

Multidimensional Modelling of Cross-Beam Energy Transfer for Direct-Drive Inertial Confinement Fusion

Philip William Moloney

I M P E R I A L

September 2024

Submitted in partial fulfilment of the requirements for the degree of
Doctor of Philosophy of Imperial College London

Department of Physics
Imperial College London
Prince Consort Road
London SW7 2AZ

Declaration

I hereby certify that the material of this thesis, which I now submit for the award of Doctor of Philosophy, is entirely my own work unless otherwise cited or acknowledged within the body of text.

Philip Moloney

September 2024

Copyright

The copyright of this thesis rests with the author. Unless otherwise indicated, its contents are licensed under a Creative Commons Attribution-Non Commercial 4.0 International Licence (CC BY-NC).

Under this licence, you may copy and redistribute the material in any medium or format. You may also create and distribute modified versions of the work. This is on the condition that: you credit the author and do not use it, or any derivative works, for a commercial purpose.

When reusing or sharing this work, ensure you make the licence terms clear to others by naming the licence and linking to the licence text. Where a work has been adapted, you should indicate that the work has been changed and describe those changes.

Please seek permission from the copyright holder for uses of this work that are not included in this licence or permitted under UK Copyright Law.

Abstract

Accurate modelling of laser propagation and energy deposition in direct-drive Inertial Confinement Fusion (ICF) is crucial to performing predictive simulations. Cross-Beam Energy Transfer (CBET) is an interaction of laser light and the plasma state of matter, which can halve the deposited power in direct-drive implosions and amplify deposition asymmetries by an order of magnitude. Computational models which do not account for CBET, can therefore not be truly predictive. This thesis presents the development, validation and use of SOLAS: a 3-D, ray-based CBET model, which has been integrated into the CHIMERA code.

The development of SOLAS and its integration into the CHIMERA code are described. Test problems were conducted which validated the ray-trace, energy deposition, electric field reconstruction and CBET solver. CHIMERA-SOLAS simulations of direct-drive targets on the OMEGA laser facility are presented, which are in excellent agreement with existing codes.

A study is presented, which investigated the role of CBET with respect to the beam radius initial condition for direct-drive implosions. The results demonstrate that in the absence of CBET, increasing the beam radius improves the stagnation state symmetry of implosions. However, a larger beam radius leads to more CBET and a subsequent symmetry degradation. These results could help to explain observed trends in statistical modelling of OMEGA implosions.

A final study is presented, which aimed to understand how the role of CBET changes when a target is magnetised prior to the laser drive. Magnetised direct-drive experiments have demonstrated that anisotropic thermal conduction in the coronal plasma can lead to asymmetry of the implosion. Simulations of magnetised direct-drive targets were conducted, to understand if coronal magnetisation affected CBET scattering. Although CBET was observed to be dynamically significant to the implosions, the interplay of CBET with the magnetic field was found to be minimal.

Role of the Author

The research presented in this manuscript has been conducted using the multidimensional CHIMERA Radiative-Magnetohydrodynamics (Rad-MHD) code. CHIMERA, a variant of which is also referred to as GORGON, is developed and maintained by the Centre for Inertial Fusion Studies (CIFS) at Imperial College London [1]. Many authors have contributed to the development of this code over two decades. All simulations discussed in this manuscript were carried out by the author of this thesis. Analyses of the results presented in this manuscript are the work of the author of this thesis. The SOLAS laser ray-tracing and CBET module was developed by the author and is presented in Chap. 3. All development, implementation and debugging of this solver was the work of the author of this thesis. Most of the test problems used to validate SOLAS are the work of other authors, who are referenced. The field reconstruction and CBET test problems, published by Russ Follett *et. al.* in Ref. [2], were particularly helpful to validate the solver. Analysis of the results from these test problems was carried out by the author of this thesis. The previous laser ray-trace module was developed by Christopher Walsh, and some functions from SOLAS are based on this implementation. The author also contributed to the development of other aspects of the code, including the thermal conduction module, which is not presented in this thesis. Initial conditions for the magnetised simulations conducted in Chap. 5 were provided by Cody Wu Chang and Johan Frenje at the Massachusetts Institute of Technology (MIT). These simulations were based on experiments conducted at the OMEGA laser facility at the Laboratory for Laser Energetics (LLE). Consultation and advice about implementation and usage of the CHIMERA code came from all members of CIFS. Permissions for any external copyrighted material used can be found in Appendix B.

Acknowledgements

The past four years have been simultaneously incredibly rewarding and challenging. Without the love and support I have received from the many friends, family members and colleagues I have spent this time with, I know that I would have achieved little, and enjoyed it far less. First, I would like to thank my supervisor, Jerry, who has provided me with much advice, support and guidance. I think that my three-month side project of making a CBET model might almost be finished.

The many members, past and present, of CIFS and the wider Imperial plasma group have made each and every day working here a pleasure, and I would like to thank them all. Particularly (in alphabetical order) Adam D. for being an excellent Griffin replacement, Adam E. for being unceasingly joyful, Aidan B. for helping me to overcome (some) of my squeamishness, Aidan C. for knowing everything about everything, Ben for giving me a good fructose and breakfast sausage contest, Brian for some amazing pre-conference trips, Dom for the best (and toughest) bike ride of my life, Nic for teaching me how to surf, Niki for many cups of tea, advice and chat, and Sam for our fruitful collaborations.

I am delighted to have such good friends outside of work as well. Huge thanks to Tom and Ronan. It scares me to think that soon we will have lived apart for as long/ longer than we lived together! I am not sure that I would have made it through my undergrad, let alone a PhD, if I didn't have you both as incredible friends. Thanks also to my many other friends and housemates from my time in London, particularly Tommy M., Greg, Silvia, Liv and Chris. I look forward to many more weekend gatherings, general silliness and trips away with Ben, Push and Robbo. Regularly seeing you guys makes me laugh and reconnects me to home. Special thanks to Griff. It still feels weird not seeing you every day, but your friendship and support has been constant through these four years, and I am incredibly grateful for it.

So much thanks to my family. I feel so privileged that there are too many of you to name each individually, but I love spending time with each and every Moloney, Cowie, Phillips and Thomas. In particular, I must mention Rory and Granny, who I miss hugely. Feeling like a big, fun cousin to Rory helped shape me as a teenager and I know that Granny would have truly loved to see me finish my PhD. Mum, dad and Liv, what can I say. I don't know what you could have done to support me more over the past four years, and the twenty-three years before that. You are the best family I could have asked for.

Finally, to Ewa. You've been by my side from first applications all the way through to now, and every day you've made me smile and feel loved. I wouldn't be where I am without you. All my love and gratitude goes to you.

Contents

1	Introduction	27
1.1	Nuclear Fusion	27
1.2	Inertial Confinement Fusion	30
1.2.1	Ignition Requirements	30
1.2.2	Central Hotspot Ignition	32
1.2.3	Alternative Ignition Approaches	34
1.2.4	Inertial Fusion Energy Considerations	35
1.3	Current Experiments and Main Approaches	36
1.3.1	Indirect-Drive	36
1.3.2	Direct-Drive	39
1.4	Objective of the work	44
2	The Interaction of Light with Plasma	47
2.1	Basic Plasma Physics	48
2.2	The Kinetic and Fluid Formulations of Plasma	49
2.2.1	The Kinetic Formulation	50
2.2.2	The Fluid Equations	51
2.2.3	Radiation Transport	53
2.2.4	Magnetohydrodynamics	54
2.2.5	Kinetic Heatflow	56
2.3	Fluid Description of Waves in Plasma	57
2.3.1	Plasma as a Dielectric Medium	57
2.3.2	Electromagnetic Waves	59
2.3.3	Plasma Waves	60
2.4	Propagation and Absorption of Light in a Plasma	64
2.4.1	The WKB Approximation	64
2.4.2	Ray-Tracing	66
2.4.3	Inverse-Bremsstrahlung	68
2.4.4	Resonance Absorption	71
2.5	Laser Plasma instabilities	72
2.5.1	The Ponderomotive Force	72
2.5.2	Three-Wave Coupling with Two Light Waves	74
2.5.3	Cross-Beam Energy Transfer	75
2.5.4	Mitigation of Cross-Beam Energy Transfer	79

3	SOLAS, a 3-D Laser Ray-Trace and Cross-Beam Energy Transfer Model	81
3.1	Ray-Tracing for Hydrodynamic Simulations of Fusion Plasmas	82
3.2	Existing Cross-Beam Energy Transfer Models	83
3.2.1	Ray-Based Models	83
3.2.2	Wave Solvers	85
3.2.3	Particle in Cell Codes	86
3.3	SOLAS 3-D Laser Ray-Trace	86
3.3.1	Outline of SOLAS-CHIMERA Interfacing	87
3.3.2	SOLAS Mesh Structure	87
3.3.3	Ray Initialisation	91
3.3.4	Equations of Rays and Adaptive Integration	94
3.4	Ray-Based Field Reconstruction and Ray Sheets	100
3.4.1	Ray-Amplitude and Field Estimate from Neighbour-Rays	101
3.4.2	Caustic Field Capping	105
3.4.3	Field Reconstruction Validation	107
3.5	Ray-Based CBET Model	110
3.5.1	Power Change of Rays due to CBET	111
3.5.2	Caustic Gain Truncation	115
3.5.3	Coherent Caustic Correction and Caustic Region Identification	116
3.5.4	Dynamic Memory for Storing Fields and CBET Gains	118
3.5.5	Pump Depletion Iterations	118
3.5.6	Energy Conservation Iterations	120
3.5.7	CBET Validation	122
3.6	Simulations of CBET for OMEGA Direct-Drive Implosions	127
3.6.1	1-D hydrodynamics, 3-D Ray-Trace Simulation of OMEGA Shot 89224	127
3.6.2	3-D Post Process of Absorption Non-Uniformity	131
3.7	Conclusions	134
4	Cylindrical Simulations to Study the Effect of Beam Radius in Direct-Drive	136
4.1	Introduction to Beam Radius in Direct-Drive Inertial Confinement Fusion	137
4.1.1	Previous Work Studying the Effect of Beam Radius on OMEGA	139
4.1.2	Statistical Modelling of OMEGA Direct-Drive Implosions	140
4.2	Cylindrical Simulation Platform for Beam Radius Parameter Scan	144
4.2.1	Advantages and Validity Considerations of the Cylindrical Simulation Platform	144
4.2.2	Pulse Shape and Target Initial Conditions	145
4.2.3	1-D Implosion Tuning	146
4.3	Asymmetry of Deposited Power	148
4.3.1	Analysis and Quantity Definitions	149
4.3.2	Deposition Asymmetries in the Absence of CBET	151
4.3.3	CBET Imprint on Incident Field	152
4.3.4	Deposition Asymmetries in the Presence of CBET	154

