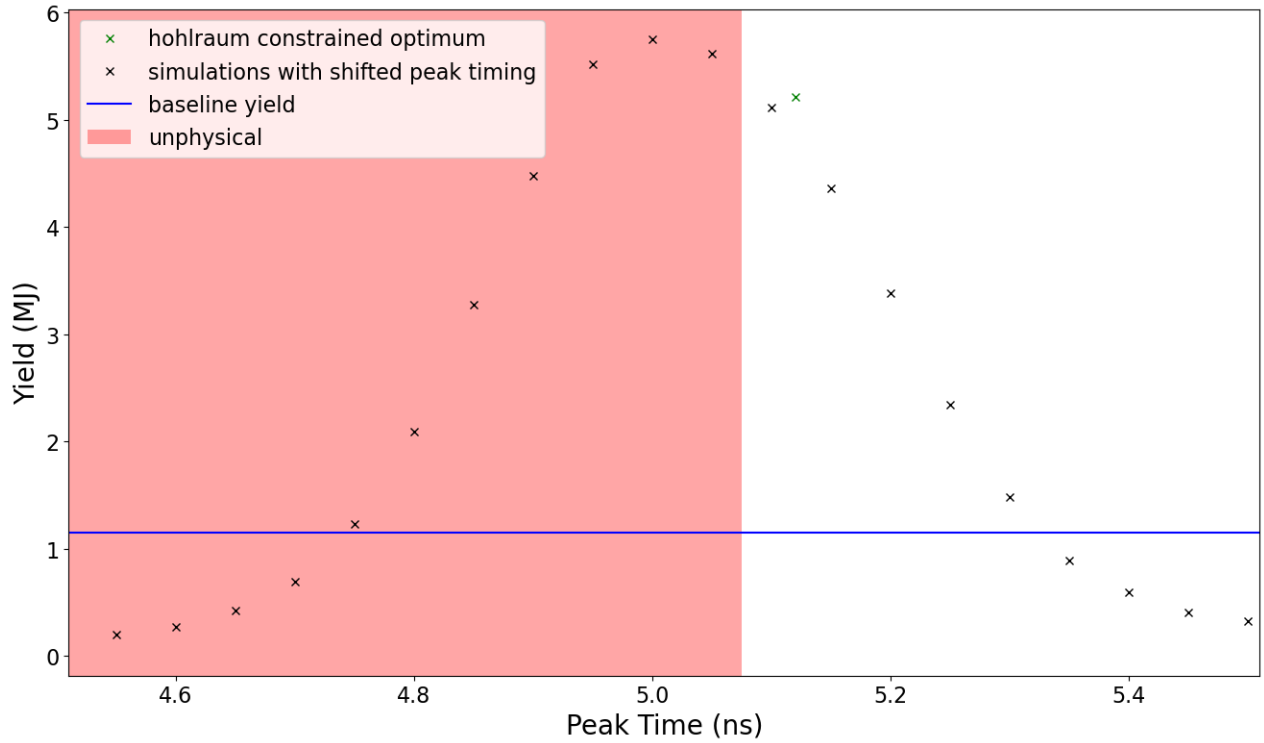


Figure 5.7: A plot of yield vs time of peak maximum

The timing scan also makes it possible to test the idea that the increased yield comes tuning the N+1 shock, rather than from some other factor (eg strengthening or launching a shock, generally introducing more energy into compression, or something else).

Three example profiles, chosen from the timing scan as examples where the peak is added too early (figure 5.8); at the right time (figure 5.9) or too late (figure 5.10), are shown. Looking at the structure of the shocks across the timing scan, the simulations suggest that adding a peak very early in time will disrupt the shock timing for the initial shocks, which results in reduced yield compared to the baseline; if the peak is slightly early, then it will partially stop N+1 shock mergers, increasing yield as the timing shifts towards the optimum timing. At the optimum timing, the peak is driving the early reflected shocks (2r,3r,1rr) and increases their shock strengths, which means that the later reflected shock(Mr) can only catch up with them and merge very close to the ice-gas interface. After the optimum timing, the peak starts to drive the later reflected shock (Mr), increasing its strength and as a result, the N+1 shock merger begins to move back into the ice layer. Eventually, the merger starts to happen earlier than it would in N210808, resulting in lower yields than the baseline again.

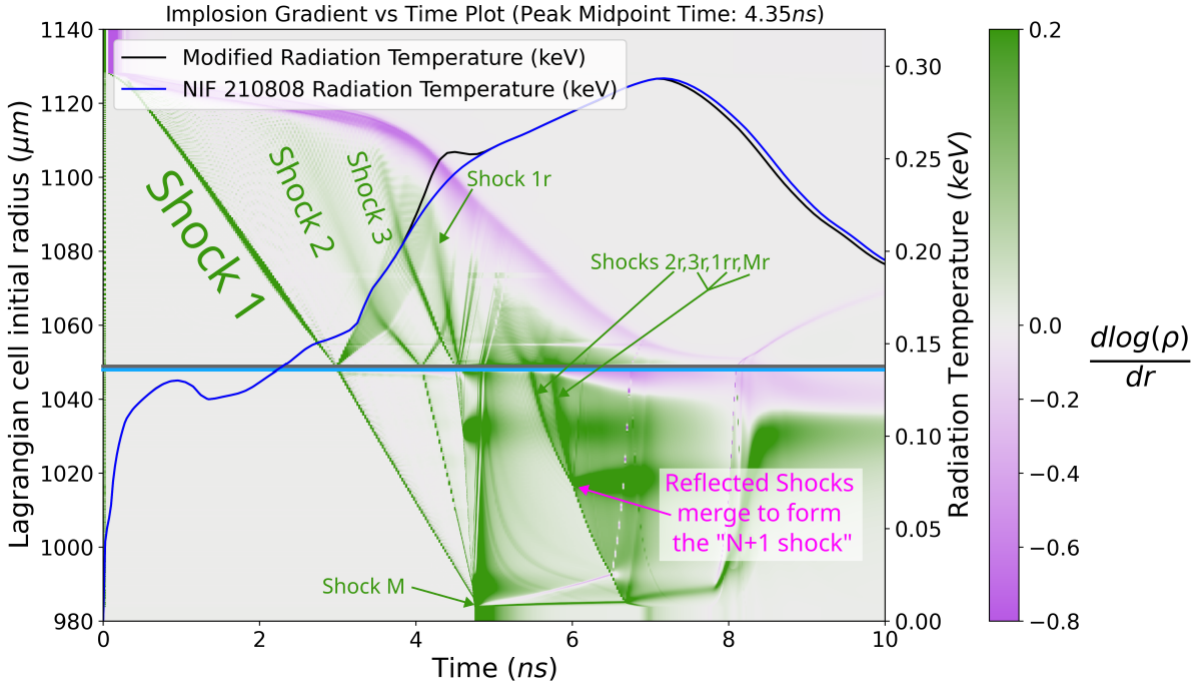


Figure 5.8: A Lagrangian shock plot for a peak midpoint time (4.35ns) which is too early to tune the N+1 shock. The peak strengthens the first shock reflection too much (resulting in an early merger at 1017μm in the lagrangian frame).

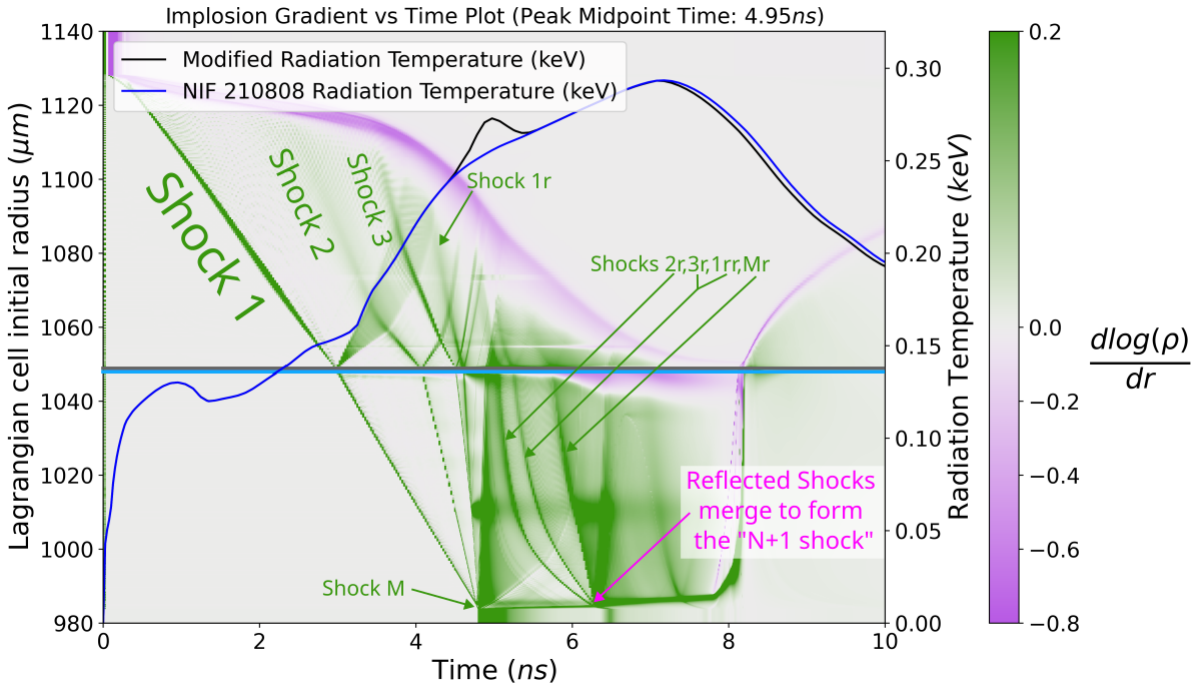


Figure 5.9: A Lagrangian shock plot for a peak midpoint time (4.95ns) which is relatively well timed to tune the N+1 shock. All 3 reflected shocks hit the ice-gas interface almost sequentially.

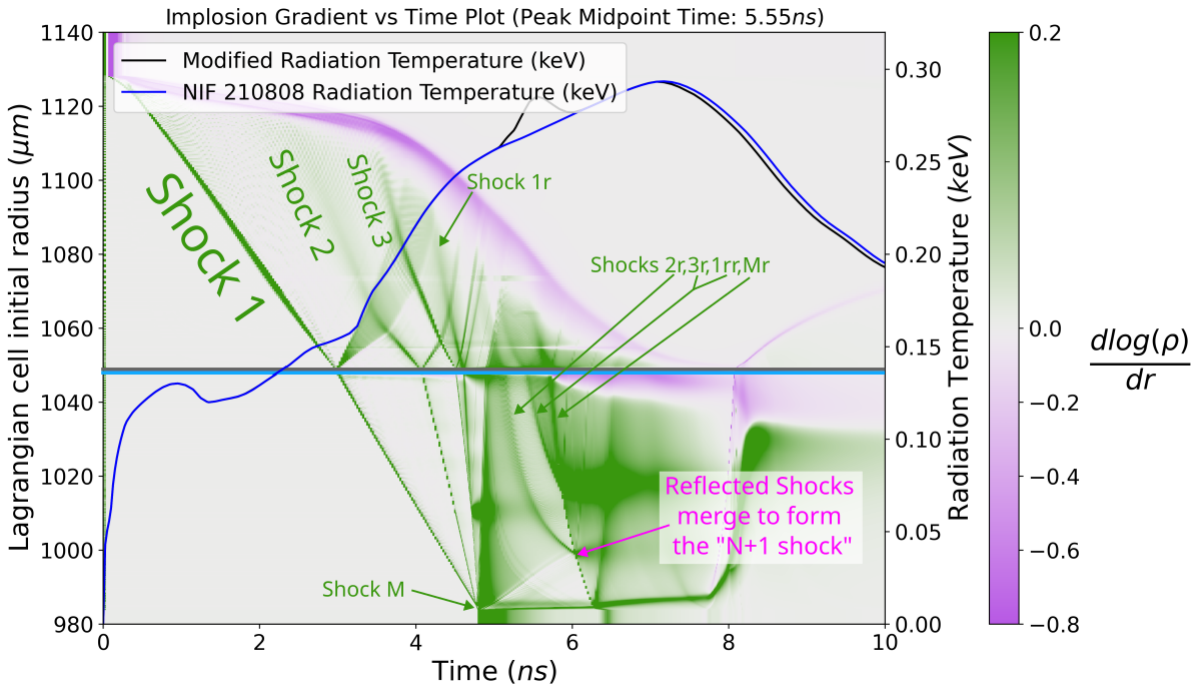


Figure 5.10: A Lagrangian shock plot for a peak midpoint time (5.55ns) which is too late to tune the N+1 shock. The early reflected shocks are able to merge with the later reflected shock, forming a strong shock in the ice layer (with shock mergers at 1017μm and 999μm in the lagrangian frame).

5.8 Peak and Dip Design

As well as introducing a peak to tune the N+1 shock by increasing the velocity of the early reflected shocks (2r,3r,1rr on figure 5.1), a small dip can be introduced to reduce the velocity of the later reflected shock (Mr), the combination of a peak and dip also moves the N+1 shock mergers lower into the ice relative to the baseline design (989μm for the 3r, 1rr and Mr merger, and 986μm for 2r to merge with the rest). The peak and dip design improves the performance more than a peak alone, and gives a lower IFAR design with lower implosion velocity. A design with a symetric peak and dip is shown in figure 5.11, which was found by a parameter scan.

The parameters chosen to characterise the peak-and-dip combination are shown in figure 5.12, and were:

1. the gradient of the downward slope from the peak to the dip

2. the time when the midpoint between the peak and dip occurs

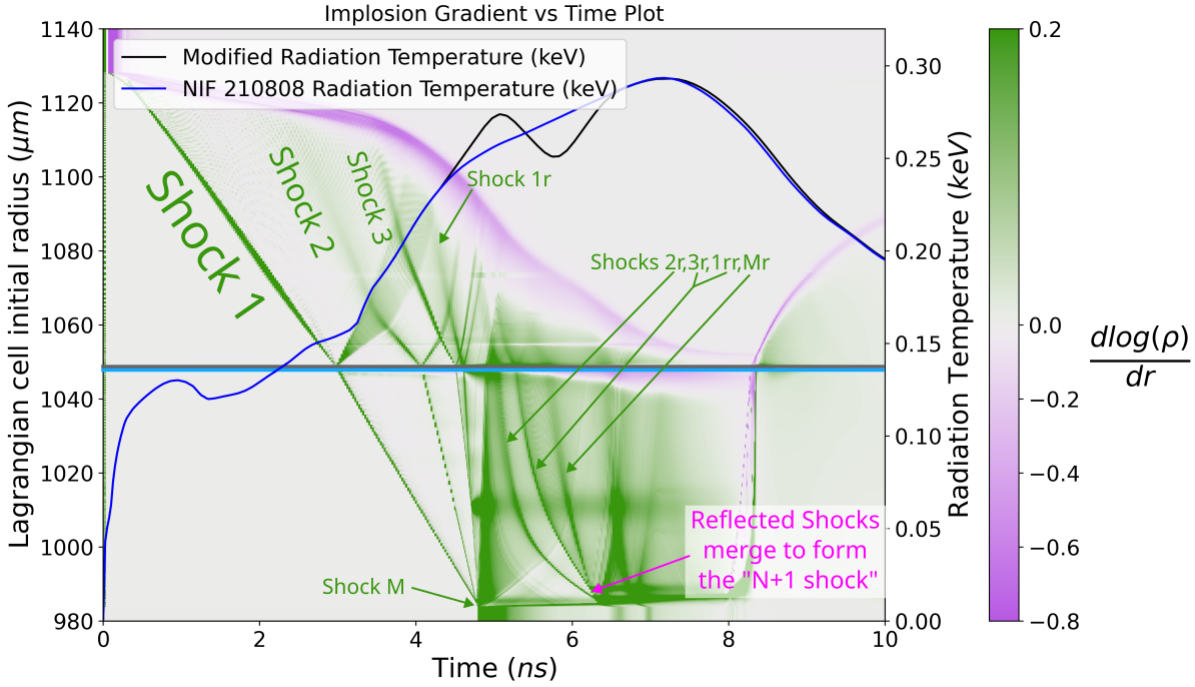


Figure 5.11: A Lagrangian shock plot for the result found by constrained optimisation of a peak and dip intorduced to the radiation temperature profile. Showing a slight improvement to the timing of the reflected shocks over a peak only design.

3. the width between the top of the peak and the bottom of the dip.

After constructing the pulse shapes as illustrated in figure 5.12, they were smoothed and passed to HYADES. The parameters were chosen primarily because they make it easy to construct the peak-and-dip structure within the hohlraum constraints whilst limiting the number of free parameters.

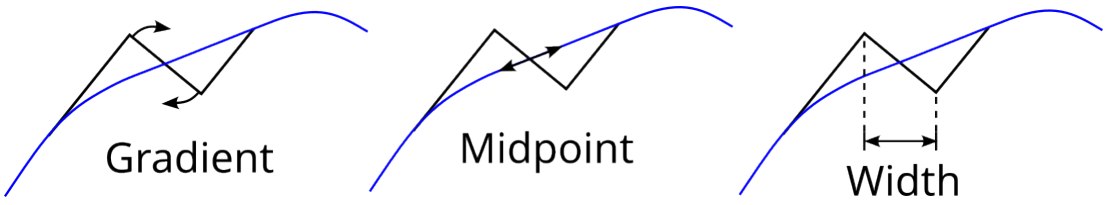


Figure 5.12: Diagram showing the 3 variables used to parameterise the peak-dip pulse shape. The baseline is shown in blue, with the modification shown in black, arrows show how the modification is changed with each variable. The pulse shape was smoothed before being passed to HYADES.

For two of the peak-to-dip widths with the highest yields, parameter scans of the other two parameters are shown in figures 5.13 and 5.14. For each parameter scan, the range of values available without including hohlraum constraints are shown on the left, while the right

shows only the set of values allowed by the hohlraum constraints. For both peak-to-dip times, it is clear that there is a large area of high yield, centred on an optimum set of parameters. The highest yield design, at just over 10MJ is in the unconstrained 1ns peak-to-dip time plot, but with hohlraum constraints so many of the high yield designs are cut off that the best result ends up being in the 0.8ns peak-to-dip time plot. If only simulations using experimental hohlraum constraints had been run, it would have appeared that simply moving to earlier times and steeper gradients is always better, but the full parameter scan reveals that there is a defined region of high yield to focus on.

It is of note that as hohlraum constraints are introduced, a large portion of the high yield shots are cut off, since this suggests that in addition to the significant design space within existing hohlraum constraints, N+1 tuning also opens up even more design space for high yield designs if hohlraums are redesigned to allow for sharper dips or rises. Sharper dips could be achieved by allowing for larger laser entrance holes (LEH) that would allow hohlraums to lose energy faster, however, this is a trade off as smaller LEH's allow for faster rises in radiation temperature. The inner surface area and material of hohlraums also play a role in the rise and drop rates of radiation temperatures, and further research using robust hohlraum models could prove productive.

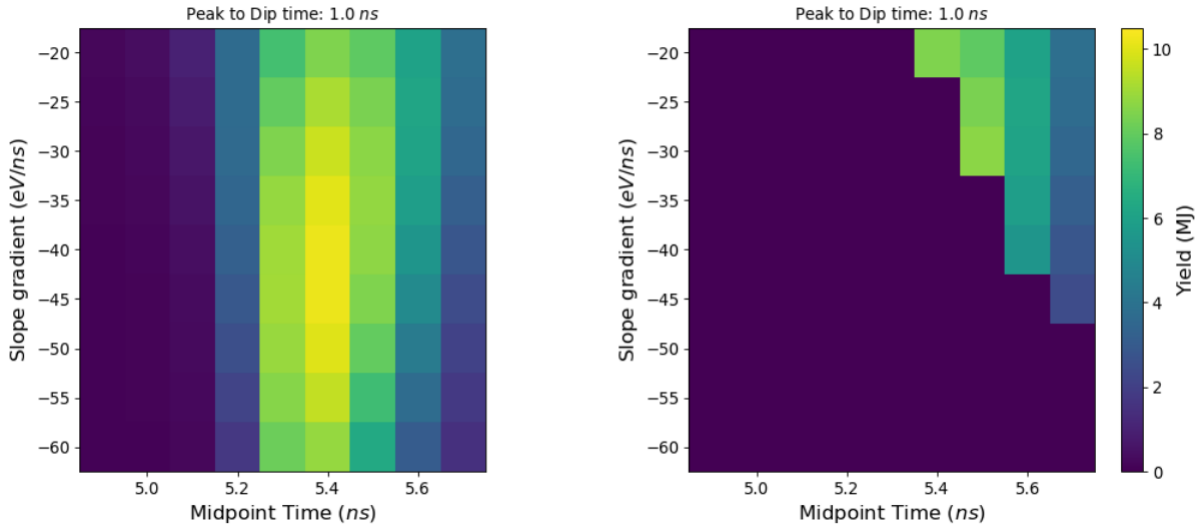


Figure 5.13: Heatmap showing the yield produced from HYADES simulations with different pulse shapes as a funciton of the slope gradient and midpoint, for a peak-to-dip time of 1ns. The left heatmap is without hohlraum constraints, and the right with hohlraum constraints.