4.4	Stagnation State Asymmetry	156
4.4.1	Stagnation State Asymmetry Trend with Beam Radius	157
4.4.2	Time Resolved Asymmetry Growth	160
4.4.3	Comparison of Results to Statistical Modelling	162
4.5	Conclusions	163
4.5.1	Summary of work	163
4.5.2	Future Work	163
5	Simulations of Cross-Beam Energy Transfer for Magnetised Direct-Drive	165
5.1	Magnetised Inertial Confinement Fusion and Exploding-Pushers	166
5.1.1	Potential Benefits of Target Magnetisation	166
5.1.2	Experimental Studies of Magnetised-ICF	168
5.1.3	Magnetised Laser-Plasma Instabilities	170
5.1.4	The Exploding-Pusher Configuration	170
5.2	Cross-Beam Energy Transfer in Unmagnetised Exploding-Pushers	171
5.2.1	Simulation Configuration	171
5.2.2	1-D Simulations	174
5.2.3	2-D Simulations	176
5.3	Magnetisation in Exploding-Pusher Implosions	178
5.3.1	In-Flight Field Structure	179
5.3.2	Resistive Diffusion and the Lorentz Force	180
5.3.3	Anisotropic Thermal Conduction	181
5.3.4	The Nernst Effect	183
5.4	The Effect of Magnetisation on Cross-Beam Energy Transfer and Stagnation . .	185
5.4.1	Analysis and Key Definitions	186
5.4.2	Spatial Change of CBET and Deposition from Magnetisation	188
5.4.3	Stagnation Profiles	190
5.5	Conclusions	194
5.5.1	Summary of Work	194
5.5.2	Future Work	195
6	Conclusions	196
6.1	The SOLAS Ray-Trace and CBET Model	197
6.1.1	Model Development and Validation	197
6.1.2	Potential SOLAS Enhancements	198
6.1.3	Potential CHIMERA Enhancements	199
6.2	Effect of Beam Radius on Beam-Mode Asymmetry	200
6.2.1	Summary of Simulations	200
6.2.2	Suggestions for Further Work	201
6.3	The Effect of Coronal Magnetisation on CBET Scattering	202
6.3.1	Summary of Simulations	202
6.3.2	Suggestions for Further Work	203

Contents	11
6.4 Final Remarks on CBET Modelling	203
A List of Acronyms	206
B Copyright Permissions	209
Bibliography	211

List of Tables

- 3.1 The ratio of the power of the seed beam as it exits the domain to its initial power.123
- 3.2 Integrated metrics from simulations of OMEGA 89224. LILAC metrics are from
Refs. [202, 203, 204]. 127
- 4.1 Results of the 1-D Tuning Simulations. 147
- 5.1 Spatially and temporally integrated results from all simulations. In the CBET
column, ‘~’ indicates partial-CBET, where CBET reduced absorption magni-
tude, but *not* spatial location, of deposition. 174
- B.1 Copyright permissions for material reproduced in this thesis. 209

List of Figures

1.1	Binding energy per nucleon for common nuclear isotopes. Binding energy peaks close to ^{56}Fe , and energy is released for reactions which increases binding energy of the products compared to the reactants. ^4He has a particularly high binding energy and therefore fusion reactions resulting in ^4He are strong candidates for fusion energy production.	28
1.2	Averaged reactivities for important fusion reactions. The deuterium-tritium reactivity is significantly larger than all other reactions up to $T \sim 500$ keV. Averaged reactivities are obtained from cross-section data, available in the ENDF/B-VII.1 library [7].	29
1.3	Four key stages of the central hotspot ignition ICF concept. The chronological order of the diagram is top-left, bottom-left, bottom-right and then top-right quadrants, which display the initial configuration, implosion, stagnation and burn propagation stages respectively.	32
1.4	Energy balance of an Inertial Fusion Energy (IFE) power plant. Based on a similar figure from Ref. [5].	35
1.5	Schematic of the indirect-drive approach to ICF. Laser light, represented as blue transparent rectangles, irradiates the interior of a high-Z (<i>e.g.</i> gold) hohlraum, producing thermal x-rays that drive the capsule implosion. A number of important physical effects, typical to indirect-drive experiments, are also illustrated.	36
1.6	G plotted against time for indirect-drive ICF shots on National Ignition Facility (NIF). The colour of each bar represents a different implosion design, and the dashed horizontal line represents $G = 1$, where the fusion energy produced is equal to the incident laser energy. Used under CC BY 4.0 from Ref. [47].	38
1.7	Schematic of the direct-drive approach to ICF. Laser light, represented by the transparent-blue rectangles is absorbed outside the critical surface and then transported in to the ablation surface by thermal conduction. Speckle from beam smoothing optics leads to short-scale non-uniformity of deposition, known as imprint. Laser-Plasma Instabilities (LPIs) degrade the performance, both by reflecting light, and by generating hot electrons, which pre-heat the fuel.	39

1.8	Schematic of beam smoothing techniques employed on the OMEGA laser system. Phase aberrations can lead to large scale modulations to the far-field intensity profile. Phase plates shift the perturbations to shorter scales, which are more efficiently smoother by thermal conduction. Smoothing by Spectral Dispersion (SSD) creates a time-varying speckle pattern, which leads to a smooth, <i>time-integrated</i> far-field profile. Polarisation smoothing splits the beam into orthogonally polarised sub-beams, which do not interfere and thus further reduce non-uniformity.	41
1.9	Important laser-plasma interactions for direct-drive ICF. Energy is absorbed by Inverse-Bremsstrahlung (Inv-Brem) close to critical, and resonance absorption at the critical surface. Hot electrons pre-heat the fuel, reducing compressibility and are generated from an Electron Plasma Wave (EPW) from Two Plasmon Decay (TPD) or Stimulated Raman Scattering (SRS). Energy is reflected by CBET, SRS and Stimulated Brillouin Scattering (SBS). The filamentation instability can cause light to self-focus and thus enhance asymmetry of the deposition. It also raises intensity of the light locally, enhancing the growth rate of other LPIs.	42
2.1	Particle motion in the electric fields of a) transverse and b) longitudinal plasma waves.	58
2.2	The dispersion relation for an Electro-Magnetic Wave (EMW) propagating in a plasma. For $\omega < \omega_p$, the oscillatory plane wave Ansatz turns into a decaying evanescent wave, thus the wave does propagate.	60
2.3	Illustration of electron (small blue circle) and ion (large red circle) motion in a) an Ion Acoustic Wave (IAW) and b) an EPW.	61
2.4	An electron-ion collision in the presence of an oscillating electric field. The collision time is much shorter than the oscillation time. The magnitude of the electron velocity just before ($\mathbf{w}$) and just after ($\mathbf{w}'$) the collision is the same, since the collision is elastic. However, the electron continues its quiver motion ($\mathbf{u}$) just after the collision and thus the time-averaged velocities before ($\mathbf{v}$) and after ($\mathbf{v}'$) are different. The figure has been reproduced with permission from Ref. [116].	68
2.5	Illustration of resonance absorption for a non-normally incident, p-polarised light wave, propagating up a density gradient. At the turning point of the light, if sufficiently close to the critical density, then the evanescent light can tunnel through to the critical surface and resonantly excite an EPW. The EPW propagates down the density ramp and transfers energy to the electrons through damping. The figure has been reproduced with permission from Ref. [116]. . .	71
2.6	Two plane light waves crossing each other and leading to a beat pattern in the summed field. When in a plasma background, the charged particles will experience a drift motion down the beat pattern envelope, known as the ponderomotive force.	74

- 2.7 The frequency and wave number matching conditions for LPIs, whereby a pump EMW, labelled 0, loses energy to two daughter waves. Here the daughter waves are another EMW with lower frequency ($\omega_1 < \omega_0$) and an IAW, which are labelled 1 and 2, respectively. EMW $\rightarrow$ EMW+IAW identifies the process as CBET or SBS. 75
- 2.8 A 1-D Illustration of the mechanism by which CBET grows. Panel a) plots the incident electric fields from two lasers, overlapping in a plasma background. This creates a beat pattern which ponderomotively pushes charges within the plasma to regions where the beat envelope is small, as is shown in panel b). If the spatial and temporal frequency of the beat oscillation matches the IAW dispersion relation, then an IAW grows, which is shown in panel c). The density modulation from the IAW acts as a grating which transfers energy from one beam to the other. 76
- 2.9 The two CBET modes which are identified from the fluid CBET gain in Eq. 2.117, in the context of direct-drive ICF. The backscatter mode (left) dominates for direct-drive ICF. In this regime, the light waves have approximately equal frequencies and therefore the required frequency difference in the plasma frame to excite an IAW is a result of the Doppler shift. In the sidescatter mode (right) which is employed in indirect-drive to tune symmetry [44], the flow is orthogonal to the wave vector and therefore CBET only occurs due to a frequency difference between the waves. This figure has been reproduced with permission from Ref. [131]. 77
- 3.1 3-D ray trajectories through a spherically symmetric, OMEGA direct-drive scale density profile. The incident rays are plotted in cyan and the reflected rays, which spread out over a large solid angle, are plotted in dark blue. The critical density is represented by the black mesh. 84
- 3.2 A flowchart illustrating the basic interfacing between CHIMERA and SOLAS. The interpolation steps refer simultaneously to the re-domain balancing and cell combination procedures described in Sec. 3.3.2. 87
- 3.3 An illustrative diagram demonstrating the MPI re-domain balanced grid employed by SOLAS compared to the CHIMERA domain balanced grid for a cylindrical geometry. 88
- 3.4 Illustrative diagrams demonstrating a) the spherical polar mesh used by CHIMERA for hydrodynamic calculations, b) the cell combination mechanism employed by SOLAS to obtain a roughly equal area grid for spherical simulations and c) the adaptive radial cell combination to reduce resolution in regions where CBET and refraction are unimportant. 90

- 3.5 Example initial locations for ~ 4000 rays for an OMEGA beam port with super-Gaussian intensity profile, with a shape defined by $n_s = 5.2$ and $\sigma = 352 \mu m$. The inverse-projection ray locations are shown in sub-figure a) which demonstrate a higher ray density for rays pointed to the pole at $x \sim -0.15 mm$. b) shows the randomly sampled rays where all rays are initialised with the same power, and c) shows uniformly sampled rays. 91
- 3.6 Shape function smearing for ray power interpolation to the SOLAS grid for a cylindrical mesh. The shape function has the same size and geometry as the SOLAS grid cell which the ray is located in. A top-hat shape function is used, so the fraction of the deposited power interpolated to a given cell is proportional to the volume overlap of the shape function with the cell. 96
- 3.7 Results of the quadratic trough ray trajectory test problem for the adaptive solver with 3 different cases. These are default settings (linear interpolation of both n_e and ∇n_e), the Kaiser algorithm (linear interpolation of n_e and uniform ∇n_e in a cell) and the adaptive algorithm with tricubic interpolation of n_e). The top plot shows ray trajectories and analytic trajectory and the bottom plot shows absolute errors for the 3 cases. 97
- 3.8 Results from the cylindrical-helix test problem, obtained using the default adaptive RK45 algorithm. a) shows the trajectory of a ray from a simulation with 100 radial cells in 3-D space. b) shows the trajectory of the ray projected on the $x-y$ plane, along with the n_e profile. c) shows the error of the test, defined as the fractional difference in radius of the ray as it exits the domain from the initial radius, as a function of number of radial cells. 99
- 3.9 Results of the blast wave Inv-Brem absorption test problem. The colour-plot shows the temperature of the 1-D simulation as a function of time (x -axis) and space (y -axis). The heatfront obtained from CHIMERA, which is the maximum x location where $T_e > T_{e0}$, is in good agreement with the analytic solution. . . . 100
- 3.10 A 3-D illustration of the ray-amplitude construction for a single ray in a spherically-symmetric, direct-drive like plasma profile. The critical surface is shown as the grey wire mesh. The main-ray trajectory is plotted in red, and its three neighbour-ray trajectories are in blue. The area of the main-ray at 5 successive phase-fronts is plotted, *i.e.* when the main- and neighbour-rays all have the same phase, φ , is illustrated by the shaded triangles. This is inversely proportional to the ray-amplitude at this location. 101
- 3.11 A demonstration of how the ray-amplitude is obtained from the area of neighbouring-rays for a beam propagating up a linear density gradient. Fig. a) explicitly shows the trajectory of a main-ray and its neighbour-rays in red and blue respectively, explicitly highlighting the beam's caustic region, where the amplitude of the light is capped. Fig. b) plots the area of the main-ray in blue (which goes through a minimum at the caustic location), as a function of the phase of the ray. Also shown are the uncapped and capped amplitudes of the ray. . . . 102