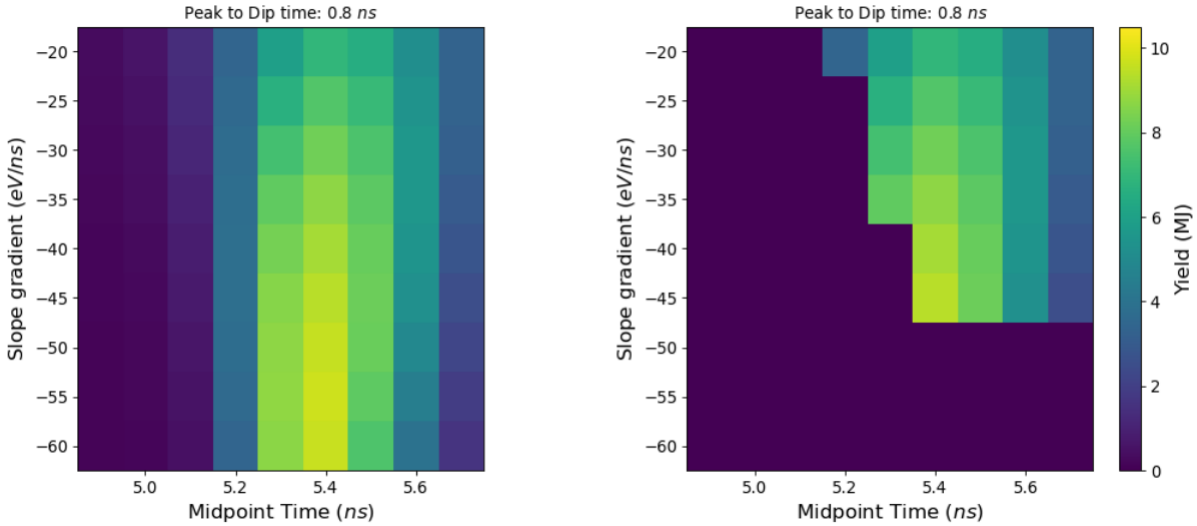


Figure 5.14: Heatmap showing the yield produced from HYADES simulations with different pulse shapes as a funciton of the slope gradient and midpoint, for a peak-to-dip time of 0.8ns. The left heatmap is without hohlraum constraints, and the right with hohlraum constraints.

5.9 N+1 Tuning and Shock Augmented Ignition

As discussed in section 5.6, adding a small dip after the peak will improve IFAR and implosion velocity; however, if the peak and dip are optimised as separate features with different parameters, the results move closer and closer to a combined N+1 tuning and shock augmented ignition like design. A pulse shape was constructed to combine N+1 tuning and SAI, and figure 5.15 shows the parameters used to build the pulse shape. The number of parameters involved is too large for a standard parameter scan, so Bayesian optimisation was used with some initial guesses provided by manually combining the peak only and SAI pulse shapes.

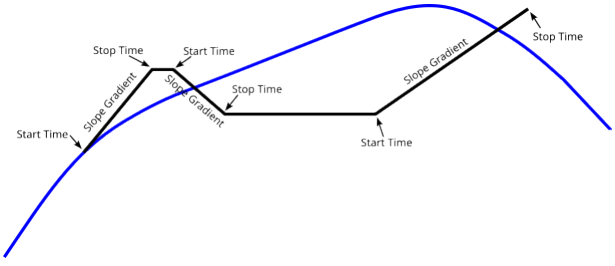


Figure 5.15: Diagram showing the many variables used to parameterise the peak-sai pulse shape. The baseline is shown in blue, with the modification shown in black, each variable used is labeled. The pulse shape was smoothed before being passed to HYADES.

The simulations suggest these schemes compliment each other, figure 5.16 shows the

Lagrangian shockplot for a pulseshape with a combined N+1 Tuning and SAI design. Here we see that the dip has helped tune the N+1 shock more than the peak alone would have, with shocks 2r, 3r, 1rr and Mr all essentially merging at the ice-gas boundary. In addition, by keeping the radiation temperature low until later in time, the rise can used to launch a shock augmented ignition style shock at the right time to get the benefits of that scheme. The result is a combined N+1 Tuning and SAI scheme which produces higher yield than either scheme alone.

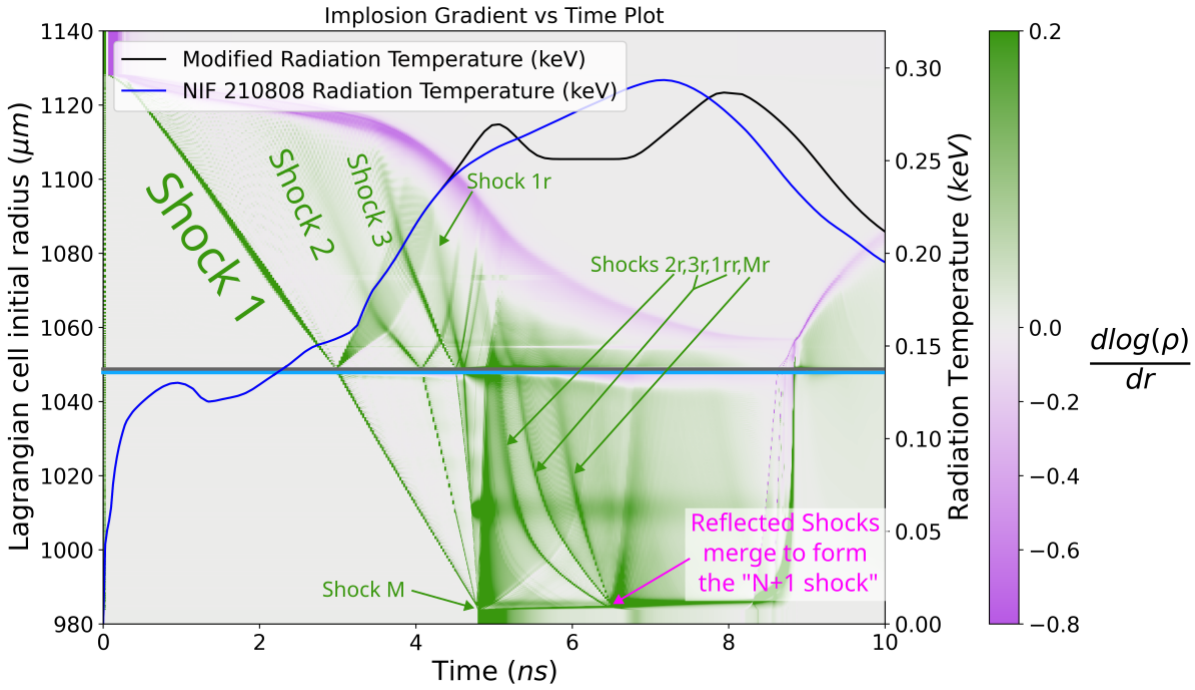


Figure 5.16: A Lagrangian shock plot showing the results of combining N+1 Tuning and shock augmented ignition.

The data in the updated table 5.3 shows a comparison of the baseline N210808 shot, N+1 tuning only, SAI only and a combined N+1 Tuning & SAI simulation. Here we see that combining the two schemes performs better than either scheme alone across all metrics except IFAR and Fall Line Radius, and even for IFAR this design is better than the baseline case.

5.10 Alpha Heating-Off Simulations

Being close to the ignition threshold means that yield is an imperfect metric, since a small change in performance gives a large change in yield. This is a limitation of the work, which

Table 5.3: Table of implosion metrics, directly comparing simulations with and without the N+1 Tuning modification. ρR, T_{ion} and $P_{hotspot}$ are burn averaged.

Implosion Metric	Baseline (N210808)	Peak Only	SAI Only	Peak + Small Dip	Peak + SAI
Yield (MJ)	1.15	5.21	6.37	9.39	11.9
ρ R (g/cm^2)	0.554	0.619	0.699	0.720	0.840
Adiabat	4.17	2.81	3.03	2.39	2.22
T_{ion} (keV)	8.3	15.5	15.4	25.2	24.3
$P_{hotspot}$ (Gbar)	253	627	548	1280	1160
max v_{shell} (km/s)	436	443	333	416	332
IFAR (75%)	25.3	34.5	16.7	25.8	21.6
Fall Line Radius (μm)	14.8	19.8	-9.8	25.5	1.8

could be addressed in future work by using a later, higher yield, implosion design[14] as the baseline. There is also a risk that the flux limiter tuning method discussed in section 3.2.1 could become less valid for higher yield implosions, artificially inflating yields.

These issues can be resolved by running simulations without alpha heating, the results of these simulations are shown in table 5.4. The metric of most note here is the ρR, which shows considerable improvement over the baseline, demonstrating that the increased yield is as a result of improved compression, which is consistent with the idea that the higher yield comes from reducing the adiabat and therefore keeping the fuel compressible. This suggests that the results are not simply a consequence of being on the threshold of ignition, these are genuine performance gains made with these methods. As for the additional metrics: removing alpha heating does not affect IFAR Or implosion velocity, so they match alpha-on simulations. Fall Line Radius is increased in all cases due to the bang time (measured as the peak yield) being significantly earlier in time. Yield, hotspot pressure and ion temperature are significantly reduced from the alpha-on simulations as there is no energy input from fusion reactions; and are included for completeness.

5.11 Pulse Shape Sensitivity Study

Simulations were undertaken to check the robustness of the modified radiation temperature profile against errors in pulse shaping. This was done by running a large number of simulations, each with small perturbations to the pulse shape; which had a similar structure and amplitudes to those described in NIF pulse shaping documents [131]. These were modelled as fractional shifts in the laser pulse shape, then a simple model mapping laser energy to

Table 5.4: Table of implosion metrics for alpha-off simulations, directly comparing simulations with and without the N+1 Tuning modification. ρR, T_{ion} and $P_{hotspot}$ are burn averaged.

Implosion Metric	Baseline	Peak	SAI	Peak + Small Dip	Peak + SAI
Yield (kJ)	85.7	102	66	131	75.1
ρ R (g/cm^2)	0.691	0.852	1.057	0.948	1.155
Adiabat	4.16	2.82	3.04	2.39	2.22
T_{ion} (keV)	4.67	4.50	4.08	5.38	4.35
$P_{hotspot}$ (Gbar)	156	215	159	234	181
max v_{shell} (km/s)	436	443	333	416	332
IFAR (75%)	25.3	34.5	16.7	25.8	21.6
Fall Line Radius (μm)	48.6	45.5	28.2	37.8	32.6

hohlraum temperature was used ($I \propto T^4$) to calculate the shifts in hohlraum temperature based on those perturbations. The effects of laser errors would likely be reduced further if a more realistic hohlraum model was used to account for effects like thermal smoothing, which adds confidence that real errors will be at least as small if not smaller than the errors given by the model used here.

Once the perturbed pulse shape had been generated, the exact error was then calculated using the following equation:

$$error\% = 100 \times \left\langle \frac{|T_p(t)^4 - T_b(t)^4|}{T_b(t)^4} \right\rangle$$

Where $T_b(t)$ is the baseline radiation temperature profile at timestep t, and $T_p(t)$ is the perturbed radiation temperature profile with errors added at timestep t. The error is averaged over the section of the radiation temperature profile where the laser would be expected to be active, so the long tail after the structured pulse does not count towards the error.

Example pulse shapes are shown in figure 5.17, labelled by their approximate error.

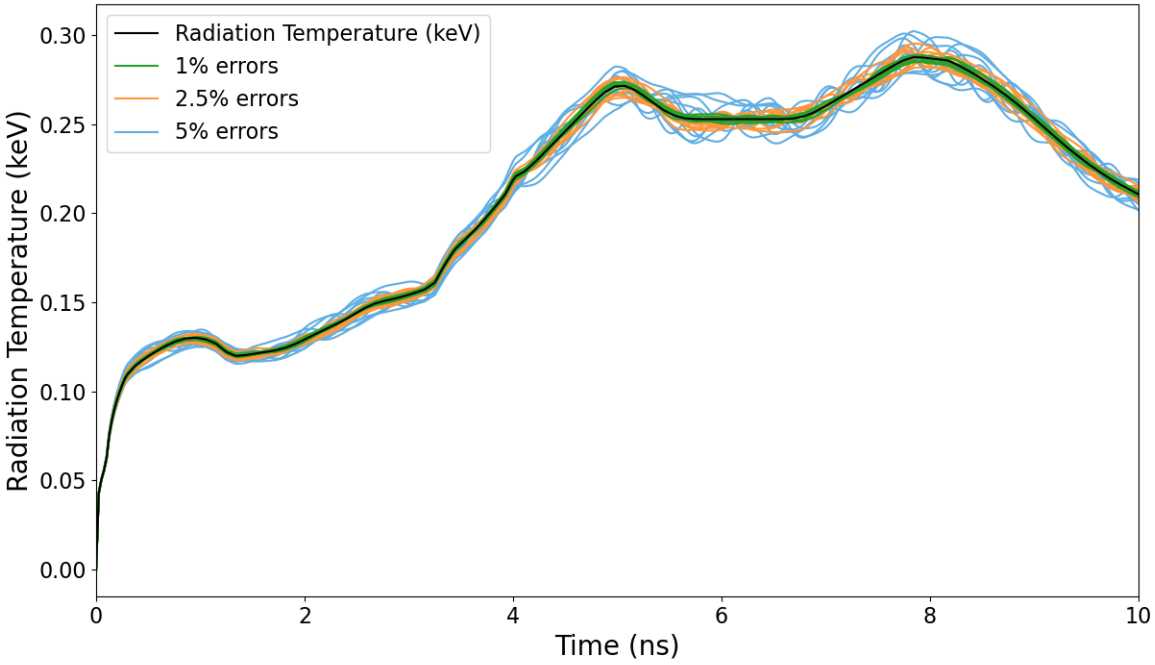


Figure 5.17: An example pulseshape, with a number of profiles with different errors applied

Simulations were ran with a range of errors, and the figures 5.18 and 5.19 show how the yield and ρR decrease with increasing pulse shaping errors for both the Peak-Only N+1 Tuning design and combined N+1 Tuning & SAI.

Since the errors on NIF's current pulse shaping system are less than 0.5% [131], and N+1 tuning still produces improved yield with errors as high as 2%, this suggests the scheme is robust against pulse shaping errors.

Some of the simulated errors result in improved performance; this is due to a combination of effects, primarily improved timing of the first three shocks, which are not perfectly retimed to converge on the ice-gas interface. Additionally, the random errors were allowed to deviate from physically achievable pulse shapes, and did not conserve total energy, so are able to increase energy in ways which the optimisation was prevented from using because they are not experimentally achievable. In the case of N+1 Tuning only, the more extreme random errors were also able to introduce small dips which improved performance as expected.

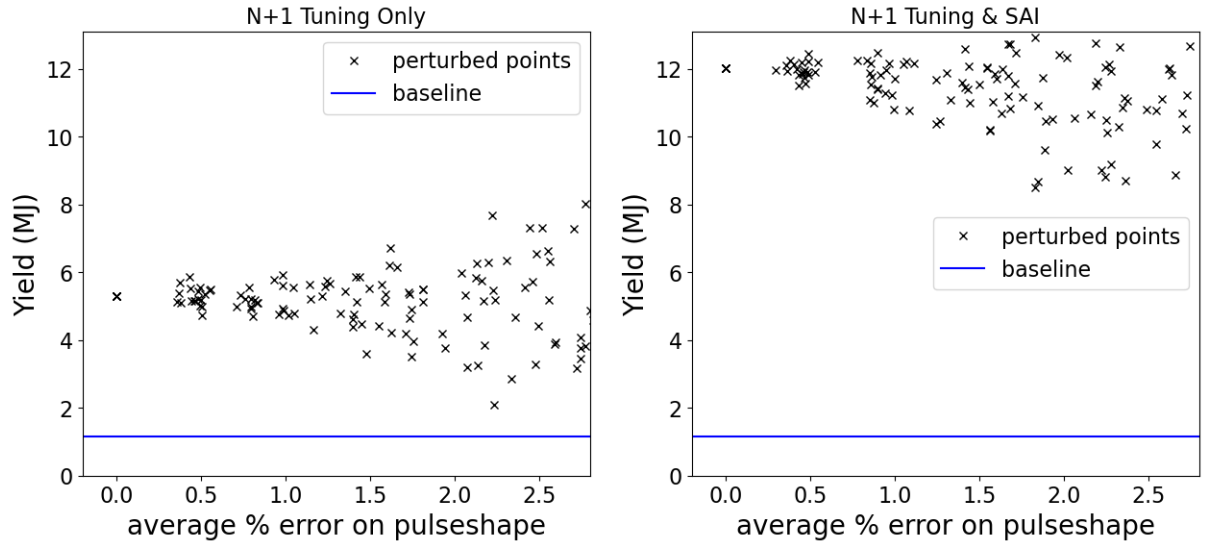


Figure 5.18: Yield(MJ) for different magnitudes of random error applied to the pulseshape

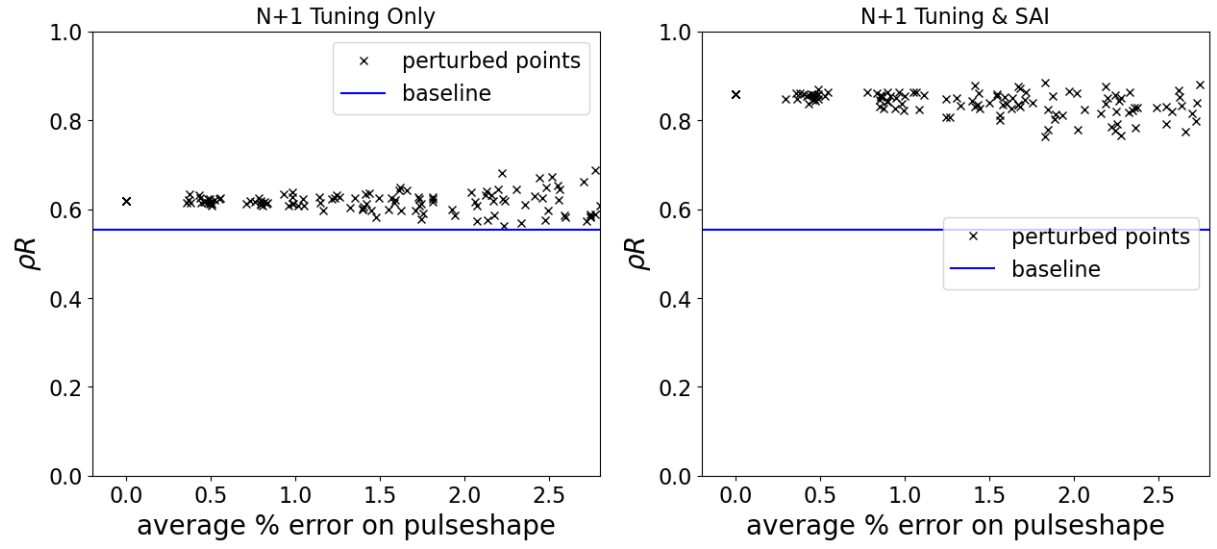


Figure 5.19: ρR for different magnitudes of random error applied to the pulseshape

5.12 Chapter Summary

This chapter has described the formation of the N+1 shock and motivation for tuning it, before describing a mechanism by which tuning the N+1 shock should increase yield. Implosion design are presented, which show reduced adiabat, increased compressability and increased yield. Along with a timing scan, these implosions support the discussed mechanism.

The N+1 Tuning designs are also explored in the context of hohlraum constraints, suggesting that redesigning hohlraums with N+1 Tuning in mind may provide additional routes to higher yield. The N+1 Tuning design is also combined with a shock augmented ignition design, and these designs are shown to be complimentary. This combination provides a 10x increase in 1D yield over the baseline simulation, as well as improvements in a range of fusion metrics.

Alpha-off simulations are also presented, demonstrating that the increase in yield is likely a result of increased compressability and not merely a consequence of being close to the threshold of ignition. Future work examining implosions further across the threshold of ignition[14] would help to verify whether this effect continues to benefit higher yield implosions. Finally, a pulse shaping sensitivity study is presented, which shows that the increase in yield from these designs is robust against pulse shaping errors typical of NIF-like laser systems.

Overall, this chapter presents a promising way forward to increase yield using existing experimental facilities.

Chapter 6

N+1 Imaging Experiment Design

6.1 Chapter Overview

So far, this thesis has presented a code which can be used to produce synthetic X-ray phase contrast images of problems in the IFE context, and HYADES simulations suggesting that tuning the N+1 shock would increase the yield of indirect drive ICF implosions. To support the work on tuning the N+1 shock, this chapter will describe the use of the GREENER code to design an experiment to image the N+1 shock using XPCI. Simulations suggest the N+1 shock exists and has significant effects, but it has not yet been imaged, and VISAR cannot detect it (as discussed in section 2.4); this means its existence is yet to be experimentally verified. By imaging the N+1 shock, its formation and evolution can be studied and understood better, allowing for increased confidence in the hypothesis that it can be tuned for higher yield. Even in the case that the N+1 shock cannot be tuned as outlined earlier, improving understanding of implosion dynamics and reflected shocks within indirect drive implosions would push the field forward, allow for improved implosion design and verification of equation of state models.

XPCI offers the best method to image the N+1 shock since, as was discussed in section 2.3.5, it highlights steep gradients in refractive index, which shocks and interfaces cause. In addition, other techniques like VISAR cannot image at the late time when the N+1 shock occurs, as discussed in section 2.4. There has already been some work on using XPCI to observe shocks in capsules by Do et al [16], but the N+1 shock is yet to be experimentally imaged.

This chapter will start by exploring the constraints on implosion design and proposing

two implosions (flat-top and big-dip). With current technology it is not possible to image the N+1 shock in a high yield implosion, due to the high shell velocity when it occurs; these designs aim to make the N+1 shock easier to image. Simulations of these implosions will then be used to produce synthetic X-ray phase contrast images of the shell during the N+1 shock's transit, which can be evaluated with a contrast based figure of merit defined in this Chapter. Material in this Chapter will show that whilst the first (flat-top) implosion design does initially appear promising, unfortunately, the implosion is too fast and the N+1 shock appears too briefly to be effectively imaged. The second design (big-dip) is better, and extends the time during which the N+1 shock is present in the ice as well as slowing the overall shell velocity, meaning that a similar experimental setup to the one built by Do et al.[16] could be used to image the N+1 shock.

6.2 Implosion Design

The N+1 shock presents a significant challenge for imaging efforts. At the time of the shocks formation, the shell velocity is very high, (~260km/s in the N210808 simulations used earlier), and the shock velocity is also high (~60km/s in the N210808 simulation). This results in significant motion blur if the detector temporal resolution is too low. For a 100ps integration time, this gives a motion blur of more than 30μm. The shock is also relatively weak until late time, which means the shock front is not sharp enough to generate significant phase contrast effects. There are two key ways to improve this, reducing the shell velocity allows for reduced blurring, making the shock easier to pick out; while strengthening the shock increases the sharpness of the gradient in density, therefore strengthening the phase contrast effect. The simpleset way to reduce the shell velocity is by reducing the radiation impinging on the target. Unfortunately, there are some constraints on what changes can be made to the radiation temperature profile. In particular, the early time structure of the profile must remain largely unchanged, both to ensure the N+1 shock forms in the first place, and to ensure the conditions of the implosion are relevant for NIF style ICF implosions. As this thesis has largely focused on N210808 as the baseline, it will be used again for the baseline design in this section.

After the initial 3 shocks are launched, the radiation temperature profile can be altered in many ways while remaining relevant for ICF. The total incident energy on the implosion can

be reduced by cutting off the rise of the radiation temperature profile after the first 3 shocks have been launched, this reduces the shell velocity making it easier to image. In addition, the N+1 tuning knowledge (see section 5.7) can be applied to design the radiation temperature profile so that the shock forms and merges earlier, extending the time available to image the shock.

In this section two designs will be presented, both preserve the first 3 shocks and structure of the N+1 shock from the N210808 shot, so are relevant for NIF style ICF implosions, and both have reduced shell velocity to make the shock easier to image.

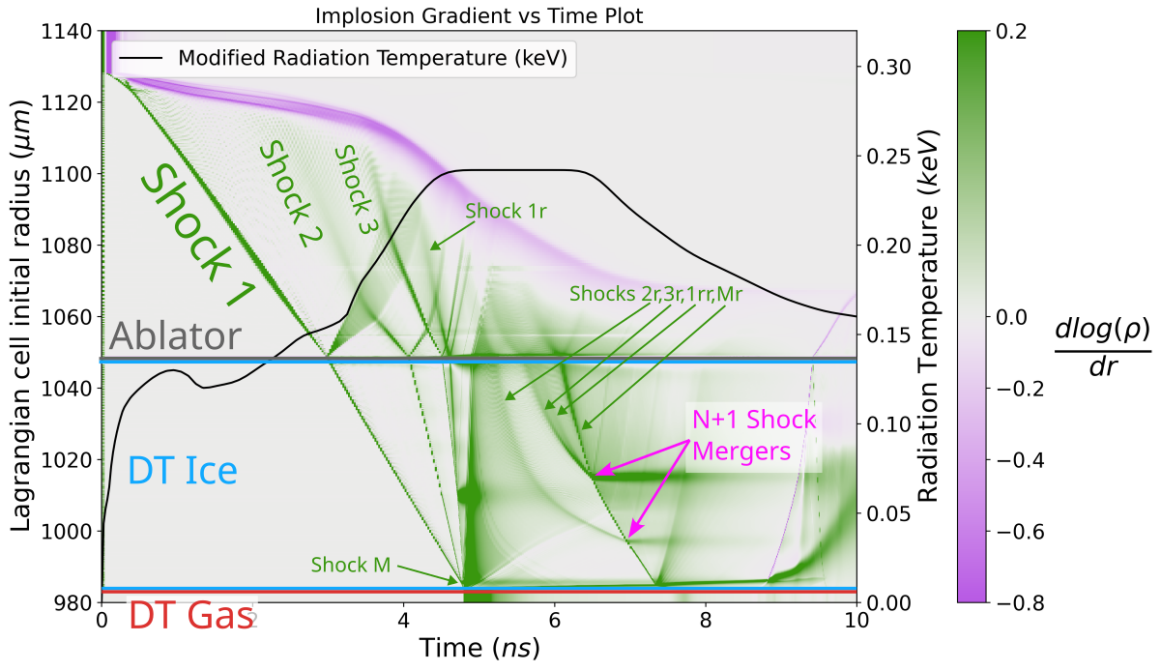


Figure 6.1: A Lagrangian shock plot showing the 'flat-top' design, which is designed to make imaging the N+1 shock easier. The radiation temperature profile is also shown, and key features are labelled.

The first (referred to as the 'flat-top' design) is shown in figure 6.1, in this case the radiation temperature profile remains unmodified until after all the main shocks have been launched and began propagating, at which point the radiation temperature profile is levelled off to minimise the additional kinetic energy introduced to the shell. This also causes the N+1 shock to merge earlier than it would otherwise, but the most significant effect is to reduce the shell velocity. Of the two designs, this is the one which most faithfully reproduces the formation of the N+1 shock from the N210808 implosion.

The second (referred to as the 'big-dip' design) is shown in figure 6.2, this is a further

modification to the profile shown in figure 6.1. The further modification is that a dip has been introduced to the radiation field between 4.5ns and 6ns, this dip is timed such that the early reflected shocks are weakened to make it easier for the later reflected shocks to catch up with them. The ramp up then strengthens the later reflected shocks, further aiding them in catching up with the earlier reflected shocks. This combined effect means that the N+1 shock merger happens much earlier than it would otherwise, this leads to a more extended period of time when there is a single clear N+1 shock travelling through the ice layer. This is intended to make the shock easier to image. The overall implosion velocity is also reduced, which acts to reduce the effect of motion blur when imaging the shock. This design does have downsides, the tuning weakens the N+1 shock slightly, making the phase contrast effect weaker, and timing the dip correctly requires knowledge of the timings of shock reflections, so this design is likely to be more experimentally challenging to produce.

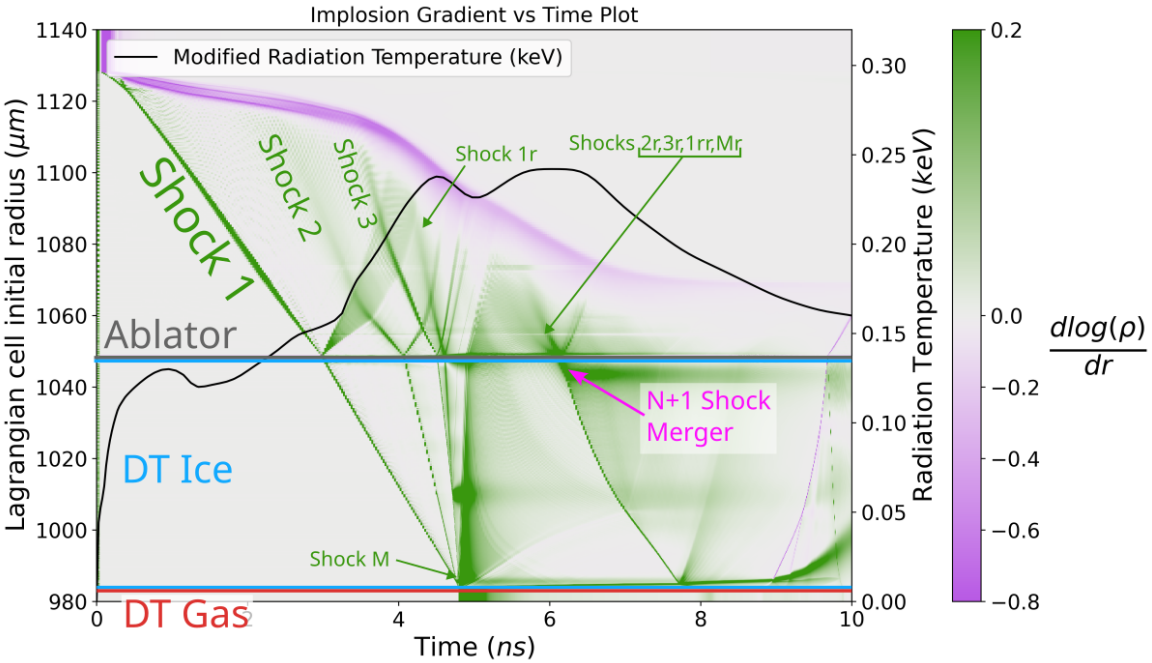


Figure 6.2: A Lagrangian shock plot showing the 'big-dip' design, which is designed to make imaging the N+1 shock easier. The radiation temperature profile is also shown, and key features are labelled.

The key differences between the flat-top and big-dip designs are that the flat-top design is more robust to equation of state uncertainties or other experimental uncertainties which might change the shock timing from the expected values. The big-dip design, while more sensitive to uncertainties, results in a shock which merges earlier and has a longer propagation

time in the ice layer, so is expected to be easier to image.

6.3 Synthetic Images and a Figure of Merit

Using the GREENER code presented in chapter 4, synthetic X-ray phase contrast images can be generated of the 'flat-top' and 'big-dip' HYADES simulations. The parameters for the imaging setup were taken from Do et al. [16, 120], along with private communications with Do and Dewald. The source was a 5μm 9keV zinc Heα source, by choosing a high energy source the implosion's self emission can more easily be filtered out, since it should be higher energy than the bulk of the self-emission spectrum. The source was placed 13mm behind the capsule, while the detector was placed 905mm away from the capsule. The detector pixels were 9x9μm, but with an extra 1.25x magnification from the streak camera, so the pixel size used in GREENER was 7.2x7.2μm to account for this.

A lineout of the synthetic image taken at 7.31ns after the laser turns on is shown in figure 6.3. The XPCI lineout is shown in blue, and the absorption only lineout is shown in orange, with a shaded region showing 3 standard deviations away form the absorption line; where the distribution was taken to be Poisson, so for N photons hitting a pixel on the detector, $\sqrt{N}$ was taken as the standard deviation. Along with the lineouts, the density profile of the implosion is shown in black, with key interfaces labelled. The density profile appears stepped as in GREENER it is constructed as a series of layered spheres, using the method discussed in section 4.3. The points chosen to construct the layers were based on the HYADES simulation, as a consequence the ablation front is slightly under sampled, but the areas of interest have high enough spatial resolution to have confidence in the results in this chapter. Appendix C discusses convergence testing in more detail.

The N+1 shock can be seen relatively clearly on the XPCI lineout as a significant deviation from the absorption only lineout around 595-615μm, although there is significant overlap in the error regions shown. It is also worth noting that this is only a single timestep snapshot and there are no motion blur effects, these will be taken into account later and reduce the contrast. The benefits of XPCI on contrast are also visible at other interfaces, which are much clearer than they are in the absorption only lineout.

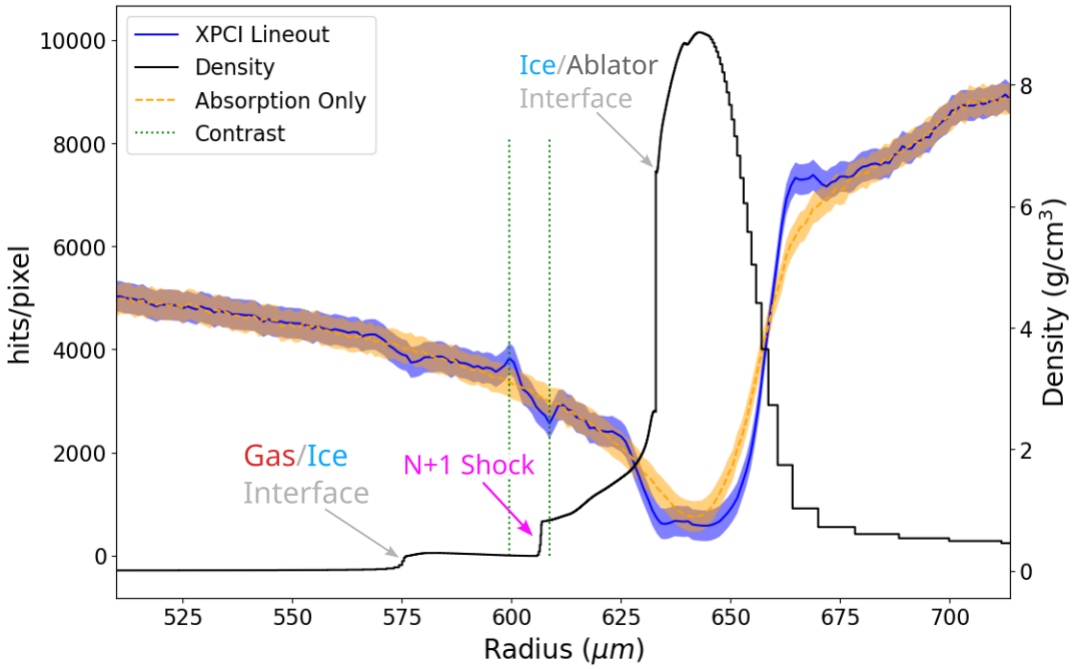


Figure 6.3: Example synthetic XPCI lineout of the big-dip design at t=7.31ns into the implosion, the density used is also plotted for comparison. The ice-gas and ice-ablator interfaces are labeled, as is the N+1 shock. An absorption only lineout is also shown. The two shaded regions represent 3 standard deviations from their respective lineouts using a poission error. The spike and dip used in contrast calculations are identified with vertical green dotted lines.

To quantitatively evaluate what is the 'best' image, a figure of merit is needed. There are a few standard options for defining signal or contrast to noise ratios. The metric needs to be easy to calculate automatically across many timesteps and different configurations, and measure how easy it is to identify the spike and dip associated with the N+1 shock interface (the spike and dip are marked with green dotted lines in figure 6.3). The simplest metric would be to take the difference in intensity between the spike and dip in the phase contrast data, the problem with this is that the spike and dip are deviations from a slope, since the absorption only image is already sloped. So for a simulation without an N+1 shock, this metric would show some contrast existing (even if the N+1 shock wasn't present). Another metric could be to fit a curve to the lineout either side of the shock, and compare the spike/dip deviation from that curve, this has the benefit that it doesn't require an absorption only image to do the calculation, but has the downside that it would be challenging to do in an automatic way across many timesteps, and that if the shock is too close to another interface there won't be sufficient points to get a good fit. The method chosen for this work was to measure how far the

XPCI spike and dip deviate from the absorption only lineout, which allows easy automatic identification of the shock, and allows the metric to be calculated across many timesteps. In an actual experiment either a fit or synthetic diagnostic data would be required to do a similar comparison.

The equation for contrast to noise ratio can be given as [132]:

$$CNR = \frac{|I_{xpci}(r_{spike}) - I_a(r_{spike})| + |I_{xpci}(r_{dip}) - I_a(r_{dip})|}{\sqrt{\sigma_a^2(r_{spike}) + \sigma_a^2(r_{dip})}}$$

where $I(r)$ refers to the intensity at a given radius, with the subscripts a and $xpci$ referring to the absorption only and XPCI quantities respectively, σ_a refers to the standard deviation of the absorption only intensity (calculated using the shot noise $\sqrt{N}$). The positions r_{dip} and r_{spike} refer to the radii where the dip and spike of the intensity profile occurs on either side of the N+1 shock, these are marked with green dotted lines on figure 6.3. Assuming the noise is approximately equal at both the spike and dip locations, this simplifies to:

$$CNR = \frac{|I_{xpci}(r_{spike}) - I_a(r_{spike})| + |I_{xpci}(r_{dip}) - I_a(r_{dip})|}{\sqrt{2}\sigma_a}$$

This metric takes advantage of one of the key benifits of XPCI: the structured nature of the signal. While it would be difficult to confidently detect a single spike that is only slightly higher than 3 sigma, it is much easier to detect a spike and a dip which are close to each other. A more advanced analysis would take advantage of the fact that the structure consistently moves over time, while shot noise is random, this would allow weaker signals to be picked out of the noise.

To automatically find the contrast to noise ratio across many timesteps, the position of the N+1 shock is found in the hyades data, then in the phase contrast image, the closest spike and dip pair to the N+1 shock are found, the contrast to noise ratio is then calculated. The criteria for a spike or dip to be detected by the algorithm is that it must deviate from the absorption lineout by at least 3 standard deviations, which means that spikes and dips with signal to noise ratio's below 3 are not detected. Unfortunately using fewer standard deviations to find the spike and dip would allow smaller structures which are not part of the N+1 shock to confuse the algorithm which looks for the nearest spike and dip.

6.4 Interpolating Synthetic Images into a Time Series

A number of data dumps were output from the HYADES simulations (big-dip and flat-top), with higher frequency data dumps while the N+1 shock forms and travels through the ice. Synthetic images were produced for each of these dump times and can be used to emulate streak images of the implosions.

The streak image for the big-dip implosion design is shown in the top left panel of figure 6.4, and the lab frame shockplot for the same region is shown in the top right panel for comparison.

The streak image has limited resolution due to only having a few dumps output, in addition, since a full XPCI simulation must be run for each dump time, getting a higher resolution image by running more timesteps is potentially computationally expensive, especially as hundreds of timesteps would be required to produce a smooth image.

Instead, higher effective resolution can be obtained by interpolating along the Lagrangian cell trajectories. Since features in the images correspond to physical material, and the trajectories of packets of material are known from the Lagrangian cell trajectories, it is reasonable to assume that the features in the images should follow the Lagrangian cell trajectories (at least over short times). By interpolating over the trajectories, and finding intensity values for intermediate times, the image in the bottom panel of figure 6.4 can be produced, which is much higher resolution.

The same can be done for the flat-top design, and is shown in figure 6.5.

With these higher resolution streaks, the effects of detectors having non-zero integration time windows can be examined, the streaks can be blurred in time to simulate an integration time. Blurring is achieved by adding the intensity value at a given radius across the multiple discrete timesteps, the higher effective resolution is needed because in reality this process would be a continuous integration, so the timesteps need to be fine enough to be a reasonable approximation. Figure 6.6 shows streaked images of the big-dip and flat-top designs, with blurs of 50ps or 100ps applied, chosen as 100ps was the blur on the streak camera used by Do. et. al. [16]. The 50ps blur still shows the shock and ice-gas interface, however for a 100ps blur the shock and ice-gas interface are much harder to pick out.

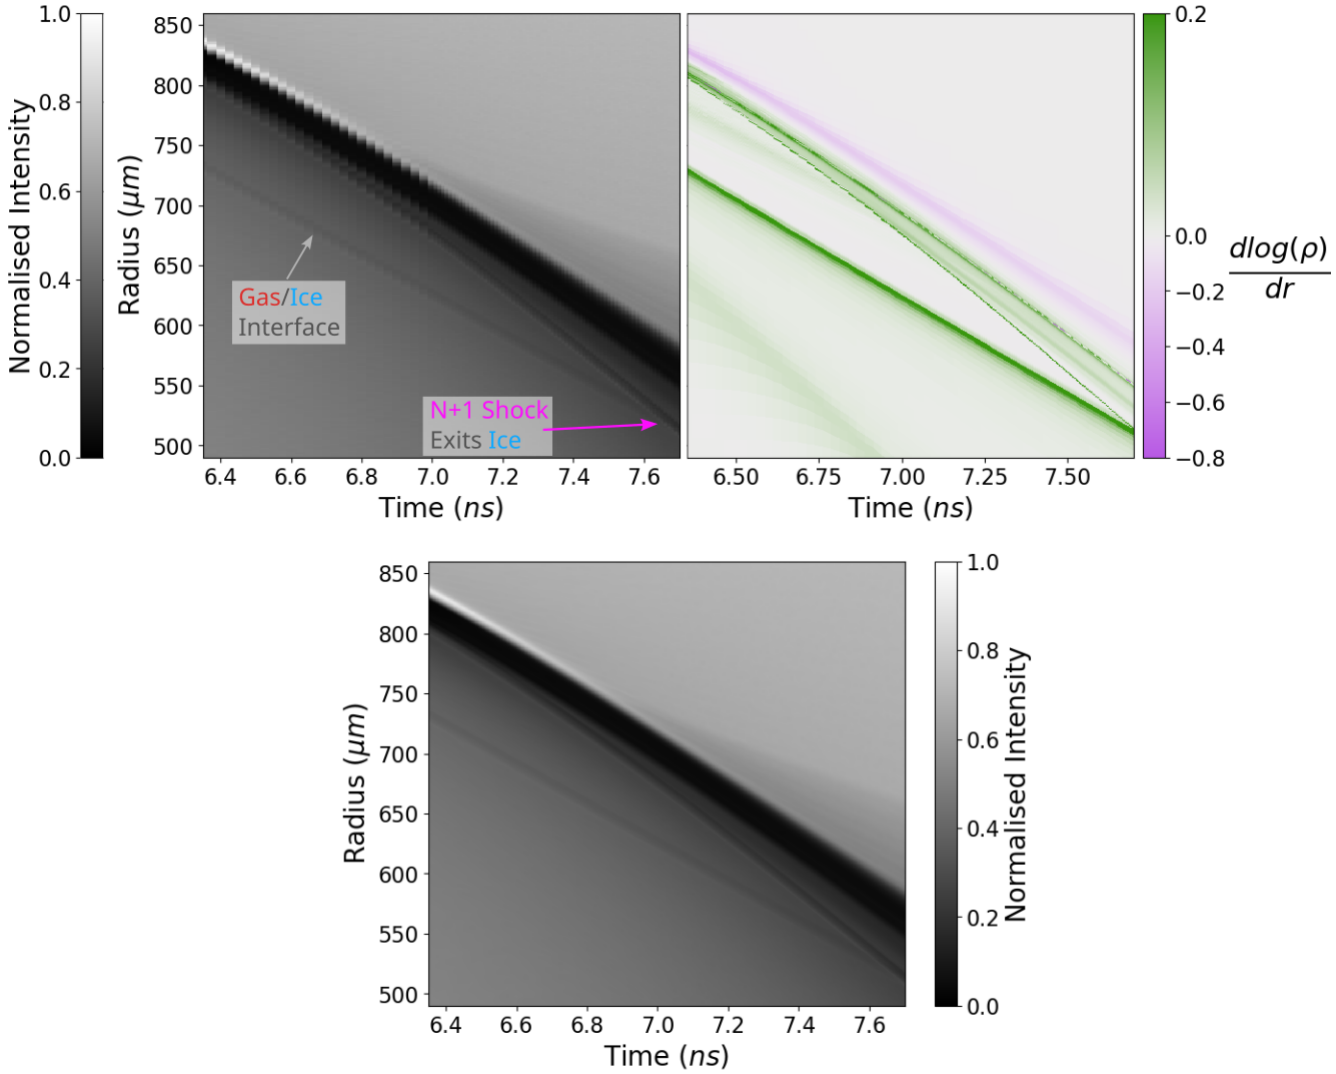


Figure 6.4: A sequence of synthetic XPCI lineouts over time for the big-dip implosion(top left), paired with the shockplot of the big-dip implosion. The trajectories of the Lagrangian cells in the shockplot are used to interpolate the intermediate values of the synthetic XPCI lineouts, to generate the bottom panel.