- 3.12 An illustration of the concept of caustics and ray sheets. In a), the trajectory of rays traversing up a linear density gradient are plotted. Separate colours are used for the incident and reflected sheets, which are separated by a caustic. In b), rays propagate up a density profile, $n_e = n_{\text{cr}} \exp [-(r_{\mu\text{m}} - 20)/100]$. Unlike in a), the ray caustics do not occur at their turning point, *i.e.* the minimum radius. 104
- 3.13 Results of the 1-D field reconstruction test, comparing SOLAS fields with different caustic field capping methods to the LPSE results. Panel a) shows the oscillatory behaviour of the field solution which arises due to the interference between the incident and reflected sheet. Panel b) plots the same results, but zoomed in on the caustic region to demonstrate the differences in fields obtained from the different methods. 108
- 3.14 Results of the 2-D field reconstruction test. Panel a) plots the field from LPSE on the left side and the SOLAS field on the right for the 1/64 OMEGA scale simulation. Panel b) shows the SOLAS field from the 1/4 scale simulation which, when compared to a), demonstrates that as scale increases, the relative size of the caustic region decreases. Panels c), d) and e) are lineouts near the caustic region along y from the 1/64 scale simulations at $x = 0, 3$ and $6 \mu\text{m}$ respectively, the positions of which are indicated by the white lines on a). Note that the circular features in the field in panels a) and b) arise due to aliasing artefacts between the true solution scale, with the resolution of the image and/ or simulation grid. 109
- 3.15 The SOLAS CBET model operation loop. F_i is shorthand for the total field from all sheets at pump depletion iteration i and f_{CP} is the modifier to the caustic pump multiplier, described in Sec. 3.5.6 of the j^{th} energy conservation iteration. ΔE denotes the energy conservation error, which is the total incident power minus deposited power and power exiting the grid. 110
- 3.16 Illustration of the Caustic Gain Truncation (CGT) and caustic region identification algorithms. Panel a) shows a laser, separated into incident and reflected sheets by a caustic, propagating up and reflecting in a linear density ramp. CGT is illustrated in b) where a cell is separated into a *lit* and *unlit* region by the reconstructed *caustic plane* (dashed magenta line). Panels c) and d) illustrate the method used to identify the *caustic region* (shaded blue), where the coherent caustic modifier is applied. c) shows that cells are geometrically separated into *capped* (blue squares) and *uncapped* (red squares) regions and the coherent caustic modifier to the gain is applied in this capped, caustic-region. 114

- 3.17 Results of successive pump depletion iterations for an illustrative simulation of two beams crossing and undergoing a CBET resonance in a uniform plasma background. Each beams' rays from the field reconstruction ray-trace are plotted in a), where more opaque rays indicate a higher power. Panel b) plots the reconstructed intensity from this ray-trace, where $i = 0$ indicates the field reconstruction ray-trace and therefore no CBET occurs. Panels c) and d) plot the intensity reconstructed from the 1st and 7th pump depletion iterations respectively, where the effect of transfer of powers between beams due to CBET is evident. Panel e) plots the convergence parameter from the simulation over successive pump depletion iterations, with convergence defined here as $K_{\text{Conv},0} = 10^{-6}$ 119
- 3.18 Energy conservation iterations for a two-beam caustic test problem, which is presented in Sec. 3.5.7.2. 4 iterations were required to achieve energy conservation, which was defined to be $\delta E < 1 \times 10^{-4}$. The order of the iterations is indicated by the text next to each point. 120
- 3.19 Results of the 2-D CBET test in the absence of laser caustics. Panels a) and b) show the electric field from the SOLAS and LPSE simulations respectively, after the pump depletion iterations have occurred. Panels c) and d) plot lineouts from the inbound fields of the pump and seed beams respectively for both SOLAS and LPSE. e) and f) plot the same for the outbound fields. 122
- 3.20 Results of the 2-D CBET test in the presence of a laser caustic. Panels a) and b) show the electric field from the LPSE and SOLAS simulations respectively. Also plotted is an inset zoom on the caustic region of the seed beam, which demonstrates that SOLAS does not capture the evanescent field. Panels c) and e) plot lineouts from the inbound fields of the pump and seed beams respectively for both SOLAS and LPSE. d) and f) plot the same for the outbound fields. 124
- 3.21 Results from a resolution scan of the two-beam caustic CBET test both with CGT and without. Plotted in a), b) and c) are the power amplification of the seed beam, initial energy conservation error before energy conservation iterations and caustic pump multiplier used to obtain energy conservation respectively. 125
- 3.22 Results from the 16 beam CBET test problem. Panel a) plots the deposited power both without (left) and with (right) CBET from the simulation ran in cylindrical coordinates at $1 \mu\text{m}$ radial resolution. Panel b) plots the results of the resolution scan for both Cartesian and cylindrical coordinates. The circles and squares are the values of power deposition without and with enforced energy conservation respectively. The shaded regions illustrate the magnitude of energy conservation error from the initial energy conservation iteration, which can be thought of as the (single-sided) error on the circular points. 125
- 3.23 The initial conditions for OMEGA shot 89224. The shot was a cryogenic, DT implosion with materials and radii shown in panel a). Panel b) plots the incident laser pulse shape. 127

- 3.24 CHIMERA-SOLAS simulation of OMEGA shot 89224. Panel a) shows the deposited power both from 2 separate simulations of 89224, one with and one without CBET as red and blue solid curves respectively. The shell radius, defined at the radius of maximum density, is also plotted on the right axis as dashed curves. Panel b) plots the time resolved partition of non-absorbed light for these simulations. Plotted in c) is the energy conservation error after energy conservation was achieved, alongside the value of f_{CP} that was used to achieve conservation. 128
- 3.25 CBET deposition diagnostics from $t = 1.6$ ns into the CHIMERA-SOLAS simulation of OMEGA 89224, with CBET. Both sub-plots plot the power deposition with and without CBET, along with the g_{CBET} , which is defined as the sum of the absolute value of all power lost or gained by rays via CBET in the cell. The electron density normalised to critical density and radial velocity normalised to the sound speed are also plotted in panel a) and b) respectively. The blue shaded region in b) represents the full-width half-maximum of the g_{CBET} curve. 130
- 3.26 3-D SOLAS post-process of spherically symmetric hydrodynamic data obtained from a LILAC simulation of an OMEGA direct-drive implosion. Figures a) and b) show the deviation in radially integrated deposited power without and with CBET respectively. Figures c) and d) plot the spherical harmonic modal decomposition for these maps alongside results from the IFRIIT and BEAMCROSSER codes without and with CBET respectively. 132
- 3.27 Electric fields from the 3-D post-process. Panels a), b), c) and d) plot the incident sheet without CBET, the reflected sheet without CBET, the incident sheet with CBET and the reflected sheet with CBET respectively. The critical surface is plotted with the inbound sheets as a transparent spherical wireframe. Note that the outbound sheets are plotted on a log scale. 134
- 4.1 The trajectory of rays from two beams through a direct-drive like coronal plasma. The density profile for both simulations is $n_e = n_{cr} \exp [-(r_{\mu m} - 20)/100]$. Panels a) and b) plot rays from beams with widths $\sigma = 10$ and $18 \mu m$, respectively. Ray trajectories are separated for each beam by colour depending on their sheet. Rays coloured in red and dark blue are from the incident sheet (before the ray caustic) and rays coloured in magenta and light blue are from the reflected sheet (after the ray caustic). 137
- 4.2 Soft x-rays emitted from the ablation surface of direct-drive implosions with various R_b/R_t values, as measured by an x-ray framing camera. All images were taken at a constant capsule radius of $R = 175 \mu m$. The figure has been reproduced with permission from Ref. [142]. 139

- 4.3 Experimentally inferred fusion yield degradation due to the finite beams on the OMEGA laser facility. The dotted orange curve is the fit obtained from just using the $\bar{R}_{b/t}$ relation, while the black dashed curve uses the full relation in Eq. 4.7. The vertical dotted line indicates the critical threshold, $\hat{R}_{b/t}^{\text{crit}} = 0.86$, after which the $\hat{R}_{b/t}$ also has an effect. The cross validation error from the $\bar{R}_{b/t}$ and full fit is -1.0% and -0.5% respectively. The figure has been reproduced with permission from Ref. [73]. 142
- 4.4 Predictions for how the absolute yield varies across OMEGA direct-drive experiments with R_b/R_t . A change in behaviour is clearly seen at $R_{b/t}^{\text{crit}} = 0.86$, at which point the yield trend with R_b/R_t flips from monotonically increasing to decreasing. The data in this figure is obtained from Ref. [223]. 144
- 4.5 The a) target initial conditions with beam geometry, and b) pulse shape used for the 2-D cylindrical simulations. All beams were polarised out of the simulation plane, in the $+\hat{z}$ direction. Initial layer radii were taken from the initial conditions for OMEGA shot 89224, presented in Fig. 3.23.a. 146
- 4.6 Streak plots from two of the 1-D tuning simulations including CBET. Panels a) and b) plot the electron density of the $R_b/R_t = 0.8$ and $R_b/R_t = 1.0$ simulations, respectively, as a function of time (x -axis) and radius (y -axis). Panels c) and d) plot the electron temperature for the $R_b/R_t = 0.8$ and $R_b/R_t = 1.0$ simulations, respectively. 148
- 4.7 Demonstration of the analysis workflow to obtain the key results for this chapter. The power deposition at $t = 1.12$ ns from the CBET (left) and no-CBET (right) simulations are plotted in panel a) for the $R_b/R_t = 0.85$ case. Panel b) plots the radially integrated deposition from the profiles in a) as a function of azimuthal angle. As can be seen from this plot, the CBET asymmetry (light-blue) is greater than the no-CBET asymmetry (red). The power spectrum of these profiles is then plotted in panel c). This demonstrates that the dominant modes in the spectrum are multiples of the number of beams. 149
- 4.8 Radially integrated deposited power from no-CBET simulations as a function of time (x -axis) and angle (y -axis), alongside amplitudes of the dominant modes from a Fourier power spectrum. Panels a) and b) plot the radially integrated deposited power and Fourier modes, respectively, for the $R_b/R_t = 0.8$ simulation. The same is plotted for the $R_b/R_t = 0.9$ simulation in c) and d), and for the $R_b/R_t = 1.0$ simulation in e) and f). The mode 10 from the number of beams is clearly visible in the radially integrated power plots as 10 peaks to troughs in angle at a given time, *i.e.* 10 cyclical perturbations along a vertical lineout. . . . 151

- 4.9 Demonstration of a mode-flip in P_{dep} for the no-CBET $R_b/R_t = 0.8$ simulation. Panel a), b) and c) plot the power deposition just before, during and just after the mode-flip. Panel d) plots the deviation from the mean of the radially integrated profiles around the azimuthal angle for all 3 times. It can be seen from this plot that the deposition is very symmetric at $t = 0.7$ ns. In all four panels, the angles of the beams are shown by the dashed magenta lines. The overlap region where deposition rises as the scale length increases, is highlighted with a dashed circle in panels a), b) and c). 152
- 4.10 Field structure which leads to CBET induced asymmetry on power deposition and its dependence on R_b/R_t and target convergence. Each panel plots the incident sheet field, which includes the effect of CBET, along with contours of the critical electron density (green) and the incident field, $|E_z^{\text{in}}| = 1 \times 10^{10} \text{ Vm}^{-1}$ contour of another beam (magenta). Panel a) and b) plot this for the $R_b/R_t = 0.9$ simulation at $t = 1.0$ ns and $t = 1.5$ ns respectively. The convergence of the target in this time interval, leads to beams wrapping around the target more, and therefore a change in the spatial location of the resonant CBET interaction. Panel c) plots the same for the $R_b/R_t = 0.80$ simulation at the same time as panel a). This demonstrates that the $R_b/R_t = 0.80$ beam is not wide enough at this time to lead to a resonant CBET interaction, unlike the wider beam in panel a). 153
- 4.11 This figure plots the same as Fig. 4.8, but now for the equivalent simulations which include the effect of CBET. Comparing these results and those in Fig. 4.8 demonstrates that CBET introduces additional modal-flips of the deposition and amplifies the magnitude of asymmetries. 155
- 4.12 Demonstration of a mode-flip in P_{dep} for the CBET, $R_b/R_t = 0.9$ simulation. Panels a), b) and c) plot the power deposition before, during and after the mode-flip. Panel d) plots the deviation from the mean of the radially integrated profiles around the azimuthal angle for all 3 simulations. Fig. 4.10.a and Fig. 4.10.b plot a single, incident sheet field from panels a) and c) respectively. It can be seen that the centre of the beam depletion at $t = 1.5$ ns leads to less deposition at the beam angles (magenta lines) on panel d). Panel d) shows that during the CBET induced mode-flip, a significant $\ell = 20$ occurs in the deposition. The reducing absorption at the centre of one beam is highlighted in by the dashed circles. 156
- 4.13 Densities of the DT fuel and ion temperatures for various R_b/R_t simulations both with and without CBET. Each row corresponds to a different R_b/R_t value; the left column contains simulations without CBET; and the right column contains simulations with CBET. It is visible from the density plots that increasing R_b/R_t improves stagnation symmetry for the no-CBET simulations, but degrades it for the CBET simulations. 158