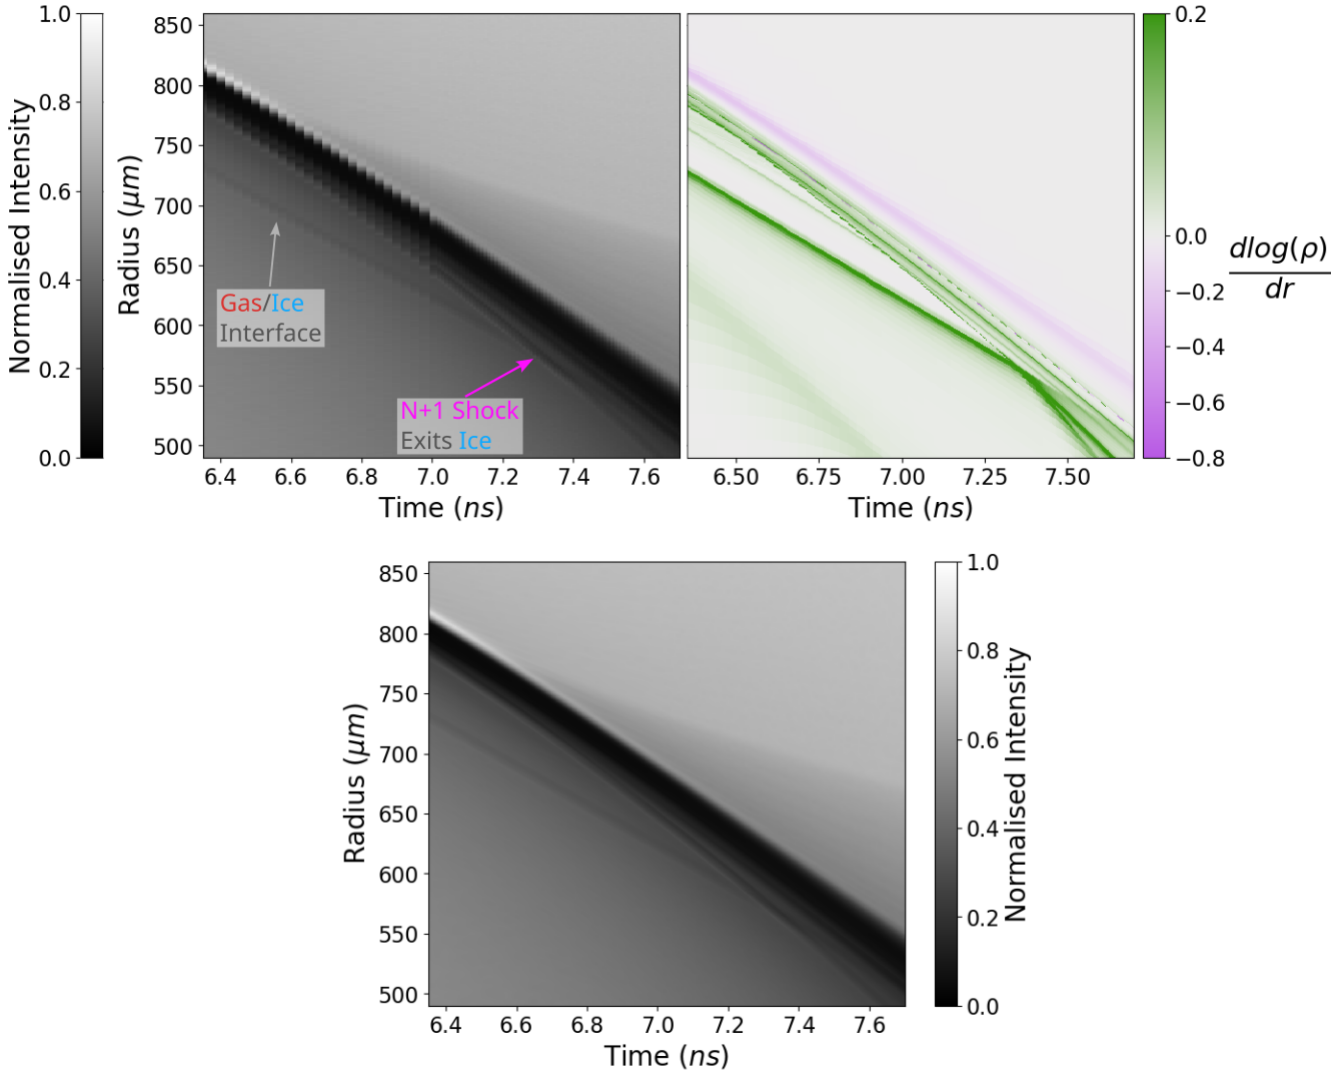


Figure 6.5: A sequence of synthetic XPCI lineouts over time for the flat-top implosion(top left), paired with the shockplot of the flat-top implosion. The trajectories of the Lagrangian cells in the shockplot are used to interpolate the intermediate values of the synthetic XPCI lineouts, to generate the bottom panel.

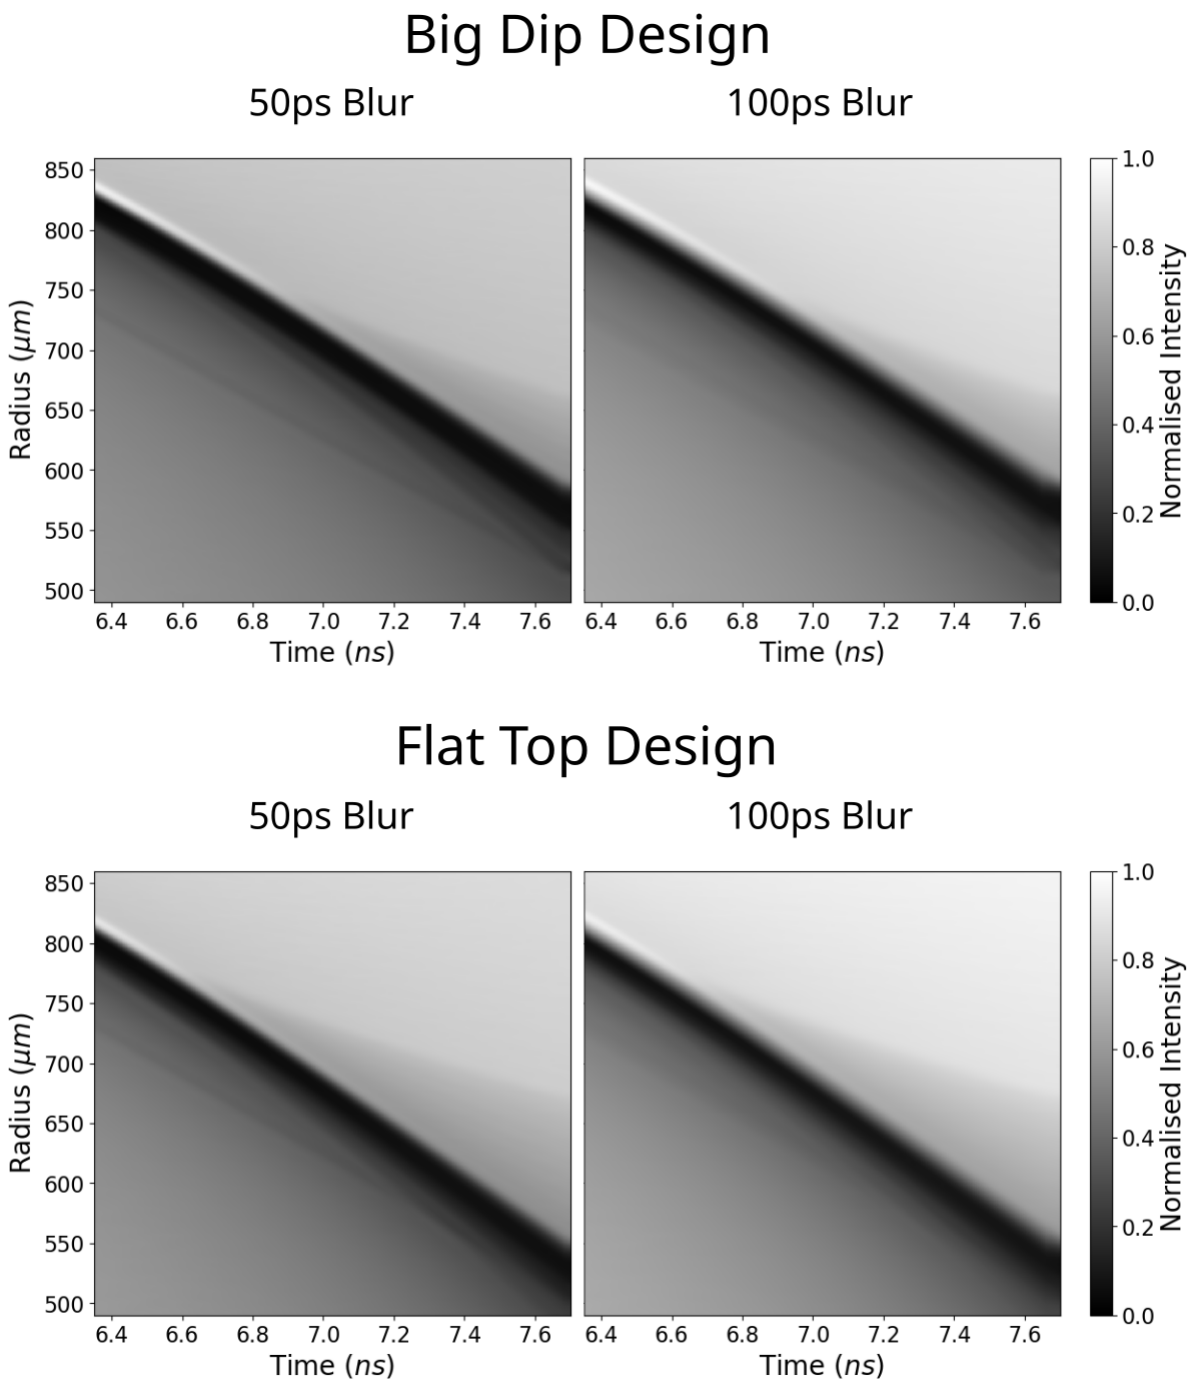


Figure 6.6: Sequences of synthetic XPCI lineouts plotted over time for the big-dip design (top panels) and the flat-top design (bottom panels). Shown with a temporal blur which simulates the effect of a detector integration time, the left panels show a 50ps blur while the right panels show a 100ps blur.

6.5 Contrast Over Time

Combining the contrast to noise metric, and the interpolated streak images, values of contrast to noise ratio over time can be plotted for both designs, where a value of 0 is used for timesteps where the shock location cannot be determined for use in the contrast to noise ratio calculation. The plot of contrast to noise ratio vs time for the big-dip design is shown in figure 6.7, along with the contrast to noise ratio vs time for streak images which have had 50ps or 100ps temporal blurs applied.

A number of features are visible, and the explanations of these features are most easily understood by comparing the contrast to noise ratio over time to the synthetic streak and shockplot shown in figure 6.4. Initially, the shock is still in the ablator, so the structure is swamped by the XPCI effects of the bulk of the shell, around 6.6ns the shock is far enough into the ice to be detectable as a separate feature, as it passes through the ice, high contrast is maintained until it reaches the ice-gas interface. During this region, the contrast fluctuates with a frequency matching the data dump times, suggesting this is a result of the interpolation rather than a real effect. As the shock approaches the ice-gas interface (around 7.5-7.6ns), the shock signal and ice-gas interface signal begin to overlap, resulting in the contrast initially reducing as the spike from the N+1 shock overlaps with the dip from the ice-gas interface before increasing again once the two features start to overlap constructively and form a sharper interface (just after 7.6ns) just before the shock breaks out into the gas layer. With a 50ps blur the story is similar, although the shock takes longer to become detectable and has lower contrast to noise throughout. There is still a region of over 0.7ns with consistent contrast, which could be used to track the shock. With a 100ps blur the contrast is too low to reliably detect.

A similar plot can be produced for the flat-top design, and is shown in figure 6.8, here the same axes have been used as the big-dip design, which makes it clear that while the flat-top design has slightly higher contrast, the time range over which the shock is detectable is severely reduced from 1ns to 0.4ns. This is made worse by blurring, for a 50ps blur the window where the contrast rises above the noise is limited to a very narrow range. Both the increased contrast and reduced time window are caused by the fact that the flat-top design drives the implosion harder than the big-dip design. This means the N+1 shock is stronger and therefore a sharper interface, so has better contrast in an instantaneous snapshot image,

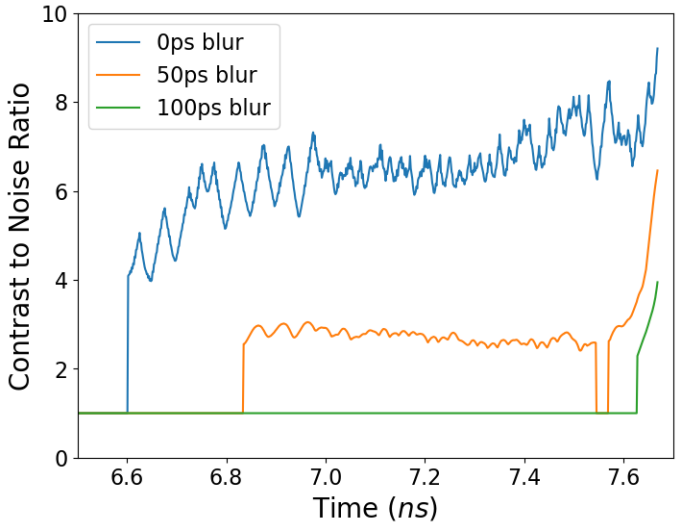


Figure 6.7: Contrast-to-noise-ratio vs time for the selected time range on the big-dip design, showing the effect of a 50ps and 100ps blur.

but since the whole implosion is going faster, the shock takes less time to go through the ice, and the effect of motion blur is more significant.

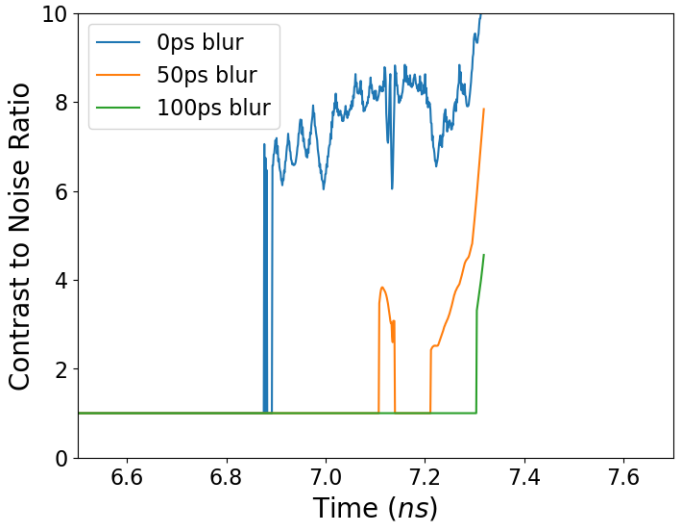


Figure 6.8: Contrast-to-noise-ratio vs time for the selected time range on the flat-top design, showing the effect of a 50ps and 100ps blur.

In both designs, the 100ps blur makes it difficult to find the shock, this is because with a large enough blur, the peaks of the signal at t=0ps will overlap with the troughs of the signal at t=100ps, such that the net effect smears the feature out of existence. This time smearing effect is not based on photon count, so lowering the Poisson noise by increasing the number of photons only helps to a limited extent, since the signal is overlapping with itself and cancelling

itself out. On the other hand, a 50ps blur on the big-dip design still shows moderate contrast over an extended window of time, suggesting that with a sufficiently fast detector it would be possible to image the N+1 shock in this configuration. While the x-ray phase contrast imaging setup used by Do et al[16] does not have sufficient temporal resolution, there are electron pulse-dilation based detectors which can achieve much higher temporal resolutions (up to 5ps)[133], suggesting that imaging the N+1 shock could be possible with a different experimental setup. Further research could look to explore this and further refine the imaging requirements.

6.6 Chapter Summary

This chapter has presented synthetic streaked images of two implosions whose radiation temperature profiles were designed to make imaging the N+1 shock easier. A contrast metric was defined and used to show that both the flat-top and big-dip designs allow the N+1 shock to be detected at single instantaneous snapshots using the experimental setup used in Do et al.[16]. Looking at the contrast over time for both designs, the big-dip design maintains moderate contrast over an extended period of more than 0.5ns, even with a detector integration time of 50ps, although longer integration times were shown to blur out the signal such that the shock could no longer be detected. This work suggests that imaging the N+1 shock could be achievable using the big-dip design and a sufficiently fast detector.

Chapter 7

Conclusion

A synthetic diagnostic code (GREENER) was presented, which was written within the GEANT4 framework to simulate X-ray Phase Contrast Imaging (XPCI). It is a Monte-Carlo based raytracer, which can handle complex geometry imported from 1D, 2D and 3D hydrodynamic simulations; surface normals are preserved such that raytracing can be performed with accurate refraction effects. Synthetic images of a range of simulations are presented; the utility of this code for experiment design and analysis is also demonstrated.

This thesis also presented implosion designs that offer the potential to significantly increase the yield of indirect drive implosions relative to the N210808 National Ignition Facility (NIF) shot, which is used as a baseline. One-dimension (1D) radiation hydrodynamic simulations were presented in which the proposed modification increased the yield by a factor of 10, this is achieved by tuning the implosion to delay the N+1 shock merger, therefore reducing the shell adiabat and as a result increasing areal density by a factor of 1.5. N+1 tuning alone increases the yield in simulations from 1.2MJ to 9.4MJ, combining N+1 tuning and the shock augmented ignition scheme gives a yield of 12MJ. Although the simulations presented are 1D, they also show equivalent or reduced in-flight aspect ratio (IFAR) and implosion velocity, indicating reduced susceptibility to 2D Rayleigh-Taylor instabilities relative to the base design of the N210808 NIF shot. Further simulations with alpha heating turned off also show improved areal density and IFAR. A pulse shaping sensitivity study is also presented indicating that this design is robust against the typical pulse shaping errors of a NIF-like laser system.

Finally, these threads were brought together by using the presented synthetic diagnostic code to design an experiment to image the N+1 shock on an indirect drive implosion, this

experiment design showed that by modifying the radiation temperature profile, the N+1 shock can be strengthened while slowing the overall implosion velocity, making it substantially easier to image. A contrast metric was defined and used to compare multiple designs and provide evidence that imaging the N+1 shock could be achievable. This proposed experiment would verify the underpinning physics of the N+1 shock and pave the way for applying the N+1 tuning scheme to ignition scale experiments.

Bibliography

[1] Methodist Church. "409: Let Us Build a House Where Love Can Dwell and All Can Safely Live". In: *Singing the Faith: Words Edition*. London: Canterbury Press Norwich, 2013, p. 701. ISBN: 978-1-84825-068-0.

[2] Quakers in Britian. "24.49". In: *Quaker Faith & Practice: The Book of Christian Discipline of the Yearly Meeting of the Religious Society of Friends (Quakers) in Britain*. Fifth edition. The Yearly Meeting of the Religious Society of Friends (Quakers) in Britain, 2013. ISBN: 978-1-907123-54-2.

[3] Methodist Church. "701: Heaven Shall Not Wait for the Poor to Lose Their Patience". In: *Singing the Faith: Words Edition*. London: Canterbury Press Norwich, 2013, p. 701. ISBN: 978-1-84825-068-0.

[4] Katherine Calvin et al. *IPCC, 2023: Climate Change 2023: Synthesis Report. Contribution of Working Groups I, II and III to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change [Core Writing Team, H. Lee and J. Romero (Eds.)]. IPCC, Geneva, Switzerland*. Tech. rep. Intergovernmental Panel on Climate Change (IPCC), July 2023. DOI: 10.59327/IPCC/AR6-9789291691647.

[5] Hyun Kim et al. "Health Risks of Climate Change in the 21 Pacific Island States and Noted Gaps in Scientific Evidence: A Scoping Review". In: *The Journal of Climate Change and Health* 8 (Oct. 2022), p. 100166. ISSN: 26672782. DOI: 10.1016/j.joclim.2022.100166.

[6] G. Supran, S. Rahmstorf, and N. Oreskes. "Assessing ExxonMobil's Global Warming Projections". In: *Science* 379.6628 (Jan. 2023), eabk0063. ISSN: 0036-8075, 1095-9203. DOI: 10.1126/science.abk0063.

[7] Mark G. Lawrence et al. "Evaluating Climate Geoengineering Proposals in the Context of the Paris Agreement Temperature Goals". In: *Nature Communications* 9.1 (Sept. 2018), p. 3734. ISSN: 2041-1723. DOI: 10.1038/s41467-018-05938-3.

[8] Hannah Ritchie. "Sector by Sector: Where Do Global Greenhouse Gas Emissions Come From?" In: *Our World in Data* (2020). https://ourworldindata.org/ghg-emissions-by-sector.

[9] Kevin A. Baumert, Timothy Herzog, and Jonathan Pershing. *Navigating the Numbers: Greenhouse Gas Data and International Climate Policy*. Washington, D.C.: World Resources Institute, 2005. ISBN: 978-1-56973-599-2.

[10] Department of Economic and Social Affairs. *The Sustainable Development Goals Report 2025*. 1st ed. The Sustainable Development Goals Report. Description based on publisher supplied metadata and other sources. Bloomfield: United Nations Research Institute for Social Development, 2025. ISBN: 978-92-1-107159-7.

[11] Hannah Ritchie. "The Number of People without Electricity More than Halved over the Last 20 Years". In: *Our World in Data* (2021). https://ourworldindata.org/number-without-electricity.

[12] Max Roser. "Why Did Renewables Become so Cheap so Fast?" In: *Our World in Data* (2020). https://ourworldindata.org/cheap-renewables-growth.