- 4.14 Trends of a) stagnation asymmetry and b) maximum (azimuthally averaged) fuel density for CBET and no-CBET simulations. The no-CBET improvement in symmetry with R_b/R_t is observed which also corresponds to improved compression. The symmetry trend including CBET is more complex, but broadly the stagnation state symmetry is worse with increasing R_b/R_t . Note that simulations had a similar stagnation areal density, $\rho_{DT}R \sim 2.5 \text{ g cm}^{-2}$. This nominal value is provided to help the reader to appreciate the magnitude of asymmetry plotted in panel a). 159
- 4.15 Time resolved $\ell = 10$ Fourier power spectrum amplitude for a) $\rho_{DT}R$ and b) $P_{\text{dep}}R$ for CBET simulations with 3 R_b/R_t values. The developing, but unrealised, modal-flip for $R_b/R_t = 0.8$ from $t \sim 0.5 \rightarrow 1.2 \text{ ns}$ reduces the $P_{\text{dep}}R_{\ell=10}$, leading to slow $\rho_{DT}R_{\ell=10}$ growth and ultimately a relatively symmetric stagnation state. Despite large values of $\rho_{DT}R_{\ell=10}$ initially, the developing modal flip of the $R_b/R_t = 1.0$ simulation from $t \sim 1.8 \rightarrow 2.1 \text{ ns}$ slows the density asymmetry growth. 160
- 5.1 Cartoon to illustrate the effect of magnetisation on collisions, and therefore transport, of a positive test charge. Particle locations after collision are represented as blue circles and the path taken by the particle is shown by the black arrows. As the Hall parameter of the particle increases, diffusion is increasingly limited, and therefore collisional transport is reduced. 166
- 5.2 Simulated fusion yields versus effective residual fuel kinetic energy under imposed low-mode radiation flux perturbations for imposed fields in the range $B_0 = 0 \rightarrow 70 \text{ T}$. The plot demonstrates, that with increasing departure from ideal compression (moving to the right on the x axis), magnetisation can enable the onset of the ignition. Reproduced with permission from Ref. [237]. . . 167
- 5.3 a) A 1.1 keV T_i increase was achieved by adding a 26 T B_0 field to a D₂ gas capsule implosion on the NIF. Also shown in the plot are the simulation results. b)→d) Equatorial shapes of the implosions. Reproduced with permission from Ref. [241]. 168
- 5.4 X-ray self-emission images of (left) an unmagnetised and (right) a magnetised implosion. The average radius of the marked contour (corresponding to 40% of peak intensity), and the oblateness parameter R_a/R_b (ratio of major-to-minor axis) are listed below each image. The polar laser-drive is indicated by the white arrows, and the axis of the initial magnetic field by the black dashed line on the right. Applying an initial magnetic field demonstrated increased oblateness of the implosion. Adapted with permission from Ref. [243]. 169
- 5.5 The initial conditions used for all simulations presented in this chapter. Panel a) plots the D₂ filled, glass shell capsule and direction of the initial magnetic field. An example field coil (illustrative and not included in simulations) is also shown, the presence of which necessitated the polar laser drive in experiments. Panel b) plots the laser pulse shape used, which had a total of 17.7 kJ laser energy. 172

- 5.6 1-D simulation results both without (left) and with (right) CBET. The top row plots the incident and absorbed energy from the simulation on the left axis and the radius of maximum density on the right. In order, the next rows plot n_e , T_e , T_i and V_r . The full-width half-maximum times of the D₂ yield are plotted as dotted vertical black lines on all panels. 175
- 5.7 n_e (right-side), T_e (top-left-side) and T_i (bottom-left-side) plots from the 2-D, 0 T simulations without (left-column) and with (right-column) CBET at a variety of in-flight times. The decreased deposited energy due to CBET, results in lower coronal electron temperatures and therefore a slower, weaker shock being driven, which is especially evident at later times. 177
- 5.8 The development of the hydrodynamic variables and magnetic field structure from the $B_{z0} = 25$ T, CBET simulation. Panels a), b) and c) plot T_e (top-left), T_i (bottom-left), n_e (right) and $\mathbf{B}$ (streamlines) at $t = 0.15, 0.30$ and 0.45 ns respectively. The approximately radially outward flowing, hot (and therefore highly conductive) ablating plasma pulled the magnetic field with it, resulting in radial $\mathbf{B}$ field lines, which were weaker at the capsule equator. Panels d), e) and f) plot n_e lineouts along the pole ($\theta = 0^\circ$) and equator ($\theta = 90^\circ$). Panels g), h) and i) plot equivalent T_e and T_i lineouts. These all show that the increased polar temperatures, partially due to beam geometry and partially due to magnetisation, led to preferential ablation along the pole. 179
- 5.9 Plasma β (left-side) and Magnetic Reynolds Number, R_m , (right-side) at various in-flight times, throughout the $B_{z0} = 25$ T, no-CBET simulation. Contours of the $n_e = n_{cr}/10$ are plotted on all panels as dashed green lines to indicate the bounding region, which contained a significant amount of plasma. Broadly, the β and R_m values are $\gg 1$ in all regions with an appreciable amount of material, which demonstrates that the Lorentz force and resistive diffusion should have minimal effect on the implosion dynamics. 180
- 5.10 In-flight a) and bangtime c) Hall parameters, from the $B_{z0} = 25$ T, no-CBET simulation. Panel b) plots the T_e from the isotropically magnetised simulation (left-side) and anisotropic conduction simulation (right-side) in-flight. Panel d) plots the same, but at bangtime. The electron Hall parameter was > 1 at the poles due to high magnetic fields and temperatures, which led to significantly restricted thermal conduction from magnetised transport. Isotropically magnetised conduction, therefore resulted in a markedly different bangtime morphology, as is shown in panels b) and d). Ion Hall parameters peaked at bangtime, when values reached about $\omega_i \tau_i \sim 0.1$ 182

- 5.11 Panel a) plots in-flight electron density profiles from the $B_{z0} = 25$ T, no-CBET simulations with (left-side) and without (right-side) Nernst-advection of the magnetic field. Panel b) plots magnetic field magnitude from the Nernst (left-side) and no-Nernst (right-side) simulation along with Nernst-advection velocity streamlines from the Nernst-on simulation. Advection of the field is important in the low Hall parameter equatorial region, pulling $\mathbf{B}$ down ∇T_e , into the dense wall. Altered field at the equator impacts on the magnetised thermal conduction, which ultimately imprints on the density, as is seen in panel a). . . 183
- 5.12 Comparison of n_e and T_e profiles from the $B_{z0} = 0$ (panel a) left-side) and 25 T (panel a) right-side), with-CBET simulations, both at $t = 0.5$ ns. Panel a) also plots streamlines of $\mathbf{B}$ for the $B_{z0} = 25$ T simulation, coloured by the field magnitude. Panels b) and c) plot n_e and T_e lineouts respectively, along both the pole and equator. It is evident from these lineouts that magnetisation anisotropically affects hydrodynamic variables, which are used to calculate the CBET gain. Therefore, it is anticipated that magnetisation could anisotropically affect the CBET scattering volume. 185
- 5.13 Various CBET diagnostics used in the analysis presented in this section. All plots are from the $B_{z0} = 0$ T, with-CBET simulation at $t = 0.30$ ns. Panel a) plots the instantaneous deposition with (right-side) and without (left-side) the effect of CBET. The ‘CBET-deficit’, ΔP_{CBET} , which is the no-CBET deposition, subtracted from the with-CBET deposition, is plotted in panel b). Panel c) plots the radially integrated ΔP_{CBET} , as a function of polar angle ($\theta = 0, 180^\circ$ are the poles). The ‘CBET-scattering’, is plotted in panel d), and the radial integral is plotted in panel e). At this time, it is evident from panel e) that more CBET occurred at the capsule poles, resulting in a CBET-reduction of deposition near the poles, seen in panel c). The critical surface is shown in all 2-D panels by a dashed magenta line. 186
- 5.14 The radially integrated CBET-deficit and CBET-scattering, plotted as a function of angle and time for the $B_{z0} = 0$ and 50 T, with-CBET simulations. Panel b) plots the CBET-deficit from the 50 T (top-half) and 0 T (bottom-half) simulations. Lineouts in θ at $t = 0.25$ (black), 0.50 (dark-green) and 0.75 ns (light-green) are plotted in panel a). The same results, but for CBET-scattering are plotted in panels d) and c). It is evident from these plots that for both simulations, as time passed and the poles of the capsule fell in faster than the equator, the region where CBET mostly occurred, shifted in angle around the capsule. Differences are visible in all plots between the $B_{z0} = 0$ and 50 T simulations, indicating that magnetisation affected CBET indirectly, via the altered hydrodynamics. 188

- 5.15 n_e , T_e and $\mathbf{B}$ bangtime profiles for different initial magnetisation (columns) and CBET treatments (rows). Panels a), b) and c) are from the full CBET simulations, which show that magnetisation increased the oblateness of the stagnation profile. Panels d), e) and f) are from the simulations where only the CBET effect on the magnitude, but not spatial location, of deposition was included. These simulations have identical coupled energy to the top row, and therefore have the same bangtimes. Panels g), h) and i) are from the no-CBET simulations, which had earlier bangtimes and increased temperatures, due to the greater coupled energy. 191
- 5.16 Difference in n_e and T_e bangtime profiles, between the full-CBET and partial-CBET a) $B_{z0} = 0$, b) $B_{z0} = 25$ and c) $B_{z0} = 50$ T simulations. The difference in variable ν , is defined as $\Delta\nu = \nu_{\text{full-CB}} - \nu_{\text{partial-CB}}$, so higher colour scale values represent regions with increased ν for the full-CBET calculation. These results primarily show that CBET slightly reduced the bangtime equatorial radius, and thus reduced the oblateness, compared to simulations where the spatial redistribution of power due to CBET was neglected. 192
- 5.17 The oblateness of bangtime density profiles for different magnetisation and CBET effects. All values and errors were obtained by fitting an ellipse (with axes orientated along $\hat{\mathbf{x}}$ and $\hat{\mathbf{z}}$), to the radius of maximum density. No-CBET simulations were consistently more round, because the initial shock was stronger, and therefore travelled ahead of the pusher material more quickly than the CBET equivalent. Thus, after rebounding off the axis, it met the infalling mass and produced thermonuclear conditions at a larger radius. As was seen in Fig. 5.16, throughout the entire implosion, when the effect of CBET on spatial location of deposition was included, it acted to slightly move energy from the pole to the waist and thus marginally reduced the oblateness of the stagnation state. Note that the dashed lines are not fits, they are simply included to guide the eye. . . 193

1 Introduction

1.1 Nuclear Fusion

Nuclear fusion is a reaction which combines multiple light nuclei together to form heavier nuclei. If the products of the reaction are more tightly bound than the reactants, excess energy is also released to the products. Broadly, the binding energy per nucleon of common nuclear isotopes increases with atomic mass up to iron, ^{56}Fe , and decreases afterwards, as shown in Fig. 1.1. The reverse process, nuclear fission, operates by splitting heavy nuclei into lighter products. Therefore, energy is mostly released via fusion, up to product masses including iron, and via fission down to iron. In 2023, fission made up approximately 9.1% of the global electricity mix [3]. Like fission, fusion energy would be carbon-free at the point of production, but it would also offer further distinct advantages. Fusion power plants would not operate in regimes with potentially dangerous chain reactions, and would also generate little-to-no long-lived nuclear waste. In comparison with other low-carbon power sources, including wind and solar, fusion plants could operate continuously, supplying base load power. Studies have also shown that fusion energy could have an economically viable Levelised Cost of Electricity (LCOE), which is a metric that compares the economic costs of a power plant over its lifetime to the value of energy produced [4]. Performing controlled nuclear fusion on Earth for energy production has been an active area of research for many decades. However, multiple scientific and engineering challenges remain to be solved in order to make it a viable energy source.

The likelihood of two reactants undergoing a specific fusion reaction is described by the cross-section of the interaction,

$$\sigma(E) = \frac{S(E)}{E} e^{-E_G/E}, \quad (1.1)$$

which is a function of the centre of mass energy, E , the ‘astrophysical S -factor’, $S(E)$, which is a weakly varying function of energy for many reactions, and the Gamow Energy, E_G [5]. The exponential term including E_G in Eq. 1.1 is related to the probability of reactants tunnelling through the energy barrier, due to electrostatic Coulomb repulsion between two positively charged reactant nuclei. One approach to achieving the energies required to overcome this barrier are ‘beam-target’ configurations, wherein a high energy beam of reactants is focussed onto a stationary target. However, this has proven unviable, due to the high likelihood of Coulomb scattering events compared to fusion reactions [6]. Thermonuclear fusion is the alternative approach, wherein the fuel is heated to sufficient temperatures that the particles in the high-energy tail of the distribution have sufficient energy to undergo fusion reactions. For a fuel in thermal equilibrium, the high particle energies required to overcome the




Figure 1.1: Binding energy per nucleon for common nuclear isotopes. Binding energy peaks close to ^{56}Fe , and energy is released for reactions which increases binding energy of the products compared to the reactants. ^4He has a particularly high binding energy and therefore fusion reactions resulting in ^4He are strong candidates for fusion energy production.

Coulomb barrier, $\mathcal{O}(100)$ keV, are well above ionisation energies (*e.g.* 13.6 eV for hydrogen), so the fuel will be in the plasma state. If the fusion products are able to deposit a sufficient fraction of their energy back into the fuel, then a self-sustaining fusion reaction is possible, where the high temperature required for the reactants to fuse is maintained. For a fusion reaction with reactants labelled by 1 and 2, the number of fusion reactions per unit time and volume is known as the ‘volumetric reaction rate’ [5],

$$R_{12} = \frac{n_1 n_2}{1 + \delta_{12}} \langle \sigma v \rangle, \quad (1.2)$$

where v is the relative velocity of a pair of reactants, δ_{12} is the Kronecker delta, which accounts for double counting of species and the ‘averaged reactivity’ $\langle \sigma v \rangle$ is defined as the integral over the velocity distribution,

$$\langle \sigma v \rangle \equiv \int_0^\infty \sigma(E) v f(v) dv. \quad (1.3)$$

Eq. 1.2 explicitly demonstrates that achieving a high fuel density, can significantly enhance reaction rates, due to the square dependence on number density.

The efficacy of a fusion fuel for energy production is dictated by the availability of the reactants, the fusion products, the averaged reactivity of the reactants and the energy released per reaction, Q , which is the difference in binding energy between the reactants and products. Most current fusion-energy experiments are focussed on demonstrating that fusion power production is possible, thus the choice of fuel is predominantly dictated by the




Figure 1.2: Averaged reactivities for important fusion reactions. The deuterium-tritium reactivity is significantly larger than all other reactions up to $T \sim 500$ keV. Averaged reactivities are obtained from cross-section data, available in the ENDF/B-VII.1 library [7].

reactivity. Hydrogen-hydrogen isotope fusion reactions have much higher reactivities than other elements, because the Coulomb repulsion scales as Z^2 , thus the Gamow energy, E_G is significantly smaller. The nuclear physics is particularly favourable for the fusion of deuterium ($D = {}^2_1\text{H}$) and tritium ($T = {}^3_1\text{H}$). This is due to a nuclear resonance at relatively modest energies for the reaction chain, which produces an excited, unstable ${}^5\text{He}$ nucleus, that subsequently decays to ${}^4\text{He}$ and a neutron [8]. Averaged reactivities of several important fusion reactions, obtained from reaction cross-sections from the ENDF/B-VII.1 library [7], are plotted in Fig. 1.2. The D-T reactivity is at least an order of magnitude larger than the other reactions plotted up to $T \sim 100$ keV, and it is thus the most commonly used fuel for high-gain fusion experiments. The reaction proceeds as,



where α is a ${}^4_2\text{He}$ nucleus, which gains 3.5 MeV from the fusion energy released, Q , and n is a neutron, which gains 14.1 MeV. The energy partition is dictated by energy-momentum conservation of the products in the centre of mass frame. The alpha particles typically couple their energy back into the fuel via Coulomb collisions, which has the effect of raising the fuel temperature and thus further raising the reactivity, at temperatures below $T \sim 60$ keV. Neutrons have a much lower reaction cross-section because they are not charged, and thus typically leave the reaction region. However, in fusion experiments with high density configurations such as Inertial Confinement Fusion (ICF), neutron heating of the fuel can also play a significant role [9].

For net energy gain, sufficient energy must be released by fusion to compensate for the energy required to perform the experiment. This requires a fuel configuration with a high

combination of temperature and density to obtain a large volumetric reaction rate. It must also remain confined for a sufficient time, so enough reactions occur to result in net energy gain. The method of confinement for fusion energy experiments has two broad streams: Magnetic Confinement Fusion (MCF), where magnetic fields are used to confine steady-state fusion-plasma over long time-scales [10], and ICF, where dense fusion fuel is assembled for a short time period and ‘confined’ by its own inertia [11]. The work conducted in this thesis is of relevance to ICF schemes, specifically those in which the plasma is produced by laser irradiation.

1.2 Inertial Confinement Fusion

Large energy releases from nuclear fusion have been achieved on Earth in the form of thermonuclear warheads, although quantity of energy released in these uncontrolled reactions would destroy any infrastructure which would attempt to harness the power for energy production. In the 1950s and 1960s, the first devices were built which demonstrated stimulated emission of microwave [12] and optical [13] light. The optical-wavelength lasers were quickly recognised as an ideal driver for a far smaller and therefore less destructive thermonuclear device than was used in warheads, which could thus be used for fusion energy generation [14]. Much of the research was declassified and subsequently published in an article by Nuckolls *et al.* in 1972 [11]. In these ICF experiments, the high temperatures and densities required to initiate an appreciable number of fusion reactions are maintained over a short timescale, which is set by the inertia of the fuel configuration, before the fuel disassembles due to large pressure gradients. The required fuel density is typically achieved by an implosion process. High temperatures are obtained either from the conversion to internal energy of the implosion kinetic energy, which is the central hotspot ignition variant, described in Sec. 1.2.2, or via an external heating source, as described in Sec. 1.2.3. Before describing these schemes in more detail however, necessary criteria for energy gain conditions shall be discussed, which are agnostic of how the fusion fuel is assembled and dictate the plasma conditions which must be achieved.

1.2.1 Ignition Requirements

The point at which alpha heating becomes the dominant term in the power balance of the fusion fuel is termed ‘ignition’ and it is a necessary, although not sufficient, condition for high gain ICF experiments. Ignition necessitates that the plasma is confined for a sufficiently long period of time, for an appreciable number of fusion reactions to occur. These reactions couple their energy back to the plasma, raising the fuel temperatures and reactivities and thus propagating fusion burn. Estimates for the required plasma conditions which must be achieved for this to occur can be estimated by considering the timescales of confinement and fusion. Initially, a uniform ion number density n_i and temperature T , spherical fuel assembly, with outer radius R , shall be considered. The ions have an average mass, m_i , such that the mass-density of the fuel, $\rho = m_i n_i$ and the volume of the fuel sphere, $V = 4\pi R^3/3$.

Dimensional considerations give an order of magnitude estimate for the timescale on which fusion reactions occur,

$$\tau_{\text{fus}} = \frac{1}{n_i \langle \sigma v \rangle}. \quad (1.5)$$

A radially inward pressure gradient will exist due to the high temperature of the plasma, which will lead to an inward propagating rarefaction wave, disassembling the fuel. This rarefaction wave will move at the isothermal sound speed, $c_s = \sqrt{2k_B T / m_i}$,¹ from $t = 0$, such that the position of the wave is given by $r = R - c_s t$. Therefore, the timescale on which the burning fuel remains un-rarefied and thus confined is,

$$\begin{aligned} \tau_{\text{conf}} &= \int_0^{R/c_s} \frac{(R - c_s t)^3}{R^3} dt, \\ &= \frac{R}{4c_s}. \end{aligned} \quad (1.6)$$

The ratio of these timescales,

$$\frac{\tau_{\text{conf}}}{\tau_{\text{fus}}} = \frac{\langle \sigma v \rangle \rho R}{4m_i c_s}, \quad (1.7)$$

illustrates that, for ignition and hence gain, ICF reactions require a large ρR value, which is typically referred to as the ‘areal-density’ [15].

The timescale ratio can be used to derive the ‘burn-up fraction’, which is defined as the ratio of fusion reactions, $N_{\text{fus}} = R_{\text{DT}} V \tau_{\text{conf}}$, to the number of DT pairs, $N_{\text{DT}} = n_i / 2$, in the fuel assembly,

$$\begin{aligned} \Phi &\equiv \frac{N_{\text{fus}}}{N_{\text{DT}}}, \\ &= \frac{\rho R}{H_B}, \end{aligned} \quad (1.8)$$

where the ‘burn-parameter’, $H_B \equiv \langle \sigma v \rangle / 8m_i c_s$ has been defined, which is simply a function of temperature and should be minimised to achieve the highest Φ possible. For DT fuel, there is a minimum of the burn-up parameter, $H_{B,\text{min}} \sim 7 \text{ g/cm}^2$, which occurs at $T \sim 40 \text{ keV}$ [5]. Low temperatures raise H_B due to low reactivities, and high temperatures increase H_B by increasing the rate of fuel disassembly rates due to increased rarefaction wave speeds. Eq. 1.8 is only valid for $\Phi \ll 1$, because the conversion of DT pairs into non-fusing products has not been considered. A modified form of Eq. 1.8, which approximately accounts for burn-depletion, was derived by Fraley *et al.* in Ref. [15],

$$\Phi = \frac{\rho R}{H_B + \rho R}. \quad (1.9)$$

Using the approximate minimum $H_{B,\text{min}} \sim 7 \text{ g/cm}^2$ in Eq. 1.9, suggests that a burn-up fraction of $\Phi = 30\%$ can be achieved with an areal density, $\rho R = 3 \text{ g/cm}^3$. If solid density DT was used, the mass of fuel required to achieve $\Phi = 30\%$ of the fuel would be 2.5 kg, which would release the same energy as 50 kilotons of TNT [16]. In order to harness the released

¹The factor of 2 in c_s is because the pressure is the sum of contributions from ions and electrons.