[13] J. Knaster, A. Moeslang, and T. Muroga. "Materials Research for Fusion". In: *Nature Physics* 12.5 (May 2016), pp. 424–434. ISSN: 1745-2473, 1745-2481. DOI: 10.1038/nphys3735.

[14] H. Abu-Shawareb et al. "Achievement of Target Gain Larger than Unity in an Inertial Fusion Experiment". In: *Physical Review Letters* 132.6 (Feb. 2024), p. 065102. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.132.065102.

[15] David S. Montgomery. "Invited Article: X-ray Phase Contrast Imaging in Inertial Confinement Fusion and High Energy Density Research". In: *Review of Scientific Instruments* 94.2 (Feb. 2023), p. 021103. ISSN: 0034-6748, 1089-7623. DOI: 10.1063/5.0127497.

[16] Alex Do et al. "Direct Measurement of Ice-Ablator Interface Motion for Instability Mitigation in Indirect Drive ICF Implosions". In: *Physical Review Letters* 129.21 (Nov. 2022), p. 215003. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.129.215003.

[17] Mordecai D. Rosen. "The Long Road to Ignition: An Eyewitness Account". In: *Physics of Plasmas* 31.9 (Sept. 2024), p. 090501. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/5.0221005.

[18] Matteo Barbarino. "A Brief History of Nuclear Fusion". In: *Nature Physics* 16.9 (Sept. 2020), pp. 890–893. ISSN: 1745-2473, 1745-2481. DOI: 10.1038/s41567-020-0940-7.

[19] William Happer. *Review of the Department of Energy's Inertial Confinement Fusion Program*. Tech. rep. National Research Council of Inertial Confinement Fusion, 1990. URL: https://fire.pppl.gov/ICF_Happer2.PDF.

[20] John Nuckolls et al. "Laser Compression of Matter to Super-High Densities: Thermonuclear (CTR) Applications". In: *Nature* 239.5368 (Sept. 1972), pp. 139–142. ISSN: 0028-0836, 1476-4687. DOI: 10.1038/239139a0.

[21] Paul-Henri Rebut. "The Joint European Torus (JET)". In: *The European Physical Journal H* 43.4-5 (Dec. 2018), pp. 459–497. ISSN: 2102-6459, 2102-6467. DOI: 10.1140/epjh/e2017-70068-y.

[22] John D. Lindl. *Inertial Confinement Fusion: The Quest for Ignition and Energy Gain Using Indirect Drive*. New York Berlin: Springer, 1998. ISBN: 978-1-56396-662-0.

[23] Stefano Atzeni and Jürgen Meyer-ter-Vehn. *The Physics of Inertial Fusion*. Oxford University Press, June 2004. ISBN: 978-0-19-856264-1. DOI: 10.1093/acprof:oso/9780198562641.001.0001.

[24] Christian Hill. "Nuclear Fusion Cross Sections and Reactivities". In: *Scientific Python* (Mar. 2021). https://scipython.com/blog/nuclear-fusion-cross-sections/.

[25] T. Giegerich et al. "Development of a Viable Route for Lithium-6 Supply of DEMO and Future Fusion Power Plants". In: *Fusion Engineering and Design* 149 (Dec. 2019), p. 111339. ISSN: 09203796. DOI: 10.1016/j.fusengdes.2019.111339.

[26] Stefano Atzeni and Debra Callahan. "Harnessing Energy from Laser Fusion". In: *Physics Today* 77.8 (Aug. 2024), pp. 44–50. ISSN: 0031-9228, 1945-0699. DOI: 10.1063/pt.zghg.fite.

[27] R. S. Craxton et al. "Direct-Drive Inertial Confinement Fusion: A Review". In: *Physics of Plasmas* 22.11 (Nov. 2015), p. 110501. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/1.4934714.

[28] O. L. Landen et al. "What next: Further Implosion Space Exploration on the Path to NIF Extended Yield Capability". In: *Physics of Plasmas* 31.6 (June 2024), p. 062712. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/5.0209027.

[29] H. Abu-Shawareb et al. "Lawson Criterion for Ignition Exceeded in an Inertial Fusion Experiment". In: *Physical Review Letters* 129.7 (Aug. 2022), p. 075001. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.129.075001.

[30] J. Zhang et al. "Double-Cone Ignition Scheme for Inertial Confinement Fusion". In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 378.2184 (Nov. 2020), p. 20200015. ISSN: 1364-503X, 1471-2962. DOI: 10.1098/rsta.2020.0015.

[31] Stefano Atzeni. "Inertial Confinement Fusion with Advanced Ignition Schemes: Fast Ignition and Shock Ignition". In: *Laser-Plasma Interactions and Applications*. Ed. by Paul McKenna et al. Heidelberg: Springer International Publishing, 2013, pp. 243–277. ISBN: 978-3-319-00037-4 978-3-319-00038-1. DOI: 10.1007/978-3-319-00038-1_10.

[32] R. H. H. Scott et al. "Shock-Augmented Ignition Approach to Laser Inertial Fusion". In: *Physical Review Letters* 129.19 (Nov. 2022), p. 195001. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.129.195001.

[33] R Paul Drake. *High-Energy-Density Physics*. Graduate Texts in Physics. Cham: Springer International Publishing, 2018. ISBN: 978-3-319-67710-1 978-3-319-67711-8. DOI: 10.1007/978-3-319-67711-8.

[34] Jeff Colvin and Jon Larsen. *Extreme Physics: Properties and Behavior of Matter at Extreme Conditions*. 1st ed. Cambridge University Press, Nov. 2013. ISBN: 978-1-107-01967-6 978-1-139-09515-0. DOI: 10.1017/CBO9781139095150.

[35] IA B. Zel'dovich and IU P. Raızer. *Physics of Shock Waves and High-Temperature Hydrodynamic Phenomena*. Originally published (in English, in 2 v.): New York : Academic Press, 1966-1967. With new pref. Mineola, N.Y: Dover Publications, 2002. ISBN: 978-0-486-42002-8.

[36] Philipp Rehner, Gernot Bauer, and Joachim Gross. "FeOs : An Open-Source Framework for Equations of State and Classical Density Functional Theory". In: *Industrial & Engineering Chemistry Research* 62.12 (Mar. 2023), pp. 5347–5357. ISSN: 0888-5885, 1520-5045. DOI: 10.1021/acs.iecr.2c04561.

[37] James D. McHardy. *An Introduction to the Theory and Use of SESAME Equations of State*. Tech. rep. United States: Los Alamos National Laboratory (LANL), Los Alamos, NM (United States), Dec. 2018. DOI: 10.2172/1487368.

[38] R. Betti, V. Lobatchev, and R. L. McCrory. "Feedout and Rayleigh-Taylor Seeding Induced by Long Wavelength Perturbations in Accelerated Planar Foils". In: *Physical Review Letters* 81.25 (Dec. 1998), pp. 5560–5563. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.81.5560.

[39] V. Lobatchev and R. Betti. "Ablative Stabilization of the Deceleration Phase Rayleigh-Taylor Instability". In: *Phys. Rev. Lett.* 85 (21 Nov. 2000), pp. 4522–4525. DOI: 10.1103/PhysRevLett.85.4522.

[40] Isaac Newton. *Philosophiae Naturalis Principia Mathematica*. London: Jussu Societatis Regiæ ac Typis Josephi Streater., 1687.

[41] R. Betti and O. A. Hurricane. "Inertial-Confinement Fusion with Lasers". In: *Nature Physics* 12.5 (May 2016), pp. 435–448. ISSN: 1745-2473, 1745-2481. DOI: 10.1038/nphys3736.

[42] M. Hohenberger et al. "First Demonstration of Improved Fusion Yield with Increased Compression through Reduced Adiabat in Inertial Confinement Fusion Experiments at the National Ignition Facility". In: *Physical Review Letters* 135.3 (July 2025), p. 035101. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/h2ss-6zpf.

[43] Robert Tipton. *Generalized Lawson Criteria for Inertial Confinement Fusion*. Tech. rep. LLNL–TR-676592, 1234606. Aug. 2015, LLNL–TR-676592, 1234606. DOI: 10.2172/1234606.

[44] H. F. Robey et al. "Precision Shock Tuning on the National Ignition Facility". In: *Physical Review Letters* 108.21 (May 2012), p. 215004. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.108.215004.

[45] K. L. Baker et al. "Adiabat-Shaping in Indirect Drive Inertial Confinement Fusion". In: *Physics of Plasmas* 22.5 (May 2015), p. 052702. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/1.4919694.

[46] R. H. H. Scott et al. "Shock Ignition Laser-Plasma Interactions in Ignition-Scale Plasmas". In: *Physical Review Letters* 127.6 (Aug. 2021), p. 065001. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.127.065001.

[47] R. Betti et al. "Shock Ignition of Thermonuclear Fuel with High Areal Density". In: *Physical Review Letters* 98.15 (Apr. 2007), p. 155001. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.98.155001.

[48] M. Khan et al. "Performance enhancement of direct-drive shock-augmented ignition inertial fusion implosions through shock timing optimization". In: *Phys. Rev. E* 113 (6 June 2026), p. 065210. DOI: 10.1103/zl4y-4blx.

[49] R.H.H. Scott et al. "First Experiments on the Shock-Augmented Ignition Approach to Indirect Drive Laser Inertial Fusion". [Unpublished Work], 2026.

[50] Wilhelm Conrad Röntgen. "On a New Kind of Rays". In: *Nature* 53.1369 (Jan. 1896), pp. 274–276. ISSN: 0028-0836, 1476-4687. DOI: 10.1038/053274b0.

[51] Egwonor Loveth Irede et al. "Medical Imaging: A Critical Review on X-ray Imaging for the Detection of Infection". In: *Biomedical Materials & Devices* (July 2024). ISSN: 2731-4812, 2731-4820. DOI: 10.1007/s44174-024-00212-1.

[52] B. Kozioziemski et al. "X-Ray Imaging Methods for High-Energy Density Physics Applications". In: *Review of Scientific Instruments* 94.4 (Apr. 2023), p. 041102. ISSN: 0034-6748, 1089-7623. DOI: 10.1063/5.0130689.

[53] David Attwood. *Soft X-Rays and Extreme Ultraviolet Radiation: Principles and Applications*. 1st ed. Cambridge University Press, Aug. 1999. ISBN: 978-0-521-65214-8 978-0-521-02997-1 978-1-139-16442-9. DOI: 10.1017/CBO9781139164429.

[54] Jens Als-Nielsen and Des McMorrow. *Elements of Modern X-ray Physics*. 1st ed. Wiley, Mar. 2011. ISBN: 978-0-470-97395-0 978-1-119-99836-5. DOI: 10.1002/9781119998365.

[55] Francesco Barbato et al. "PhaseX: An X-ray Phase-Contrast Imaging Simulation Code for Matter under Extreme Conditions". In: *Optics Express* 30.3 (Jan. 2022), p. 3388. ISSN: 1094-4087. DOI: 10.1364/OE.448479.

[56] Tudor Wyatt Johnston and John M. Dawson. "Correct Values for High-Frequency Power Absorption by Inverse Bremsstrahlung in Plasmas". In: *The Physics of Fluids* 16.5 (May 1973), pp. 722–722. ISSN: 0031-9171. DOI: 10.1063/1.1694419.

[57] Xizeng Wu and Hong Liu. "Clinical Implementation of X-ray Phase-contrast Imaging: Theoretical Foundations and Design Considerations". In: *Medical Physics* 30.8 (Aug. 2003), pp. 2169–2179. ISSN: 0094-2405, 2473-4209. DOI: 10.1118/1.1593836.

[58] C. Stoeckl et al. "Monochromatic Backlighting of Direct-Drive Cryogenic DT Implosions on OMEGA". In: *Physics of Plasmas* 24.5 (May 2017), p. 056304. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/1.4977918.

[59] O. L. Landen et al. "X-Ray Backlighting for the National Ignition Facility (Invited)". In: *Review of Scientific Instruments* 72.1 (Jan. 2001), pp. 627–634. ISSN: 0034-6748, 1089-7623. DOI: 10.1063/1.1315641.

[60] A. Do et al. "X-Ray Imaging of Rayleigh–Taylor Instabilities Using Fresnel Zone Plate at the National Ignition Facility". In: *Review of Scientific Instruments* 92.5 (May 2021), p. 053511. ISSN: 0034-6748, 1089-7623. DOI: 10.1063/5.0043682.

[61] G. N. Hall et al. "The Crystal Backlighter Imager: A Spherically Bent Crystal Imager for Radiography on the National Ignition Facility". In: *Review of Scientific Instruments* 90.1 (Jan. 2019), p. 013702. ISSN: 0034-6748, 1089-7623. DOI: 10.1063/1.5058700.

[62] P. C. Diemoz, A. Bravin, and P. Coan. "Theoretical Comparison of Three X-ray Phase-Contrast Imaging Techniques: Propagation-Based Imaging, Analyzer-Based Imaging and Grating Interferometry". In: *Optics Express* 20.3 (Jan. 2012), p. 2789. ISSN: 1094-4087. DOI: 10.1364/OE.20.002789.

[63] Paolo Russo. *Handbook of X-ray Imaging: Physics and Technology*. Series in Medical Physics and Biomedical Engineering. Boca Raton: CRC press, 2017. ISBN: 978-1-4987-4152-1.

[64] A. Pogany, D. Gao, and S. W. Wilkins. "Contrast and Resolution in Imaging with a Microfocus X-Ray Source". In: *Review of Scientific Instruments* 68.7 (July 1997), pp. 2774–2782. ISSN: 0034-6748, 1089-7623. DOI: 10.1063/1.1148194.

[65] John M. Cowley. *Diffraction Physics*. Elsevier, 1995. ISBN: 978-0-444-82218-5. DOI: 10.1016/B978-0-444-82218-5.X5000-7.

[66] Jeffrey A. Koch et al. "Refraction-Enhanced x-Ray Radiography for Inertial Confinement Fusion and Laser-Produced Plasma Applications". In: *Journal of Applied Physics* 105.11 (June 2009), p. 113112. ISSN: 0021-8979, 1089-7550. DOI: 10.1063/1.3133092.

[67] Adithya Pediredla et al. "Path Tracing Estimators for Refractive Radiative Transfer". In: *ACM Transactions on Graphics* 39.6 (Dec. 2020), pp. 1–15. ISSN: 0730-0301, 1557-7368. DOI: 10.1145/3414685.3417793.

[68] J. A. Koch et al. "Quantitative Analysis of Backlit Shadowgraphy as a Diagnostic of Hydrogen Ice Surface Quality in ICF Capsules". In: *Fusion Technology* 38.1 (July 2000), pp. 123–131. ISSN: 0748-1896. DOI: 10.13182/FST00-A36128.

[69] D. S. Montgomery et al. "Characterization of D-T Cryogenic Layer Formation in a Beryllium Capsule Using X-ray Phase Contrast Imaging". In: *Journal de Physique IV (Proceedings)* 133 (June 2006), pp. 869–873. ISSN: 1155-4339, 1764-7177. DOI: 10.1051/jp4:2006133175.

[70] Jeffrey A. Koch et al. "Refraction-Enhanced Backlit Imaging of Axially Symmetric Inertial Confinement Fusion Plasmas". In: *Applied Optics* 52.15 (May 2013), p. 3538. ISSN: 1559-128X, 2155-3165. DOI: 10.1364/AO.52.003538.

[71] L. Antonelli et al. "X-Ray Phase-Contrast Imaging for Laser-Induced Shock Waves". In: *EPL (Europhysics Letters)* 125.3 (Mar. 2019), p. 35002. ISSN: 1286-4854. DOI: 10.1209/0295-5075/125/35002.

[72] L. Antonelli et al. "X-Ray Phase-Contrast Imaging of Strong Shocks on OMEGA EP". In: *Review of Scientific Instruments* 95.11 (Nov. 2024), p. 113504. ISSN: 0034-6748, 1089-7623. DOI: 10.1063/5.0168059.

[73] J. Workman et al. "Phase-Contrast Imaging Using Ultrafast x-Rays in Laser-Shocked Materials". In: *Review of Scientific Instruments* 81.10 (Oct. 2010), 10E520. ISSN: 0034-6748, 1089-7623. DOI: 10.1063/1.3485109.

[74] E.L. Dewald et al. "Direct Observation of Density Gradients in ICF Capsule Implosions via Streaked Refraction Enhanced Radiography (RER)". In: *High Energy Density Physics* 36 (Aug. 2020), p. 100795. ISSN: 15741818. DOI: 10.1016/j.hedp.2020.100795.

[75] C. H. Allen et al. "Toward an Integrated Platform for Characterizing Laser-Driven, Isochorically Heated Plasmas with 1 µm Spatial Resolution". In: *Applied Optics* 61.8 (Mar. 2022), p. 1987. ISSN: 1559-128X, 2155-3165. DOI: 10.1364/AO.446182.

[76] F Barbato et al. "Feasibility Study of an XPCI Diagnostic to Observe the Evolution of Micro-Voids in an ICF Target". In: *Plasma Physics and Controlled Fusion* 66.2 (Feb. 2024), p. 025017. ISSN: 0741-3335, 1361-6587. DOI: 10.1088/1361-6587/ad1a42.

[77] S. H. Glenzer et al. "Symmetric Inertial Confinement Fusion Implosions at Ultra-High Laser Energies". In: *Science* 327.5970 (Mar. 2010), pp. 1228–1231. ISSN: 0036-8075, 1095-9203. DOI: 10.1126/science.1185634.

[78] A. B. Zylstra et al. "Hot-Spot Mix in Large-Scale HDC Implosions at NIF". In: *Physics of Plasmas* 27.9 (Sept. 2020), p. 092709. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/5.0003779.

[79] L. A. Pickworth et al. "Development of New Platforms for Hydrodynamic Instability and Asymmetry Measurements in Deceleration Phase of Indirectly Driven Implosions on NIF". In: *Physics of Plasmas* 25.8 (Aug. 2018), p. 082705. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/1.5039744.

[80] L. A. Welser et al. "Processing of Multi-Monochromatic x-Ray Images from Indirect Drive Implosions at OMEGA". In: *Review of Scientific Instruments* 74.3 (Mar. 2003), pp. 1951–1953. ISSN: 0034-6748, 1089-7623. DOI: 10.1063/1.1535278.

[81] Dylan T. Cliche and Roberto C. Mancini. "Impact of 3D Effects on the Characteristics of a Multi-Monochromatic x-Ray Imager". In: *Applied Optics* 58.17 (June 2019), p. 4753. ISSN: 1559-128X, 2155-3165. DOI: 10.1364/AO.58.004753.

[82] J. A. Koch et al. "Multispectral X-Ray Imaging with a Pinhole Array and a Flat Bragg Mirror". In: *Review of Scientific Instruments* 76.7 (July 2005), p. 073708. ISSN: 0034-6748, 1089-7623. DOI: 10.1063/1.1940614.

[83] E. L. Dewald et al. "Dante Soft X-Ray Power Diagnostic for National Ignition Facility". In: *Review of Scientific Instruments* 75.10 (Oct. 2004), pp. 3759–3761. ISSN: 0034-6748, 1089-7623. DOI: 10.1063/1.1788872.

[84] Cristina Serbanescu et al. "K-Alpha x-Ray Source Using High Energy and High Repetition Rate Laser System for Phase Contrast Imaging". In: *SPIE Optical Engineering + Applications*. Ed. by James Dunn and Gregory J. Tallents. San Diego, CA, Aug. 2009, p. 745115. DOI: 10.1117/12.826646.

[85] S. Kneip et al. "X-Ray Phase Contrast Imaging of Biological Specimens with Femtosecond Pulses of Betatron Radiation from a Compact Laser Plasma Wakefield Accelerator". In: *Applied Physics Letters* 99.9 (Aug. 2011), p. 093701. ISSN: 0003-6951, 1077-3118. DOI: 10.1063/1.3627216.

[86] R. J. Shalloo et al. "Automation and Control of Laser Wakefield Accelerators Using Bayesian Optimization". In: *Nature Communications* 11.1 (Dec. 2020), p. 6355. ISSN: 2041-1723. DOI: 10.1038/s41467-020-20245-6.

[87] J. D. Kilkenny et al. "National Diagnostic Working Group (NDWG) for Inertial Confinement Fusion (ICF)/High-Energy Density (HED) Science: The Whole Exceeds the Sum of Its Parts". In: *Review of Scientific Instruments* 94.8 (Aug. 2023), p. 081101. ISSN: 0034-6748, 1089-7623. DOI: 10.1063/5.0128650.

[88] Daniel Dolan. *Foundations of VISAR Analysis*. Tech. rep. SAND2006-1950, 886901. June 2006, SAND2006–1950, 886901. DOI: 10.2172/886901.

[89] H. F. Robey et al. "Shock Timing Experiments on the National Ignition Facility: Initial Results and Comparison with Simulation". In: *Physics of Plasmas* 19.4 (Apr. 2012), p. 042706. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/1.3694122.

[90] Marco Ament, Christoph Bergmann, and Daniel Weiskopf. "Refractive Radiative Transfer Equation". In: *ACM Transactions on Graphics* 33.2 (Mar. 2014), pp. 1–22. ISSN: 0730-0301, 1557-7368. DOI: 10.1145/2557605.

[91] A.M. Smith, M. Goldberg, and E.S.K. Liu. "Numerical Ray Tracing in Media Involving Continuous and Discrete Refractive Boundaries". In: *Ultrasonic Imaging* 2.4 (Oct. 1980), pp. 291–301. ISSN: 0161-7346, 1096-0910. DOI: 10.1177/016173468000200401.

[92] Luca Brombal et al. "X-Ray Differential Phase-Contrast Imaging Simulations with Geant4". In: *Journal of Physics D: Applied Physics* (Oct. 2021). ISSN: 0022-3727, 1361-6463. DOI: 10.1088/1361-6463/ac2e8a.

[93] J. Allison et al. "Recent Developments in Geant4". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 835 (Nov. 2016), pp. 186–225. ISSN: 01689002. DOI: 10.1016/j.nima.2016.06.125.