Figure 1.3: Four key stages of the central hotspot ignition ICF concept. The chronological order of the diagram is top-left, bottom-left, bottom-right and then top-right quadrants, which display the initial configuration, implosion, stagnation and burn propagation stages respectively.

fusion energy without destroying the surrounding infrastructure, a smaller energy release and therefore mass is required, which necessitates compression of the fuel to achieve the ρR constraint. A pellet containing 1 mg, compressed to densities $\mathcal{O}(10^3)$ times solid density, could release 100 MJ of fusion energy at $\Phi = 30\%$. Spherical compression is optimal because the fuel compresses in three dimensions. This minimises the required convergence ratio, $CR = R_{\text{init}}/R_{\text{final}}$, which increases the tolerance to hydrodynamic instabilities.

1.2.2 Central Hotspot Ignition

Increasing the fuel mass in an ICF experiment requires larger and larger driver energy to heat the entire fuel to the required fusion temperatures, $T \sim 5$ keV, at which an appreciable population of energetic ions are able to overcome the Coulomb repulsion. This necessitates larger and more expensive drivers, limiting both the upfront driver cost and potential target gain. Therefore, experimental configurations which can minimise the driver energy required to ignite a given fuel mass are more practical for current experiments and future power plants. The ‘central hotspot ignition’ method is one such configuration, where a thin shell of fuel containing a low density fuel-fill is approximately isentropically compressed by a carefully constructed driver pulse shape. Isentropic compression maintains a cold and dense, imploding fuel shell, while compressing the gas fill. Only a small mass of fuel is heated to the required fusion temperatures on the stagnation axis. The energy required for this heating is

provided from implosion kinetic energy of the dense shell, which transfers this energy to the gas analogously to a stiff piston compressing a gas. Fig. 1.3 shows a schematic representation of a central hotspot ignition implosion at four important stages. The top left quadrant shows the initial, spherical target configuration, which is comprised of an outer ablator, a solid DT ice layer and an inner DT gas fill. The choice of ablator material is driven by a number of considerations, including mass ablation rate, minimisation of instability growth, radiation coupling and ease of manufacturing [17, 18, 19, 20]. Cryogenic temperatures are required to form the ice layer from the DT fuel. The driver should, ideally, uniformly irradiate the outer surface to maintain the optimal spherical compression and to minimise instability growth. Specific radiation drives shall be discussed in Sec. 1.3, but for now, the scheme shall be discussed, while remaining agnostic of the driver.

The bottom left quadrant of Fig. 1.3 shows the implosion phase of the scheme. As the radiation from the driver heats the outer layer of the ablator, it heats up and expands radially outward, which imparts a rocket force on the interior target, propelling it inward. A gas fill makes compression of the target easier, such that high convergence ratios can be achieved. Greatest convergence can be achieved if the pressure of the gas fill is minimised and thus high-gain ICF targets aim to limit ‘preheat’ of the interior gas fill ahead of the shock. For example, a poorly pulse-shaped driver will lead to undesired shock heating of the interior fuel ahead of the converging shell, which raises its pressure and thus limits compressibility. Assuming perfectly spherical targets and driver radiation, isentropic pulses maximise the target gains, which are designed to limit shock preheating of the gas-fill. In reality, target defects and short wavelength perturbations in the driver radiation seed the Rayleigh–Taylor Instability (RTI) on the ablation surface, which can grow due to the misaligned density and pressure gradients. Realistic pulse shapes often slightly preheat the target, which increases stability to hydrodynamic instability growth [21, 22]. The preheat and thus stability of the design, is typically parameterised by the ‘adiabat’ parameter,

$$\alpha = \frac{P_{\text{shell}}}{P_{\text{F}}}, \quad (1.10)$$

where P_{shell} is the pressure of the DT shell and P_{F} is the Fermi degenerate pressure, *i.e.* truly isentropic compression of a $T \sim 0$ K fuel has $\alpha = 1$. Higher adiabat designs are more stable, and the parameter is set in each shot by launching a weak shock through the capsule before the main pulse.

The stagnation phase of the implosion is shown in the bottom right quadrant of Fig. 1.3. As the remaining relatively cold and dense shell converges on the axis, PdV heating raises the temperature of the compressed interior gas fill. If the implosion velocity is sufficiently high, this piston heating of the fuel results in fusion relevant temperatures. The deceleration phase is also hydrodynamically unstable, because the high pressure, low density interior pushes on the high density, low pressure shell, which results in further RTI growth. Long wavelength perturbations to the radiation drive result in unstagnated kinetic energy in the hotspot, limiting maximum achievable temperatures [23]. Short wavelength perturbations from the RTI can introduce high-Z ablator material to the hotspot, which enhances radia-

tive losses. If the perturbations are sufficiently large amplitude, they can puncture the shell, severely degrading confinement [24].

The upper-right quadrant of Fig. 1.3 shows the final ignition and burn propagation phase of the design. If the density and temperature of the fuel hotspot are sufficient, then fusion reactions begin to occur, generating energetic alpha fusion products. For sufficiently high areal density shells, the alpha particles deposit their energy in the dense fuel shell, raising its temperature and ablating additional fusion fuel radially inward to the burn process. This leads to ignition if the confinement holds the fuel together for a sufficiently long period of time.

1.2.3 Alternative Ignition Approaches

In the central hotspot ignition scheme, the energy required to heat the fuel to fusion temperatures is supplied by the kinetic energy of the imploding shell. The high shell velocity required to obtain hotspot fusion temperatures is ~ 300 km/s, which is typically large enough to result in significant instability growth. Alternative schemes exist, namely the shock- and fast-ignition variants, where the required shell density is obtained via a much slower implosion process, limiting instability growth. Temperatures are then achieved by applying a heating source, separate from the implosion dynamics. For shock ignition, the driver pulse shape is sharply ramped up at the end of the pulse, in order to generate a strong, spherically converging shock, heating the fuel to fusion temperatures [25, 26]. The main issue with this scheme for laser-driven ICF has been uncertainty surrounding laser coupling at the high intensities required to generate the ignitor shock. At high laser intensities, a mostly deleterious class of laser-plasma interaction, known as Laser-Plasma Instabilities (LPIs), dominate. Recent work has explored the possibility of augmenting the laser pulse shape to add a dip in power prior to the ignitor pulse rise, which conditions the plasma such that it is more amenable to strong shock generation. Therefore, the intensities required to generate the shock is reduced [27].

Fast-ignition is an alternative concept, which utilises an external ignitor-beam of charged particles to provide the hotspot heating [28]. Scientific research related to this scheme includes studying how the high energy charged particles are generated via ultra-intense laser interactions and if they can be focused to efficiently couple their energy to the fuel [29, 30]. The overall complexity of many proposed driver and target configurations would also likely have to be reduced in order to make Inertial Fusion Energy (IFE) relevant targets.

All of the ignition concepts introduced so far focus on igniting a small central mass, which then initiates a propagating burn wave into surrounding fuel. By contrast, the volume-ignition variant proposes igniting the entire fuel volume simultaneously [31, 32]. This is achieved by assembling a large, dense fuel mass, kept at a much lower temperature than the hotspot in other schemes. A large and dense fuel assembly has minimal radiative losses and can thus ignite at $T \sim 1.5$ keV, as opposed to the ideal ignition temperature of hotspot schemes, $T = 4.3$ keV [5]². One issue with this scheme is that significantly more laser energy

²The ideal ignition temperature is the value at which alpha heating balances radiative losses.




Figure 1.4: Energy balance of an IFE power plant. Based on a similar figure from Ref. [5].

is required to ignite a given fuel mass, because the entire volume must be brought to ignition conditions at the same time. The high driver energies required to ignite the fuel also limit the gain to $G \sim 20$, and it is thus not particularly relevant to IFE [5].

1.2.4 Inertial Fusion Energy Considerations

For a power plant to produce net energy from ICF implosions, the energy from each target must naturally be greater than the energy to drive the implosion. In reality, additional inefficiencies in power plants set more stringent constraints on the energy which must be produced. The energy balance of an IFE power plant is shown in Fig. 1.4. A driver, such as a laser system (which is the driver considered for the work conducted in this thesis), converts an input energy, E_{in} , to a driver energy, $E_d = \eta_d E_{in}$, where η_d is the energy efficiency of the driver. The driver energy initiates a fusion reaction, which releases E_{fus} , with a gain defined by,

$$G = E_{fus} / E_d. \quad (1.11)$$

The released fusion energy, for example in the form of energetic neutrons, is converted into thermal energy in an encompassing ‘blanket’, and then into output electrical energy, E_{out} , by a thermal-cycle (such as steam turbines), with efficiency, η_{th} . A fraction, f , of this generated energy is recycled back into the plant to power the driver, such that the remaining fraction is sent to the grid,

$$E_{grid} = (1 - f) \eta_{th} \eta_d G E_{in}. \quad (1.12)$$

In order for the reaction to be self-sustaining, f must be set such that E_{in} is sufficient to initiate another reaction. This leads to the energy balance for the cycle, $f \eta_d \eta_{th} G = 1$. By definition, $0 \leq f \leq 1$, so if $\eta_d \eta_{th} G \leq 1$, then net energy production is not possible. Power plants must also produce a sufficiently large volume of energy to balance the capital cost required to build them. Therefore, several ~ 100 MJ reactions must occur each second for a 100 MW plant, which necessitates a driver that can operate at ~ 10 Hz.

The Diode-Pumped Solid State Laser (DPSSL) design could feasibly produce laser energy at ~ 10 Hz repetition rate, with an efficiency $\eta_d \sim 10\%$ [33, 34]. Assuming a thermal cycle




Figure 1.5: Schematic of the indirect-drive approach to ICF. Laser light, represented as blue transparent rectangles, irradiates the interior of a high-Z (*e.g.* gold) hohlraum, producing thermal x-rays that drive the capsule implosion. A number of important physical effects, typical to indirect-drive experiments, are also illustrated.

efficiency, $\eta_{\text{th}} = 40\%$ and a recycled energy fraction, $f = 1/4$, a target gain of approximately $G \sim 100$ would be required for power production. Additionally, the plant must be economically profitable, thus for a $G = 100$ reaction, which releases $E_{\text{fus}} = 100 \text{ MJ} = 28 \text{ kWh}$, the energy sold to the grid would make about $\text{£}1.50^3$. Therefore, targets must cost approximately $\text{£}0.10$, which limits the tolerable manufacturing complexity. IFE target concepts exist, which could feasibly be mass-produced for an acceptable cost, for example initially uniform density liquid spheres, known as ‘dynamic shell’ targets [35, 36].

1.3 Current Experiments and Main Approaches

In Sec. 1.2, the basic principles of ICF were introduced without specifying the driver technology. Although heavy ion beams have been proposed as a driver technology [37], most current research focuses on the direct-drive and indirect-drive approaches, which use laser light to directly and indirectly irradiate the target, respectively. For the direct-drive approach, lasers are focussed upon the outer surface of the target, whereas in indirect-drive, they illuminate the interior surface of a high-Z material ‘hohlraum’, which generates a thermal bath of x-rays. The indirect-drive approach was developed in order to relax requirements on laser beam uniformity and sensitivity to hydrodynamic instabilities in direct-drive [38]. Sources of degradation and implosion dynamics are somewhat different in each of these approaches. Each approach is introduced below, and progress on the main experimental facilities is summarised.

1.3.1 Indirect-Drive

1.3.1.1 Indirect-Drive Physics

A schematic of an indirect-drive experiment is shown in Fig. 1.5. The National Ignition Facility (NIF) at the Lawrence Livermore National Laboratory (LLNL) in the United States, is the

³The UK April 2024 energy price of $\text{£}0.25$ per kWh was used.

largest ICF facility and high-power laser system in the world, and it mainly focuses on the indirect-drive approach to ICF [39, 40]. The laser system is composed of 192 beams, which are clustered around two poles, so that they can enter the small Laser Entrance Hole (LEH) of the hohlraum, as is demonstrated in Fig. 1.5. Approximately 2 MJ of laser energy can be delivered to the target, usually on a timescale of $10 \rightarrow 20$ ns. The Laser Mégajoule (LMJ) is a newer facility at Commissariat à l'Énergie Atomique et aux Energies Alternatives (CEA) in France, which is the largest ICF experiment outside the United States. It also predominantly focuses on the indirect-drive approach and, although currently still in construction, it is intended to have 176 beams that will deliver approximately 1 MJ of laser energy to the hohlraum [41].

The main benefit of the indirect-drive approach compared to direct-drive is the uniformity of the driving radiation⁴. High-power laser focal spots are instantaneously highly non-uniform, which can seed instabilities, and is discussed in more detail in Sec. 1.3.2. Sphericity of compression is also limited by the number of available beams. Conversion of laser light to x-rays in a hohlraum creates a highly uniform radiation field, which has limited short wavelength perturbations to seed RTI growth. Instability growth is therefore mostly seeded by engineering features, such as the tent used to hold the capsule in place, defects in the capsule, surface roughness, and the fill tube used to insert the fuel [42]. A high-Z material such as gold is used for the hohlraum, due to high laser absorption and emissivity. Longer wavelength modes can be seeded if the laser heating of the hohlraum is asymmetric. For example, Fig. 1.5 shows that the hohlraum wall can heat and expand, blocking the 'inner-cone' beams which heat the material near the capsule equator. Filling the hohlraum with a low-Z gas fill⁵ to limit the 'gold bubble' expansion, however, this allows for deleterious LPIs to occur, which reflect laser energy back out of the LEH [43]. Current experiments on the NIF use a relatively low gas fill to reduce backscatter-LPIs (specifically Stimulated Brillouin Scattering (SBS) and Stimulated Raman Scattering (SRS)). They also introduce a wavelength shift between the inner and outer cones of laser beams, which allows for an instability known as Cross-Beam Energy Transfer (CBET) to transfer power to the outer beams late in the implosion [44, 45, 46]. This compensates for absorption in the expanding gold bubble and can maintain symmetry of the implosion. The main drawback of the approach, compared to the direct-drive, is that the conversion of laser light to capsule kinetic energy via the hohlraum introduces an additional $\sim 10\%$ efficiency decrease, limiting the maximum gains that can be achieved. Increasing the size of the capsule relative to the hohlraum increases the fraction of x-ray energy which is coupled to the target, but can introduce long-wavelength instabilities if made too large by blocking laser propagation. Direct-drive is thus often considered the preferred scheme for IFE.

⁴Although higher frequency radiation also leads to higher ablation pressures for a given power absorption, the higher coupling efficiency of direct-drive means that a given laser power has a higher ablation pressure for direct-drive than indirect-drive.

⁵A low-Z hohlraum fill is used to limit laser absorption away from the hohlraum wall.




Figure 1.6: G plotted against time for indirect-drive ICF shots on NIF. The colour of each bar represents a different implosion design, and the dashed horizontal line represents $G = 1$, where the fusion energy produced is equal to the incident laser energy. Used under CC BY 4.0 from Ref. [47].

1.3.1.2 Indirect-Drive Experimental Progress

Although early experiments on the NIF dramatically underachieved the main aim of demonstrating ignition, a tremendous amount of progress has been made in recent years [48]. Fig. 1.6 plots G for indirect-drive ICF shots on the NIF, up to December 2022 [47]. The initial target designs, up to 2013, utilised a low adiabat design, with a plastic (CH) ablator, known as 'Low Foot', which proved very sensitive to hydrodynamic instabilities [49]. The results of this campaign are plotted in shaded grey in Fig. 1.6. This led to the subsequent, 'High Foot' campaign, where the target was kept similar, but the adiabat of the design was raised to create a more stable implosion, which improved yields, plotted in red [22].

The 'Hybrid' campaigns which came after these experiments made several changes to improve performance. The gas fill-density of the hohlraum was lowered to reduce the loss of laser energy to backscatter LPIs. To limit large growth-times of gold bubbles while the laser fired and thus a loss of drive symmetry the laser pulse duration had to be shortened. This necessitated deployment of a new ablator material, High Density Carbon (HDC), which has a higher density than CH, and therefore the shock transit times and overall pulse lengths were reduced [50]. CBET was also employed to compensate for the loss of symmetry re-




Figure 1.7: Schematic of the direct-drive approach to ICF. Laser light, represented by the transparent-blue rectangles is absorbed outside the critical surface and then transported in to the ablation surface by thermal conduction. Speckle from beam smoothing optics leads to short-scale non-uniformity of deposition, known as imprint. LPs degrade the performance, both by reflecting light, and by generating hot electrons, which pre-heat the fuel.

sulting from the gold-bubble expansion in the low-fill hohlraum, larger capsules increased coupling efficiency and thicker shells mitigated instabilities [51]. Additional improvements were made to capsule quality, which minimised sources of instability [52]. Shots from these campaigns are shown in the light and darker blue in Fig. 1.6 and have led to experiments entering the burning-plasma⁶ [53, 54] and ignition regimes [55]. This ultimately resulted a shot $G > 1$ in December 2022, which was the first fusion experiment to produce more energy than the driver delivered [47]. Although not plotted in Fig. 1.6, the gain record at the time of writing stands at $G = 2.4$ [56].

1.3.2 Direct-Drive

1.3.2.1 Direct-Drive Physics

A schematic of a direct-drive experiment is shown in Fig. 1.7. Although a limited number of direct-drive implosions are performed on the NIF, the largest dedicated direct-drive facility in the world is the OMEGA laser facility at the Laboratory for Laser Energetics (LLE) in the United States [57, 58]. Work presented in this thesis is most relevant to direct-drive implo-

⁶The burning-plasma regime is when alpha heating is larger than PdV heating. This is opposed to ignition, where alpha heating is greater than loss terms.

sions on this facility. The OMEGA laser has a total of 60 beams, arranged to provide approximately symmetric radiation to a spherical target. In total, beams can deliver approximately 30 kJ of laser energy to the target, usually in a timescale of $2 \rightarrow 3$ ns. As previously stated, high driver efficiencies are desirable for IFE, and the relative inefficiency of the hohlraum light to x-ray conversion means that direct-drive is the assumed method for a future IFE power plant. As opposed to indirectly-driven NIF experiments, where the capsule is held in place by a thin tent attached to the hohlraum, directly-driven OMEGA targets are glued to a mounting stalk to hold them in place. This can lead to low-mode asymmetries and unstagnated flows in the hotspot, which degrade the performance [59].

In plasmas, electro-magnetic radiation is absorbed beyond the critical surface, interior to which the electron density is sufficiently high, such that light of a given wavelength cannot propagate. For indirect-drive, high frequency x-rays are absorbed close to the ablation surface, whereas for direct-drive, longer wavelength laser light has a critical surface which is exterior to the ablation surface. Thus, energy is absorbed beyond the critical surface, heating this ‘coronal plasma’⁷ to high temperatures. The absorbed laser energy is then transported radially inward via thermal conduction in the ‘conduction-zone’, to the higher density, lower temperature ablation surface, as is shown in Fig. 1.7. Non-uniform absorption at the critical surface is partially smoothed by non-radial thermal conduction in this region, which is more efficient for shorter wavelength perturbations [60].

Despite this non-radial smoothing in the conduction zone, non-uniformity of the laser drive has proven to be a significant difficulty for direct-drive ICF implosions. The beam layout and their widths lead to a ‘beam-mode’ asymmetry, which for OMEGA has a distinctive Legendre-mode, $\ell = 10$ pattern, which can be seen in Fig. 3.26. Focal spots of high-power lasers are also often highly non-uniform when so-called ‘beam-smoothing’ techniques are not employed. As light propagates from the beam port to the far-field, random ‘phase-aberrations’ are picked up by the beam due to distortions in the medium through which it travels. These then result in large scale modulations of the focal spot intensity profile. When the target is illuminated by these modulated beams, significant asymmetry is imparted on the drive, severely limiting the overall performance. A variety of beam smoothing techniques are thus employed which lead to a significantly more uniform beam profile. Fig. 1.8 shows a schematic of several important beam smoothing techniques employed by the OMEGA laser system. The diagram demonstrates that without beam smoothing, a uniform near-field intensity profile picks up phase aberrations along the laser chain, leading to interference and thus large scale modulations at focus. Phase plates are an array of small optical elements which offset the phase of the beam at discrete points in the near-field, acting to shift the scale of modulation in the far-field to a much smaller spatial scale [61]. These short-scale modulations, known as ‘laser-speckle’, are much more easily smoothed in the conduction zone than the large modulations from beams without smoothing techniques applied. The static phase plate far-field profile can be further improved by using SSD, which creates a temporally varying speckle pattern and thus the time-integrated profile is much smoother. Fig. 1.8

⁷The term coronal plasma shall be used throughout this thesis to refer to the approximately isothermal, outward flowing, low density plasma blow-off, which is exterior to the critical surface.




Figure 1.8: Schematic of beam smoothing techniques employed on the OMEGA laser system. Phase aberrations can lead to large scale modulations to the far-field intensity profile. Phase plates shift the perturbations to shorter scales, which are more efficiently smoother by thermal conduction. Smoothing by Spectral Dispersion (SSD) creates a time-varying speckle pattern, which leads to a smooth, *time-integrated* far-field profile. Polarisation smoothing splits the beam into orthogonally polarised sub-beams, which do not interfere and thus further reduce non-uniformity.

compactly displays how SSD operates, which is to apply a temporally varying bandwidth to the light signal, then pass it through a diffraction grating, such that different wavelengths are dispersed in different directions [62, 63]. Finally, polarisation smoothing is employed on OMEGA, which works by splitting each beam into two slightly spatially-offset, orthogonally polarised sub-beams. These sub-beams do not interfere with each other, and thus when added together the root-mean-square intensity variation of the whole beam is improved by a factor $\sqrt{2}$ [64, 65].

Although the combination of these optics creates an intensity spot which is smooth when time-integrated, instantaneously large amplitude, short wavelength perturbations exist. When an extended plasma corona has formed, these perturbations are effectively smoothed by non-radial thermal conduction in the conduction zone, but early in the implosion, the small amplitude perturbations are deposited close to the ablation surface. This is known as ‘laser imprint’ and unless mitigated by sufficiently raising the adiabat of the implosion, can break up the shell. The adiabat is typically set for direct-drive implosion by including a small ‘picket’ pulse in the laser temporal profile, which weakly shocks the shell. The laser turns off after this picket pulse, leading to an unsupported shock, which decays in amplitude as it travels through the shell, such that only the outer extent of the shell, where the imprint occurs, is on a high adiabat. Imprint does not occur in an indirectly-driven implosion, and thus the highest performing direct-drive experiments normally have much higher adiabats.

LPIs also act to significantly degrade direct-drive ICF implosions. A schematic diagram




Figure 1.9: Important laser-plasma interactions for direct-drive ICF. Energy is absorbed by Inverse-Bremsstrahlung (Inv-Brem) close to critical, and resonance absorption at the critical surface. Hot electrons pre-heat the fuel, reducing compressibility and are generated from an Electron Plasma Wave (EPW) from Two Plasmon Decay (TPD) or SRS. Energy is reflected by CBET, SRS and SBS. The filamentation instability can cause light to self-focus and thus enhance asymmetry of the deposition. It also raises intensity of the light locally, enhancing the growth rate of other LPIs.

of LPIs relevant to direct-drive is shown in Fig. 1.9. Energy is absorbed either via Inv-Brem, which is the absorption of laser energy by an electron while it undergoes a collision with an ion, or by resonance absorption, which is the excitation of a resonant EPW at the critical surface. Inv-Brem is the dominant absorption mechanism for OMEGA experiments, because frequency-tripled⁸ light is employed for the lasers, and the Inv-Brem rate scales with frequency of the light [66]. Similarly to indirect-drive, SBS and SRS act to reflect laser light away from the capsule. However, unlike for indirect-drive, where CBET is used to tune the symmetry of the implosion, in direct-drive configurations CBET acts to reflect significant amounts of energy. For direct-drive, the dominant CBET interaction occurs as light from the edge of beams reflects outward and gains energy from inward travelling beams. This mechanism is responsible for a $\sim 20\%$ reduction in energy for typical OMEGA implosions over the whole implosion, and can instantaneously reduce power deposition by up to $\sim 50\%$ [67]. Additionally, LPIs which excite an EPW, specifically TPD and SRS, can trap and accelerate a population of electrons to high energies ($\gtrsim 50$ keV), which travel through the ablator material and preheat the fuel, reducing compressibility and thus performance. Finally, filamentation of the beam can occur, where the field self-focuses through low-density channels in the plasma, which increases laser non-uniformities [68].