[94] Manuel Sanchez Del Rio et al. "*SHADOW3* : A New Version of the Synchrotron X-ray Optics Modelling Package". In: *Journal of Synchrotron Radiation* 18.5 (Sept. 2011), pp. 708–716. ISSN: 0909-0495, 1600-5775. DOI: 10.1107/S0909049511026306.

[95] Luca Rebuffi and Manuel Sánchez Del Río. "*ShadowOui* : A New Visual Environment for X-ray Optics and Synchrotron Beamline Simulations". In: *Journal of Synchrotron Radiation* 23.6 (Nov. 2016), pp. 1357–1367. ISSN: 1600-5775. DOI: 10.1107/S1600577516013837.

[96] T.T. Böhlen et al. "The FLUKA Code: Developments and Challenges for High Energy and Medical Applications". In: *Nuclear Data Sheets* 120 (June 2014), pp. 211–214. ISSN: 00903752. DOI: 10.1016/j.nds.2014.07.049.

[97] Silvia Cipiccia et al. "Inclusion of Coherence in Monte Carlo Models for Simulation of X-Ray Phase Contrast Imaging". In: *Optics Express* 22.19 (Sept. 2014), p. 23480. ISSN: 1094-4087. DOI: 10.1364/OE.22.023480.

[98] B. Fryxell et al. "FLASH: An Adaptive Mesh Hydrodynamics Code for Modeling Astrophysical Thermonuclear Flashes". In: *The Astrophysical Journal Supplement Series* 131.1 (Nov. 2000), pp. 273–334. ISSN: 0067-0049, 1538-4365. DOI: 10.1086/317361.

[99] Thomas Goffrey. "A Cylindrical Magnetohydrodynamic Arbitrary Lagrangian Eulerian Code". PhD thesis. University of Warwick, Jan. 2014. URL: https://wrap.warwick.ac.uk/id/eprint/63693/1/WRAP_THESIS_Goffrey_2014.pdf.

[100] Amaresh Sahu et al. "Arbitrary Lagrangian–Eulerian finite element method for curved and deforming surfaces. I. General theory and application to fluid interfaces". eng. In: *arXiv.org* (2020). ISSN: 2331-8422.

[101] Lyman Spitzer and Richard Härm. "Transport Phenomena in a Completely Ionized Gas". In: *Physical Review* 89.5 (Mar. 1953), pp. 977–981. ISSN: 0031-899X. DOI: 10.1103/PhysRev.89.977.

[102] A. R. Bell. "Non-Spitzer Heat Flow in a Steadily Ablating Laser-Produced Plasma". In: *The Physics of Fluids* 28.6 (June 1985), pp. 2007–2014. ISSN: 0031-9171. DOI: 10.1063/1.865378.

[103] A. R. Bell, R. G. Evans, and D. J. Nicholas. "Elecron Energy Transport in Steep Temperature Gradients in Laser-Produced Plasmas". In: *Physical Review Letters* 46.4 (Jan. 1981), pp. 243–246. ISSN: 0031-9007. DOI: 10.1103/PhysRevLett.46.243.

[104] T D Arber et al. "Contemporary Particle-in-Cell Approach to Laser-Plasma Modelling". In: *Plasma Physics and Controlled Fusion* 57.11 (Nov. 2015), p. 113001. ISSN: 0741-3335, 1361-6587. DOI: 10.1088/0741-3335/57/11/113001.

[105] A R Bell and M Sherlock. "The fastVFP Code for Solution of the Vlasov–Fokker–Planck Equation". In: *Plasma Physics and Controlled Fusion* 66.3 (Mar. 2024), p. 035014. ISSN: 0741-3335, 1361-6587. DOI: 10.1088/1361-6587/ad2278.

[106] Jon T. Larsen and Stephen M. Lane. "HYADES—A Plasma Hydrodynamics Code for Dense Plasma Studies". In: *Journal of Quantitative Spectroscopy and Radiative Transfer* 51.1-2 (Jan. 1994), pp. 179–186. ISSN: 00224073. DOI: 10.1016/0022-4073(94)90078-7.

[107] Jon Larsen. *Hyades User's Guide*. Cascade Applied Sciences, Nov. 2013.

[108] Jon Larsen. *Hyades Physics Manual, Chapter E: Thermal Energy Transport*. Cascade Applied Sciences.

[109] Robbie Scott. *Private Communication*. 2026.

[110] S. Atzeni and A. Caruso. "A diffusive model for alpha-particle energy transport in a laser plasma". eng. In: *Nuovo cimento della Società italiana di fisica. Sezione B* 64.2 (1981), pp. 383–395. ISSN: 1826-9877.

[111] R. M. More et al. "A New Quotidian Equation of State (QEOS) for Hot Dense Matter". In: *The Physics of Fluids* 31.10 (Oct. 1988), pp. 3059–3078. ISSN: 0031-9171. DOI: 10.1063/1.866963.

[112] John David Jackson. *Classical Electrodynamics*. 3. ed., [Nachdr.] Hoboken, NY: Wiley, 2009. ISBN: 978-0-471-30932-1.

[113] S. Fourmaux et al. "Laser-Based Synchrotron X-ray Radiation Experimental Scaling". In: *Optics Express* 28.3 (Feb. 2020), p. 3147. ISSN: 1094-4087. DOI: 10.1364/OE.383818.

[114] L. Brombal et al. "A Geant4 Tool for Edge-Illumination X-ray Phase-Contrast Imaging". In: *Journal of Instrumentation* 17.01 (Jan. 2022), p. C01043. ISSN: 1748-0221. DOI: 10.1088/1748-0221/17/01/C01043.

[115] Zhentian Wang et al. "Implement X-ray Refraction Effect in Geant4 for Phase Contrast Imaging". In: *2009 IEEE Nuclear Science Symposium Conference Record (NSS/MIC)*. Orlando, FL: IEEE, Oct. 2009, pp. 2395–2398. ISBN: 978-1-4244-3961-4. DOI: 10.1109/NSSMIC.2009.5402180.

[116] B.L. Henke, E.M. Gullikson, and J.C. Davis. "X-Ray Interactions: Photoabsorption, Scattering, Transmission, and Reflection at E = 50-30,000 eV, Z = 1-92". In: *Atomic Data and Nuclear Data Tables* 54.2 (July 1993), pp. 181–342. ISSN: 0092640X. DOI: 10.1006/adnd.1993.1013.

[117] C. M. Poole et al. "A CAD Interface for GEANT4". In: *Australasian Physical & Engineering Sciences in Medicine* 35.3 (Sept. 2012), pp. 329–334. ISSN: 0158-9938, 1879-5447. DOI: 10.1007/s13246-012-0159-8.

[118] Pauli Virtanen et al. "SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python". In: *Nature Methods* 17 (2020), pp. 261–272. DOI: 10.1038/s41592-019-0686-2.

[119] Will Schroeder, Ken Martin, and Bill Lorensen. *The Visualization Toolkit: An Object-Oriented Approach to 3D Graphics*. 4th ed. Kitware, 2006. ISBN: 978-1-930934-19-1.

[120] Alex Do. *Private Communication*. May 2024.

[121] S. Fourmaux et al. "X-Ray Phase Contrast Imaging of Spherical Capsules". In: *Optics Express* 28.9 (Apr. 2020), p. 13978. ISSN: 1094-4087. DOI: 10.1364/OE.386618.

[122] Peng Zhang, C P Ridgers, and A G R Thomas. "The Effect of Nonlinear Quantum Electrodynamics on Relativistic Transparency and Laser Absorption in Ultra-

Relativistic Plasmas". In: *New Journal of Physics* 17.4 (Apr. 2015), p. 043051. ISSN: 1367-2630. DOI: 10.1088/1367-2630/17/4/043051.

[123] B. Kettle et al. "Single-Shot Multi-keV X-Ray Absorption Spectroscopy Using an Ultrashort Laser-Wakefield Accelerator Source". In: *Physical Review Letters* 123.25 (Dec. 2019), p. 254801. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.123.254801.

[124] H. F. Robey et al. "Performance of Indirectly Driven Capsule Implosions on the National Ignition Facility Using Adiabat-Shaping". In: *Physics of Plasmas* 23.5 (May 2016), p. 056303. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/1.4944821.

[125] A. L. Kritcher et al. "Design of an Inertial Fusion Experiment Exceeding the Lawson Criterion for Ignition". In: *Physical Review E* 106.2 (Aug. 2022), p. 025201. ISSN: 2470-0045, 2470-0053. DOI: 10.1103/PhysRevE.106.025201.

[126] C. D. Harris et al. "Spectral sensor error analysis for measuring x-ray radiation drive using the DANTE diagnostic toward inertial confinement fusion experiments". In: *Review of Scientific Instruments* 92.3 (Mar. 2021), p. 033505. ISSN: 0034-6748. DOI: 10.1063/5.0035584.

[127] scikit-optimize contributors. *Bayesian Optimisation with Skopt*. URL: https://scikit-optimize.readthedocs.io/en/latest/auto_examples/bayesian-optimization.html.

[128] Peter I. Frazier. *A Tutorial on Bayesian Optimization*. 2018. DOI: 10.48550/ARXIV.1807.02811.

[129] Mikhail Basko. "An Improved Version of the View Factor Method for Simulating Inertial Confinement Fusion Hohlraums". In: *Physics of Plasmas* 3.11 (Nov. 1996), pp. 4148–4155. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/1.871565.

[130] D. A. Callahan et al. "A Simple Model to Scope out Parameter Space for Indirect Drive Designs on NIF". In: *Physics of Plasmas* 27.7 (July 2020), p. 072704. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/5.0006217.

[131] A Gowda et al. *High Fidelity Pulse Shaping for the National Ignition Facility*. Tech. rep. LLNL-PROC-855426; 1083859. United States: Lawrence Livermore National Laboratory, Sept. 2023. URL: https://www.osti.gov/biblio/2202546.

[132] F. Timischl. "The contrast-to-noise ratio for image quality evaluation in scanning electron microscopy". In: *Scanning* 37.1 (2015), pp. 54–62. DOI: 10.1002/sca.21179.

[133] C. Trosseille, S. R. Nagel, and T. J. Hilsabeck. "Electron pulse-dilation diagnostic instruments". eng. In: *Review of scientific instruments* 94.2 (2023), pp. 021102–. ISSN: 0034-6748.

[134] Paul D. Mason et al. "High Energy Diode-Pumped Solid-State Laser Development at the Central Laser Facility". In: *SPIE Photonics Europe*. Ed. by Jacob I. Mackenzie et al. Brussels, Belgium, Apr. 2016, p. 989309. DOI: 10.1117/12.2230317.

Appendix A

Gain Requirements

This appendix is based on Lindl [22], with a modern laser efficiency.

To design an IFE reactor, there are a number of additional considerations, which constrain implosion designs. One constraint is that fundamentally a reactor must produce more electricity ($E_{electricity}$) than it uses to drive an implosion (E_{driver})) and to run auxiliary processes (E_{aux}, eg powering the control room, tritium recovery, etc). So that net electricity onto the grid (E_{grid}) is produced. The electricity onto the grid is given by:

$$E_{grid} = E_{electricity} - E_{driver} - E_{aux}$$

The electricity produced (for a DT implosion) will come from a turbine, which turns the energy released from the implosion (E_{fusion}) into electricity, and has some limited efficiency ($\eta_{turbine}$). Which gives:

$$E_{grid} = E_{fusion}\eta_{turbine} - E_{driver} - E_{aux}$$

The fusion energy produced is based on the gain (G) of the implosion and the laser energy (E_{laser}) in, meanwhile the required driver energy can be found from the laser energy and laser efficieny (η_{laser}).

$$E_{grid} = GE_{laser}\eta_{turbine} - \frac{E_{laser}}{\eta_{laser}} - E_{aux}$$

In addition, a reactor design must also be economical, if an alternative system (such as solar panels + battery storage) is cheaper it is likely a more efficient use of resources. To keep the cost down, the recirculating energy ($E_{driver} + E_{aux}$ should be less than 25% of the

power to the grid.

$$E_{grid} = GE_{laser}\eta_{turbine} - \frac{E_{laser}}{\eta_{laser}} - E_{aux} > 4(\frac{E_{laser}}{\eta_{laser}} + E_{aux})$$

For a typical reactor design, the auxiliary processes are a few percent of the gross electricity generated ($E_{electricity}$), for this calculation 5% will be used. Turbines typically are 35-45% efficient [23, 27, 22], so 35% will be used to be conservative. A diode pumped solid state laser system like DiPOLE [134] can have more than 10% wall plug efficiency. This gives a gain requirement [22]:

$$GE_{laser}\eta_{turbine} - \frac{E_{laser}}{\eta_{laser}} - 2E_{aux} > 4(\frac{E_{laser}}{\eta_{laser}})$$

$$(1 - 0.1)GE_{laser}0.35 - 10E_{laser} > 40E_{laser}$$

$$G > 50/0.315 = 159$$

The highest gain shot achieved by NIF (as of April 2025) had a gain of just over 4, so this gain requirement emphasizes the reality that there is much more work to be done.

Appendix B

Tuning Reflected Shocks

This appendix will describe the physical mechanisms by which modifying the radiation temperature profile of the drive that ablates a capsule can alter the strength and velocity of pressure waves that form when rarefactions reflect off the ablation front.

The simulations from sections 5.5 and 5.6 will be compared at a number of key timesteps to illustrated the effects of a modified radiation drive. While this example will be framed as illustrating the effect of an increased radiation drive relative to the baseline case, of course the effects of reducing the radiation drive can be seen by simply reframing and thinking of the "modified" drive as the baseline and the "baseline" as the modification.

The first point in time shown is at 4.38ns into the implosion, before any deviation has occurred. Figure B.1 has three panels, the top panel shows the lagrangian frame shockplot for the baseline NIF N210808 shot, and the second panel shows the lagrangian frame shockplot for the modified case. Both panels have a marker line showing the point in time that is being examined in the lowest panel. The lowest panel shows the density, electron temperature (Te) and pressures for both the baseline and modified simulations. Since there has been no deviation yet the lines overlap perfectly.

Figure B.2 shows the second point in time (at 4.58ns), here the modified radiation temperature profile has increased relative to the baseline slightly, the result is that the pressure due to ablation is higher in the modified case

Figure B.3 shows the third point (at 4.62ns), which shows what happens as the rarefaction from the second shock approaches the ablation front. Since the ablation pressure is determined by the radiation drive, as the rarefaction wave approaches the ablation front, the relative pressure difference between the pressure in the rarefied material and the ablation

pressure is greater in the modified case. This means that when this rarefaction wave reflects and begins to work its way back in towards the core as a pressure wave, the gradient of the pressure wave is stronger in the modified case. So it is a stronger (and therefore faster) pressure wave in the modified case.

Figure B.4 shows the fourth point (at 5.00ns), which shows the point in time when the pressure wave is about to cross the ablator-ice interface in the modified case. Here it is clear that the pressure gradient is stronger for the modified case, and that the pressure wave has travelled through the ablator faster. In the fifth point (at 5.10ns, shown in figure B.5), for the modified case the pressure wave has made it through the ice-ablator interface and is moving rapidly towards the ice-gas interface, while the pressure wave in the baseline case is still in the ablator.

These plots show that the increased ablation pressure caused by an increased radiation temperature results in stronger pressure gradients when a rarefaction gets to the ablation front, which results in a stronger (and faster) pressure wave when the rarefaction reflects. If the pressure wave steepens up into a shock, the shock will be stronger and faster as a result.

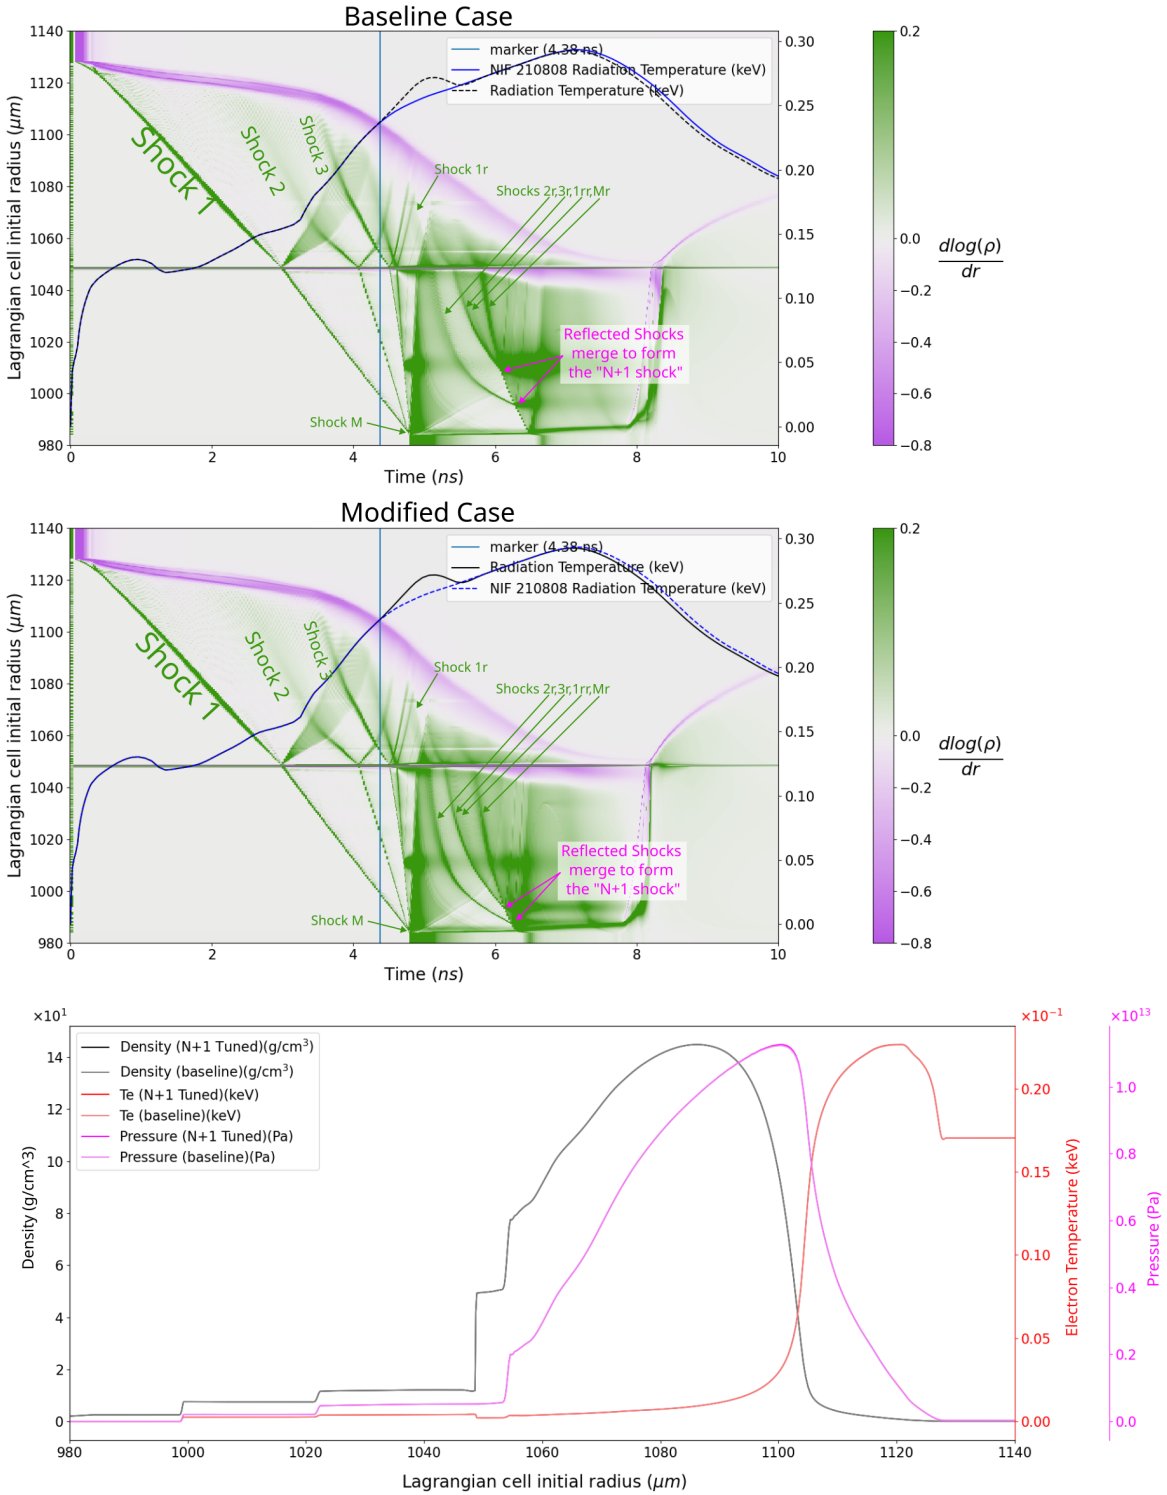


Figure B.1: Two lagrangian shockplots, showing a baseline (first panel) and modified (second panel) implosion design, with a lineout of density, electron temperature and pressure taken at the point marked in the plots. The lineouts are shown in the third panel, and taken at 4.38ns.

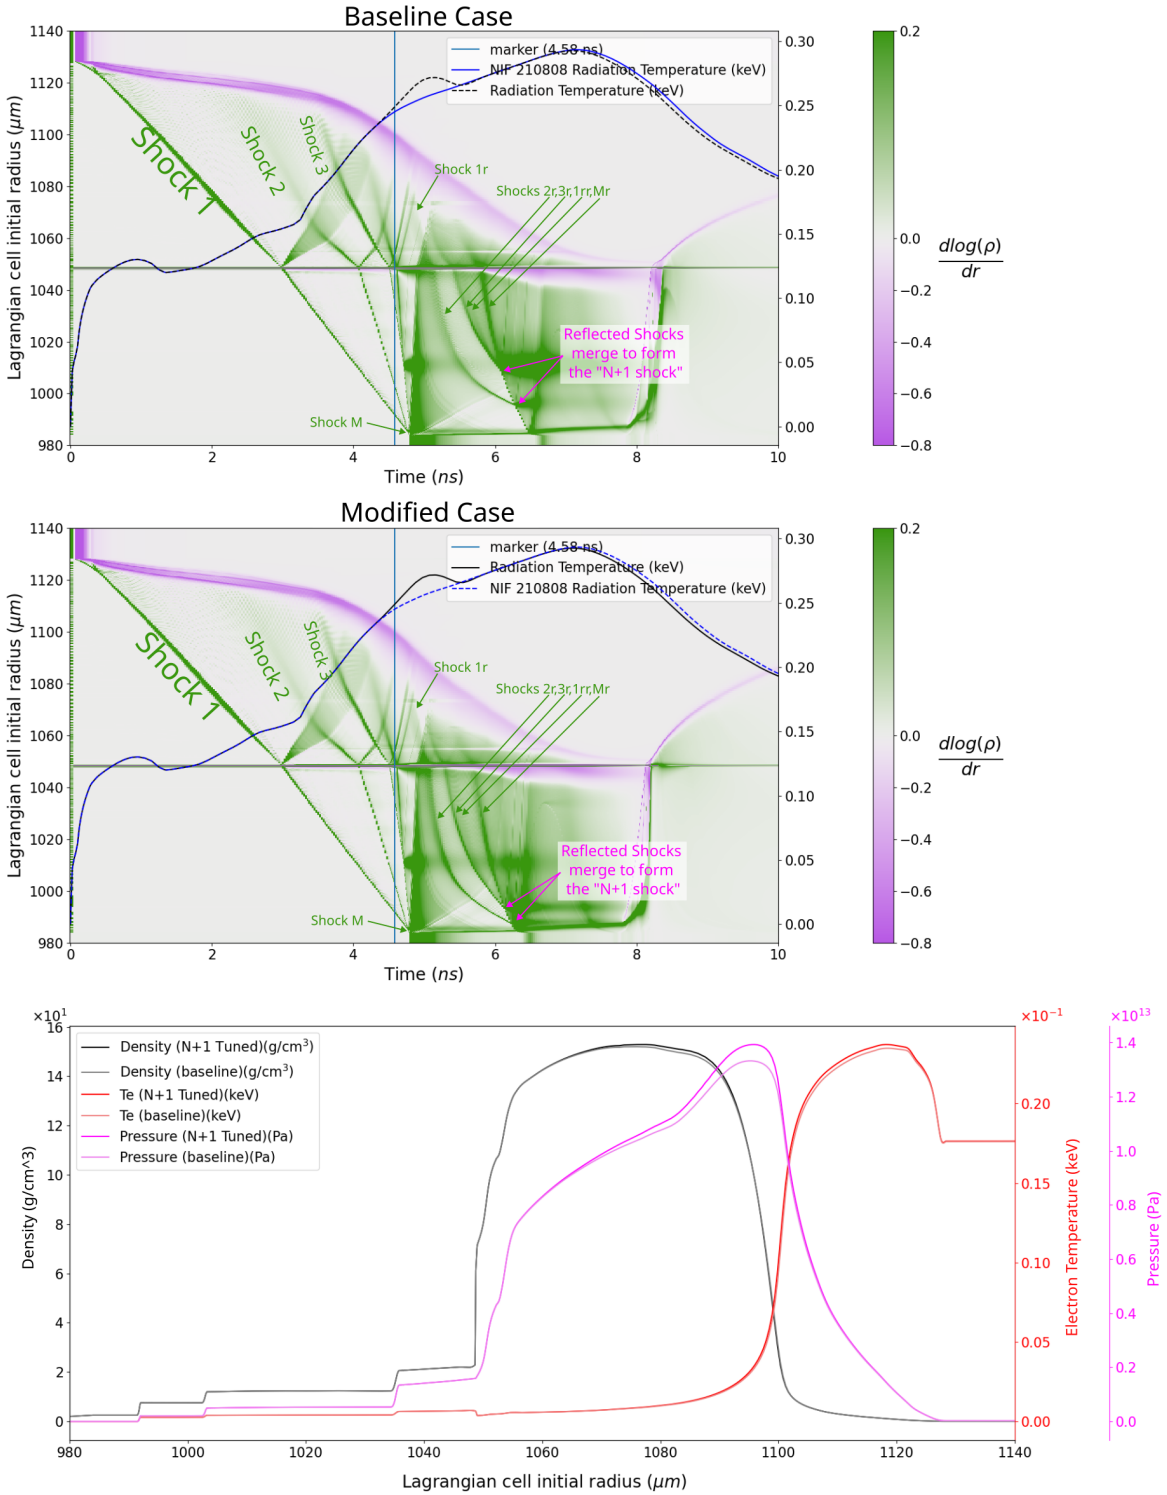


Figure B.2: Two lagrangian shockplots, showing a baseline (first panel) and modified (second panel) implosion design, with a lineout of density, electron temperature and pressure taken at the point marked in the plots. The lineouts are shown in the third panel, and taken at 4.58ns.

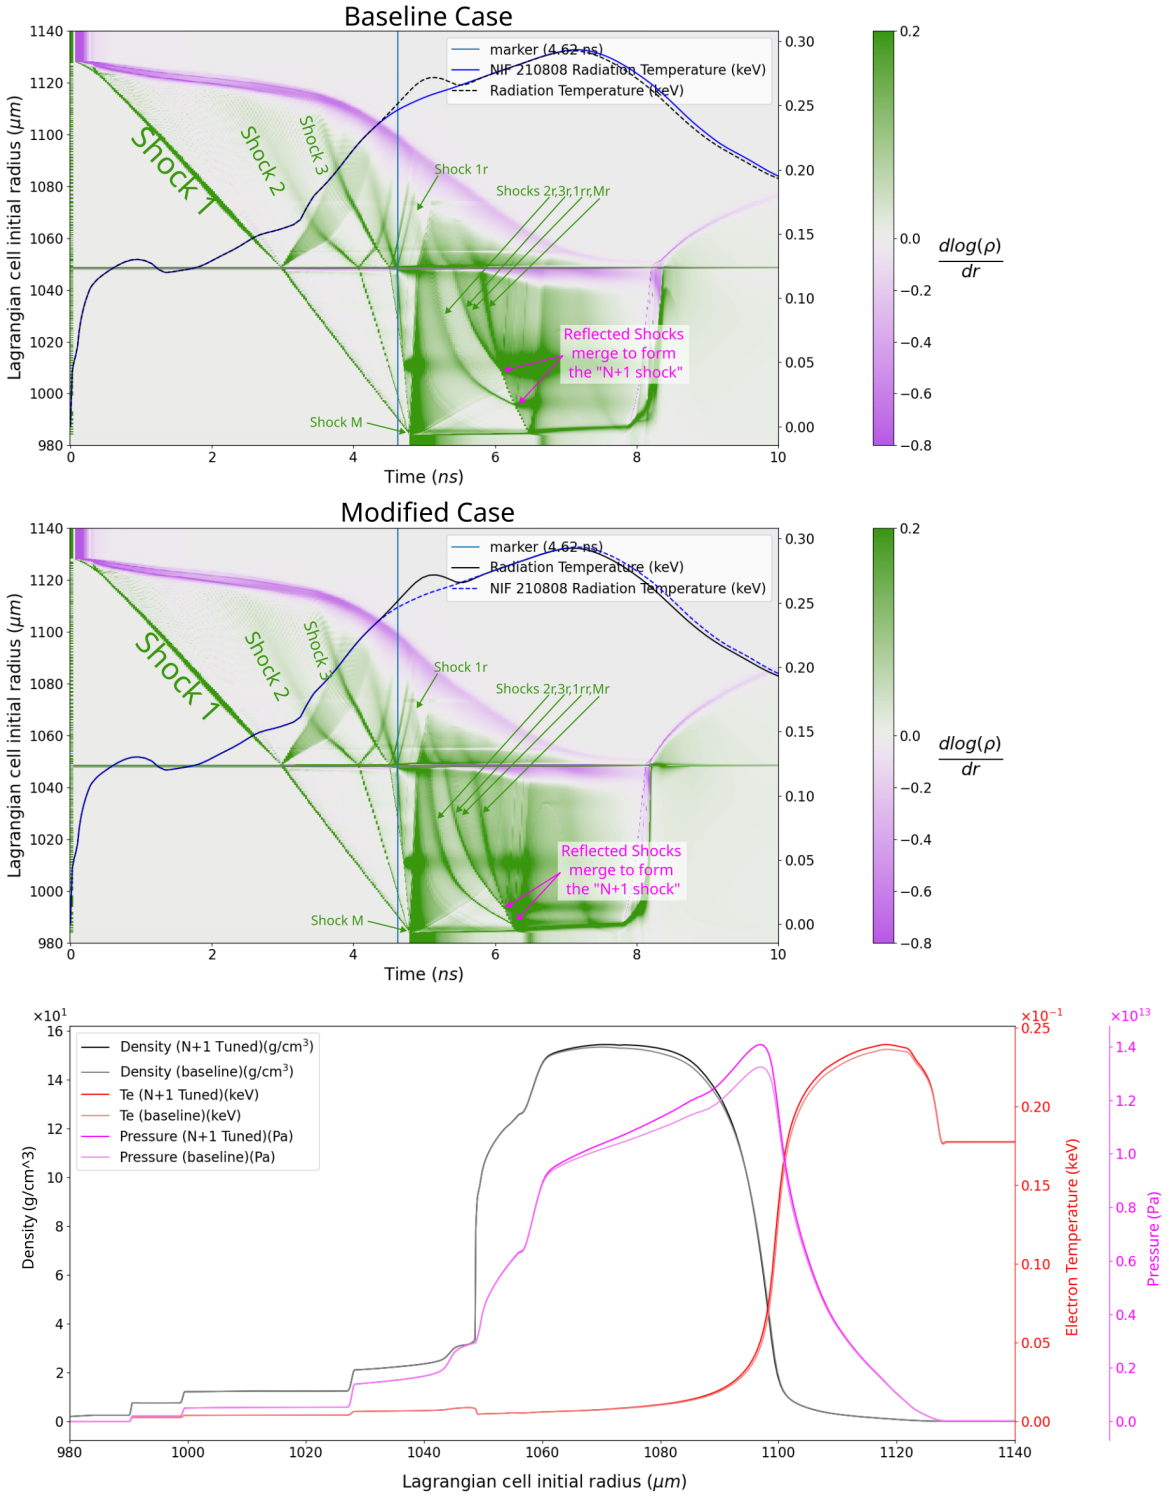


Figure B.3: Two lagrangian shockplots, showing a baseline (first panel) and modified (second panel) implosion design, with a lineout of density, electron temperature and pressure taken at the point marked in the plots. The lineouts are shown in the third panel, and taken at 4.62ns.

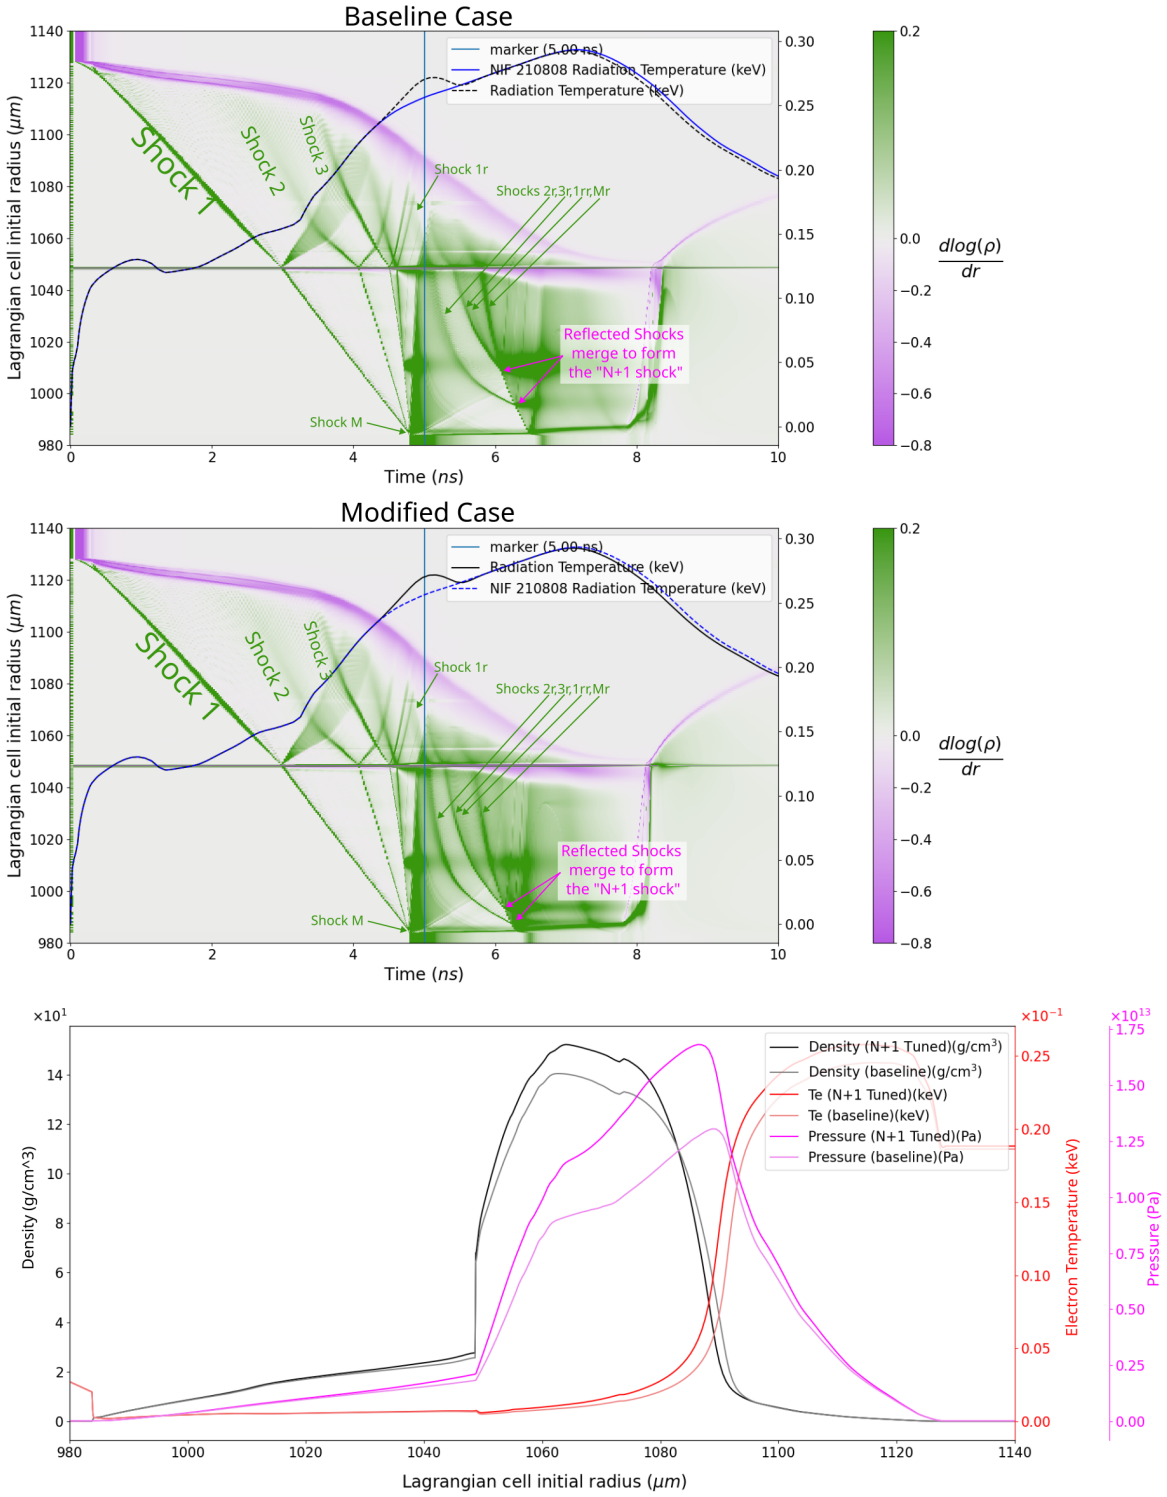


Figure B.4: Two lagrangian shockplots, showing a baseline (first panel) and modified (second panel) implosion design, with a lineout of density, electron temperature and pressure taken at the point marked in the plots. The lineouts are shown in the third panel, and taken at 5.00ns.

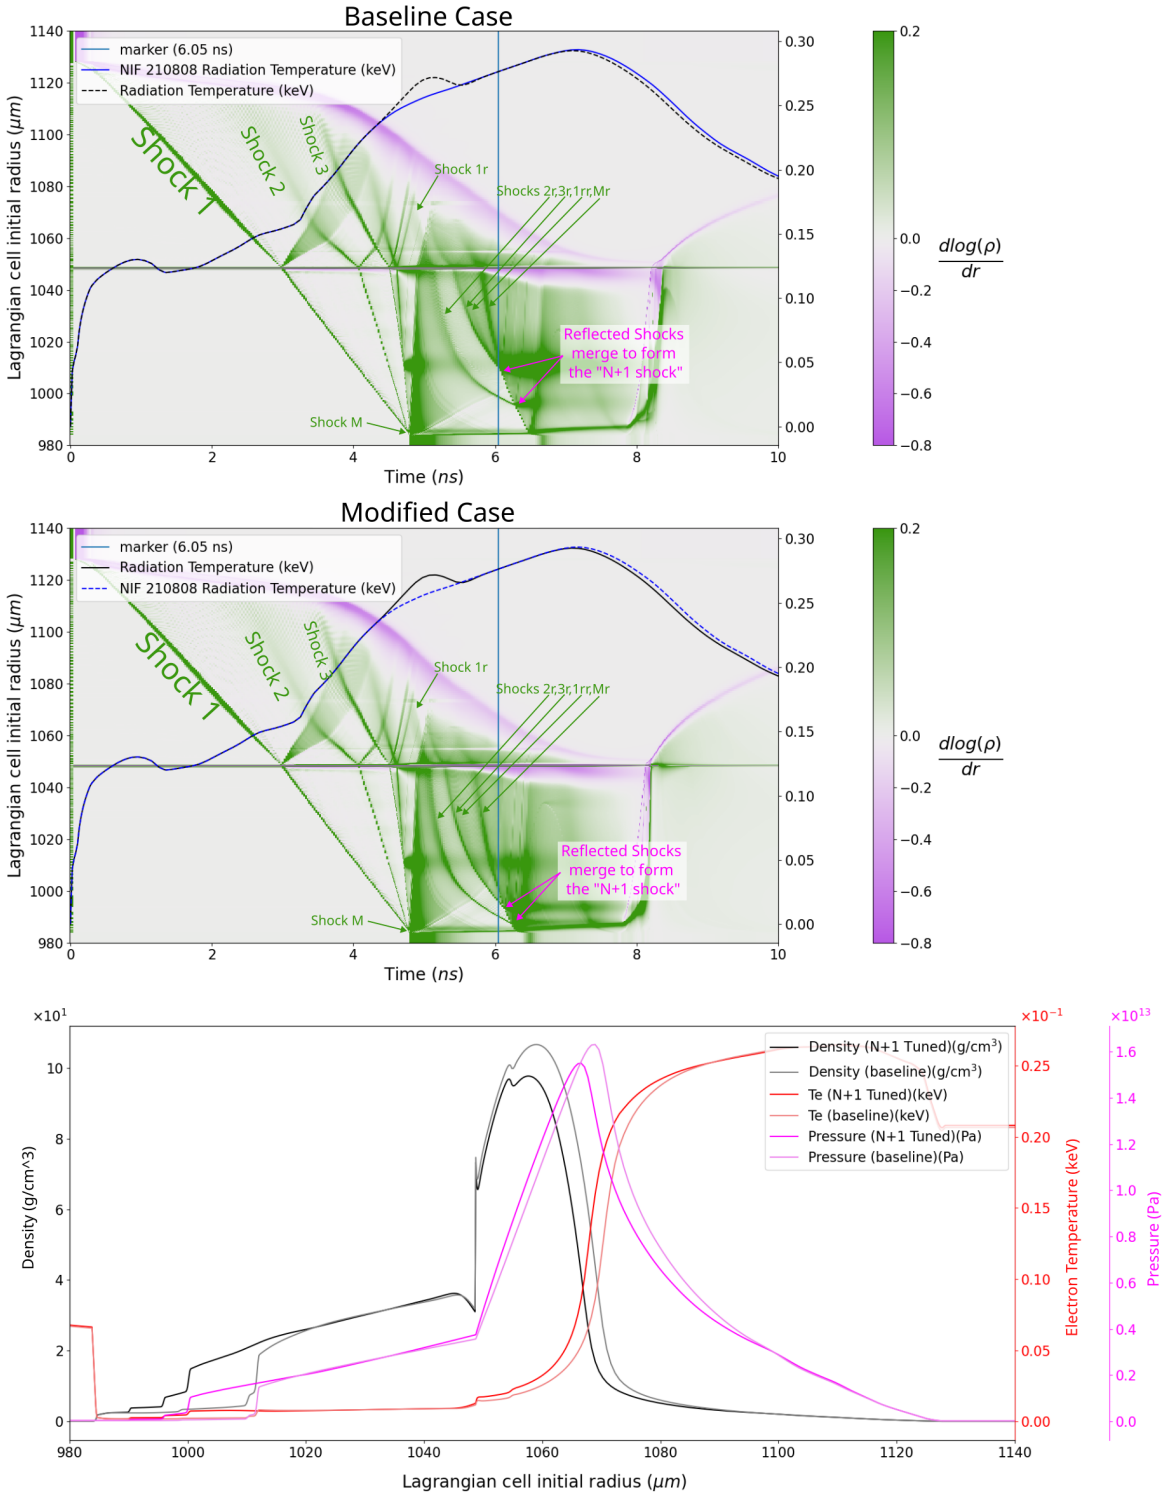


Figure B.5: Two lagrangian shockplots, showing a baseline (first panel) and modified (second panel) implosion design, with a lineout of density, electron temperature and pressure taken at the point marked in the plots. The lineouts are shown in the third panel, and taken at 5.10ns.

Appendix C

Convergence Testing GREENER

The requirements for appropriately resolved spatial data required to generate x-ray phase contrast images with greener vary significantly between geometries, imaging parameters and simulation setups. To have confidence in a result, the user must perform some convergence tests relevant for their particular simulations. This appendix shows the details of convergence tests performed to check the accuracy of the results in chapter 6, however, in an attempt to be more general: these results should be applicable for all similar spherical capsules.

Density profiles were produced which radially ramp up in density, then back down again; two examples with different gradients are shown in the left panel of figure C.1. Similar ramps were generated for a range of different ramp gradients which were chosen to be of relevance to the simulations in chapter 6. The chosen density ramp gradients were, with units of $\frac{g/cm^3}{\mu m}$: 0.001, 0.02, 0.05, 0.1, 0.15, 0.30 and 100, where $100\frac{g/cm^3}{\mu m}$ is used to represent a shock. For each of these density ramps, sampled density distributions were produced with different spatial resolutions: 80μm, 40μm 20μm 10μm 5μm 2.5μm 1.25μm 0.625μm 0.3125μm 0.1μm. Four example density profiles are shown in the right panel of figure C.1, with two different sampling resolutions for each of the example ramps. Maximum densities were capped at 8g/cm^3 to prevent too much absorption from making it difficult to resolve the inner edge. In the case of $100\frac{g/cm^3}{\mu m}$ at all step sizes, and 0.3 $\frac{g/cm^3}{\mu m}$ with step sizes greater than 20μm: the gradient was steep enough that with the limited maximum density the sampled density profiles were step functions, effectively ideal shocks. In this case spatial resolution directly correlated to the location of the step in density, so directly translates to the accuracy of the shocks location.

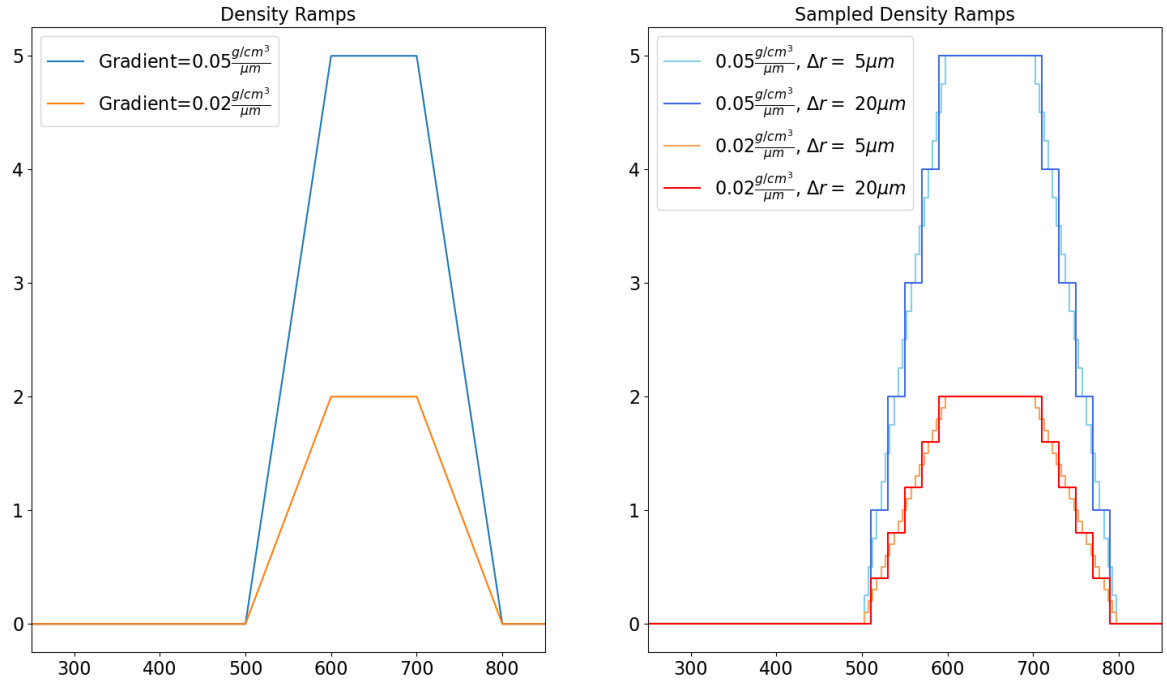


Figure C.1: Density ramps of two different gradients (left panel), along with the same two density ramps each sampled at different spatial resolutions (right panel).

XPCI simulations were performed for each gradient at each spatial resolution, these simulations used the same parameters as were used in chapter 6 However, for these simulations the seed used for Geant4's Monte-Carlo particle tracking was kept fixed, meaning that each image will have identical shot noise. This makes the simulations fully deterministic, so that two images of the same density distribution will be identical. This allows convergence with respect to the spatial resolution of the density profiles to be tested, by comparing different spatial resolutions of the same density profile until increasing the spatial resolution does not discernibly change the resulting XPCI lineout.

A plot showing XPCI lineouts and density profiles for a gradient of 0.05 $\frac{g/cm^3}{\mu m}$ is shown in figure C.2, with a few different spatial resolutions. Qualitatively it is clear that for a low resolution (40μm), the density ramp is significantly under sampled, but for higher resolutions (1.25μm), the difference from the highest simulated resolution (0.1μm) is not discernible. Quantitatively, the average percentage difference from the ideal case can be calculated to determine which spatial resolutions are high enough to converge:

$$NME\% = 100 \times \left\langle \frac{|I_s(r) - I_{ideal}(r)|}{I_{ideal}(r)} \right\rangle$$

where $NME\%$ is the mean percentage deviation (or Normalised Mean Error) form the 0.1μm case, $I_s(r)$ is the Intensity (in photons/pixel) for the XPCI lineout of the density profile with spatial resolution s, at radius r, and $I_{ideal}(r)$ is the Intensity at radius r for the ideal case.

In the lineouts produced for this convergence test, and the ones in chapter 6, with the Monte-Carlo element turned back on: the shot noise from finite numbers of particles being used in the simulations is expected to produce a typical error of at least 1%. This means that any deviations less than this will not be detectable when comparing lineouts, and therefore a mean deviation from the ideal case of less than 1% can be considered to have converged within the expected errors. For this convergence test, the 0.1μm case was treated as the ideal case, all gradients converged within 1% before this spatial resolution was required, and simulating with higher resolution proved computationally expensive.

The regions used to calculate the average deviation are labelled at the bottom of the figure, these were chosen to exclude all the areas where the distribution is not altered by the step size (eg before and after each ramp).

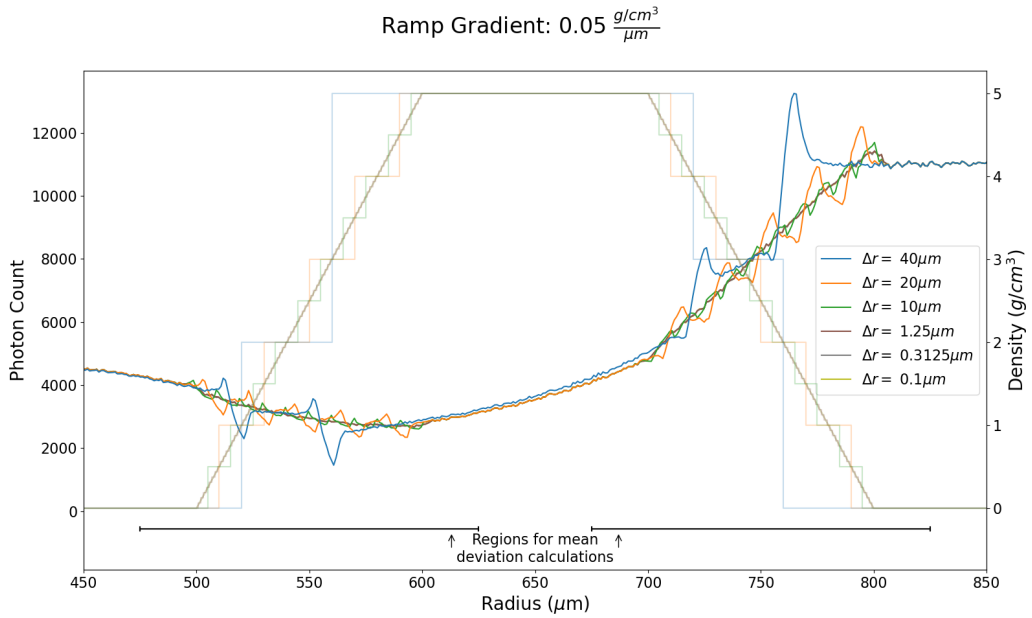


Figure C.2: XPCI lineouts for a density ramp of $0.05\frac{g/cm^3}{\mu m}$, with a range of spatial resolutions. Faint lines with the same colour as the lineout they correspond to show the density profiles used. The regions used for mean deviation calculations are indicated at the bottom of the plot

Figure C.3 shows the average percentage deviation from the simulation with 0.1μm spatial

resolution for a gradient of 0.05 $\frac{g/cm^3}{\mu m}$, here the error is less than 1% even for a 5μm step size. The error decreases further until 0.6μm, beyond which there is minimal improvement, this would be a stricter definition of true convergence, but is not necessary for this work.

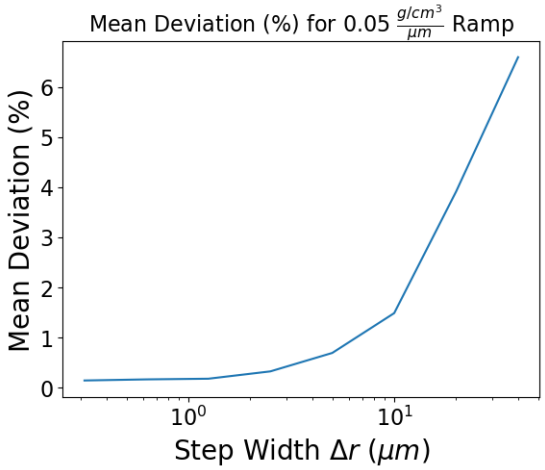


Figure C.3: Mean percentage deviation at different spatial resolutions (step widths) for a density ramp of 0.05 $\frac{g/cm^3}{\mu m}$.

Figure C.4 shows the results of repeating this process for each gradient examined, here the dependence on both spatial resolution (step width) and the gradient of the ramp being sampled can be seen.

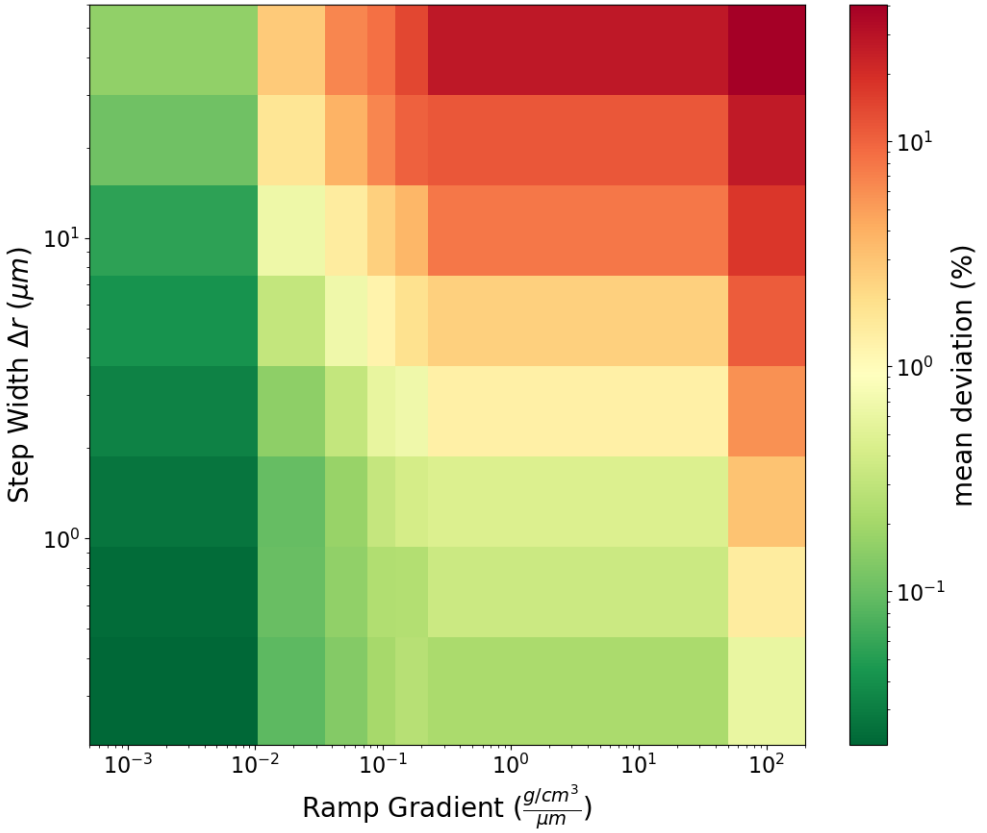


Figure C.4: Convergence graph for 0.05 \$g/cm$^3_{\{\mu}$m\}, mean percentage deviation is plotted as a funciton of spatial resolution

Returning to the results from chapter 6, and taking the timestep used in figure 6.3 (7.31ns into the implosion) as an example. The spatial resolution here varies across the lineout, since the points used were taken directly from the cell sizes that HYADES uses, and these vary across the density profile. The spatial resolution of this density profile can be doubled by adding points half way between each set of points on the original profile. Using the same Monte-Carlo seed, XPCI lineouts were produced for the original density profile and the profile with double the resolution, shown in figures C.5 and C.6 respectively. In these figures the different materials of the shell are shown in different colours.

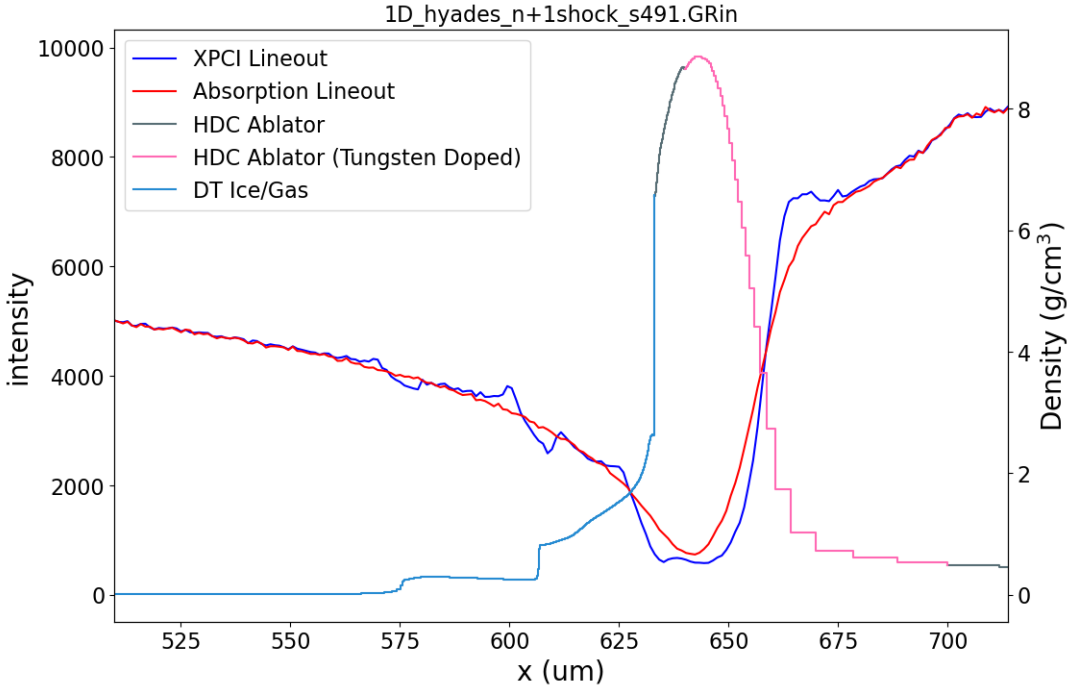


Figure C.5: Example synthetic XPCI lineout of the big-dip design at t=7.31ns into the implosion, the density used is also plotted at the resolution used in chapter 6, with different colours for different materials. An absorption only lineout is also shown.

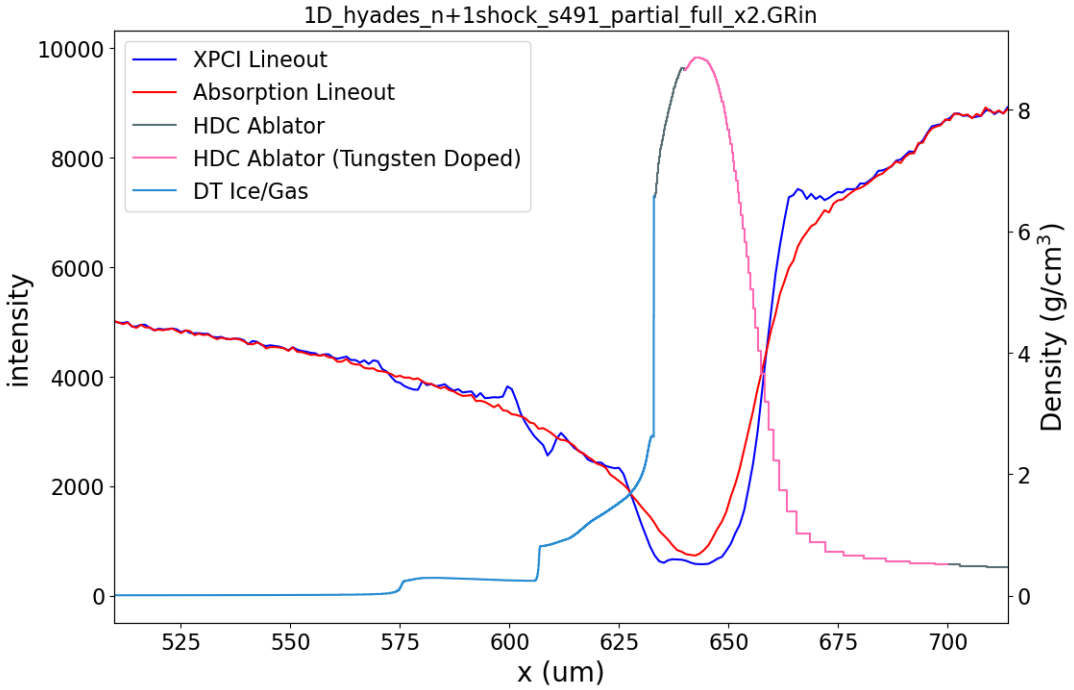


Figure C.6: Example synthetic XPCI lineout of the big-dip design at t=7.31ns into the implosion, the density used is also plotted at double the resolution used in chapter 6, with different colours for different materials. An absorption only lineout is also shown.

Here it is clear that the area around the N+1 shock is well resolved, but the area around the ablation front, in the doped layer of the shell, is slightly under resolved. Under resolving the ablation front is acceptable in this case, as it is not a feature of interest, and under resolving it does not affect the lineout near the N+1, which is the feature of interest. While it would be possible to increase the resolution so that this area was better resolved, this would come with the cost of requiring significantly more computational time, which is problematic given the large number of images produced in chapter 6.

To look at the convergence of the full range of density profiles used in chapter 6, the points from the density profiles can be overlayed over figure C.4. This can be done by taking each point in each density profile used, and calculating the gradient from one point to the next, as well as the spatial resolution of that point, then plotting it on the heatmap. The left panel of figure C.7 shows this overplotting with the points from the doped layer included, while the right panel shows the same plot with the doped layer points removed.

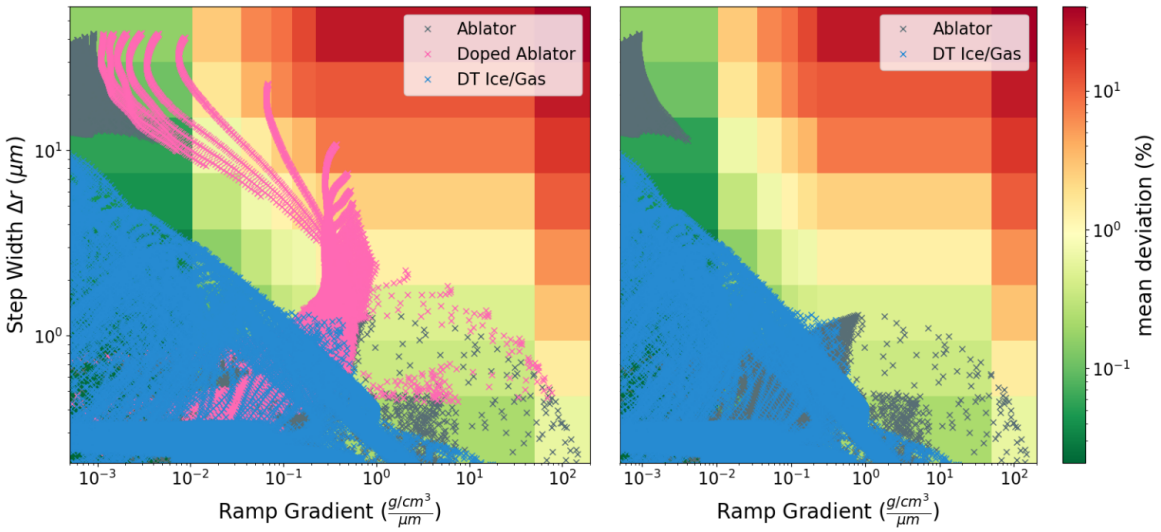


Figure C.7: Convergence plot for the full range of density ramp gradients and spatial resolutions. The left panel has all the points for all the density profiles used in chapter 6 plotted at their respective density gradient and step width, the right panel excludes the doped ablator points as they are far from the region of interest.

The figures show that all the points from the DT ice and non-doped ablator density profiles are well resolved, while the doped layer is under resolved in some cases. As the doped layer will always be far from the ice layer where the feature of interest (the N+1 shock) is present, this shows that the simulations are sufficiently well resolved to be used in chapter 6.