⁸The OMEGA laser-beams are initially generated at a vacuum wavelength, $\lambda_0 = 1064$ nm, by a Nd:YAG crystal and then passed through a non-linear optic, which converts the wavelength to $\lambda_0^{3\omega} = 351$ nm [58].

On the OMEGA laser facility, direct-drive implosion designs typically operate with peak intensities that are kept below $I \sim 10^{19} \text{ Wm}^{-2}$. This is primarily in order to limit the growth rate of LPIs, which increase in significance at greater intensities. Considering just the energy reflection caused by LPIs and not hot electron generation, CBET is the dominant instability on OMEGA. Future direct-drive facilities with greater laser energy would lead to longer plasma scale-lengths. In these long scale-length plasmas additional instabilities, namely SRS, are observed to become more significant [69]. Additionally, laser filamentation is not a significant concern for current OMEGA direct-drive designs due to the limited laser intensity and beam smoothing techniques employed. Smoothing techniques act to limit the available seeds for filament growth. The work in this thesis related to LPIs aimed to improve modelling of current-scale direct-drive designs, primarily on the OMEGA laser facility. Model development therefore focussed on accurately modelling the dominant mechanism for energy reflection, CBET.

1.3.2.2 Direct-Drive Experimental Progress

In a central hotspot ICF implosion, fusion reactions are made possible by a sufficiently hot and dense plasma hotspot, which is confined by the inertia of the target. The energy required to assemble this hot and dense configuration occurs through a series of energy transfers:

- Laser energy is converted to internal energy by absorption in the coronal plasma.
- This energy is coupled to kinetic energy in the shell by thermal conduction in the conduction zone.
- Kinetic energy of the shell stagnates on the axis, resulting in internal energy of the hotspot.
- This internal energy allows fuel reactants to overcome the Coulomb barrier and fuse, releasing binding energy as product kinetic energy.

The overall performance of the implosion scales with the internal energy of the hotspot and can be improved, either by increasing the available driver energy, or by increasing the efficiency of these energy conversions. The OMEGA laser facility has insufficient energy to achieve ignition conditions, thus inferring the performance of equivalent implosions on a larger laser direct-drive facility relies on techniques known as hydrodynamic scaling. A combination of analytical methods and simulations is typically used to extrapolate the performance of OMEGA implosions to NIF energies, $E_d \sim 2 \text{ MJ}$, in order to compare performance of indirect- and direct-drive experiments [70].

Better hydrodynamic stability of the implosion improves the integrity of the imploding shell. In analogy to a stiff piston compressing a gas, a denser and more uniform shell efficiently converts kinetic energy to internal energy of the hotspot gas, whereas a shell with short wavelength perturbations is effectively less dense and more compressible, decreasing conversion efficiency [71]. Over the past decade, implosion performance improved markedly

by using statistical modelling techniques [72], to identify and reduce sources of yield degradation [73]. Primarily by focusing on eliminating these short wavelength instabilities using the statistical model, fusion yield was tripled in a short period of time from 5×10^{13} in 2016, to 1.5×10^{14} in 2018 [74]⁹. The details of the statistical model used are discussed in Sec. 4.1.2.

More recently, focus has also turned to increasing the efficiency of the energy conversion through other means. Specifically, a significant fraction of laser energy remains unabsorbed, mostly due to CBET, which can reduce power deposition instantaneously by $\sim 50\%$. CBET has been reduced in recent campaigns via two means. Firstly, ‘DT-liners’ are an experimental configuration which have increased the outer diameter of the target, reducing the beam overlap and thus CBET, which leads to high implosion velocities [75]. This campaign recently reported for the first time that more fusion energy was produced, than was coupled to the hotspot. [76]. Secondly, a small amount ($\sim 5\%$ by atomic number density) of silicon is now routinely added to the outer CH ablator. This enhances the Inv-Brem absorption of light, reducing the quantity of light that refracts away from the target, to act as a seed for CBET. Additionally, addition of a higher-Z dopant into the ablator, reduces growth rates of many LPs, limiting the hot electron production. A campaign which optimised the addition of silicon dopant recently demonstrated a hydrodynamically equivalent burning plasma on OMEGA, which is an important milestone on the path to hydrodynamically scaled ignition of a direct-drive target [77].

1.4 Objective of the work

Important physical processes in ICF implosions occur on short time and length-scales, and have high temperatures and densities. Inference of plasma conditions from escaping particles or radiation is thus typically used to understand the implosion dynamics. However, these methods are limited to measuring emissive volumes at certain times throughout the experiments. Simulations are an important tool, and are used to study ICF experiments at greater spatial and temporal resolution than is allowed for by many diagnostic tools. They can also be used to, for example, extrapolate the performance of experiments to larger facilities, or to help design new experiments [54]. The applicability of simulations does of course depend upon the validity of the models used. Hydrodynamic codes are the workhorse tool, used to simulate entire ICF implosions. The models employed in these codes must therefore be capable of accurately including relevant physical effects on long ($\sim \text{ns}$) timescales, which precludes many high fidelity, although expensive tools.

The aim of the work conducted in this thesis was to improve the modelling of the laser-plasma interactions in the 3-D Magnetohydrodynamics (MHD) code CHIMERA, particularly with a view to more accurately simulating direct-drive ICF experiments. Prior to the work conducted in this thesis, a simple ray-tracing algorithm existed in CHIMERA to model the laser-plasma interaction, although it did not accurately account for refraction of the light. Therefore, a simplified 1-D ray-trace was typically used for direct-drive calculations, where

⁹Improved instability robustness was achieved in these implosions, mostly by limiting the convergence ratio of the implosions.

rays travelled radially inward and did not refract away from the target. To improve predictive capability, it was deemed necessary to include a 3-D ray-trace, which accurately modelled the trajectory of laser light through the plasma and accounted for refractive coupling losses.

The discussion of indirect-drive in Sec. 1.3.1 and direct-drive in Sec. 1.3.2 hopefully illustrated that LPIs are a significant concern on current ICF experiments. Particularly CBET in direct-drive is highly energetically significant, reducing the efficiency of laser absorption by $\sim 20\%$ over the whole experiment. Additionally, CBET significantly alters deposition asymmetries [67], and therefore it is crucial to accurately model the interaction, in order to understand the degradation of performance from multi-dimensional effects. Despite this, few multidimensional CBET models exist, suitable for direct-drive calculations, and only a single code, IFRIIT, had been integrated to run inline with a 3-D hydrodynamics code [67]. The second aim of the work was therefore to develop a CBET model for the ray-trace which could be employed in direct-drive simulations and run inline with the hydrodynamics.

The remainder of this manuscript is organised as followed:

- Chapter 2 introduces background theoretical material, relevant to the interaction of lasers with plasma in the regime of ICF, with a particular focus on the techniques used in the computational models employed in this thesis. Equations of the Radiative-Hydrodynamics (Rad-Hydro) framework are presented, alongside a discussion of their implementation in the CHIMERA code. A summary of the theory of the interaction of laser light with plasma is then presented. The ray-tracing equations are introduced and their validity domain is discussed. Parametric instabilities, also known as LPIs, are finally discussed, with a particular emphasis on CBET.
- Chapter 3 presents the 3-D laser ray-trace and CBET module, SOLAS, which was developed for the CHIMERA code. Initially, the development of the 3-D ray-trace is presented, which includes details on the initialisation of rays, the grid and the numerical methods. The field reconstruction algorithm is then detailed, which is necessary to obtain the electric field used for calculating the CBET gain, before the CBET algorithm is introduced. Validation problems are presented throughout, which compare SOLAS to both analytic solutions and higher-fidelity solvers. The chapter concludes with an in-line CBET calculation for a 1-D Rad-Hydro simulation of an OMEGA implosion and a 3-D post-process of CBET induced non-uniformity of deposition.
- Chapter 4 presents the results of series of CHIMERA-SOLAS simulations, which were performed to study the effect of the beam-to-target radius ratio of the uniformity of direct-drive implosions. The simulation configuration used for the study is detailed alongside a set of 1-D tuning simulations to obtain the initial conditions for the 2-D simulations. Trends of stagnation state symmetry both in the presence and absence of CBET are presented and the time-resolved growth of the beam-mode asymmetry is studied.
- Chapter 5 gives results from a set 2-D CHIMERA-SOLAS simulations, which were performed to understand the indirect effect of magnetisation on CBET in direct-drive im-

plosions. Both magnetised and unmagnetised simulations are presented, with and without the effects of CBET on laser deposition. The main results presented are the asymmetries of the stagnation state, for various levels of magnetisation and treatments of CBET.

- Chapter 6 summarises the work conducted in this manuscript and outlines plans for future work, such as potential extensions of the model.

2 The Interaction of Light with Plasma

This chapter shall introduce the theoretical background relevant to the work conducted in this thesis. The main focus of the work is centred on improving the laser modelling in the CHIMERA Radiative-Magnetohydrodynamics (Rad-MHD) code. Therefore, the theoretical framework for modelling both the plasma and the propagation and absorption of light in it is introduced.

To begin, the plasma state is defined and important length- and time-scales are provided in Sec. 2.1. The kinetic and fluid descriptions of plasma are introduced in Sec. 2.2 and the validity domain of each framework is discussed, with particular reference to typical conditions for laser-plasma interactions. Additional physics packages of the fluid code, CHIMERA, which are utilised in later chapters, are introduced. A model for kinetic heatflow is desirable in fluid codes which model laser-plasma interactions. Therefore, although CHIMERA does not have this capability, some basic theory of kinetic heatflow modelling is highlighted.

An additional aim of the work was to include a Cross-Beam Energy Transfer (CBET) model into the new CHIMERA laser package. Laser-Plasma Instabilities (LPIs) such as CBET are multi-wave coupling phenomena, and therefore a basic description of waves in plasma is provided in Sec. 2.3. The dispersion relation of the light waves in plasma is also derived. In Sec. 2.4, beginning with the full wave equation and then introducing successive assumptions which are broadly satisfied in typical laser-produced Inertial Confinement Fusion (ICF) plasmas, the equations of ray-tracing are derived. Important absorption processes are then outlined, particularly Inverse-Bremsstrahlung (Inv-Brem), which is the dominant mechanism on the largest ICF facilities in the world today. Finally, the basic theory of LPIs is provided in Sec. 2.5, particularly CBET and its relevance for direct-drive ICF.

2.1 Basic Plasma Physics

As stated in Chap. 1, thermonuclear fusion requires fuel to exist at significant temperatures, which are well above ionisation energies. Therefore, the fuel configuration in these fusion experiments is a plasma. Formally, a plasma is defined as a quasi-neutral, ionised gas, which exhibits collective behaviour [78]. The charged particles within a plasma interact via the long-range Coulomb force, and thus undergo many simultaneous interactions with the other particles. This leads to a variety of collective phenomena, such as the plasma waves described in Sec. 2.3. Quasineutrality of the plasma means that, when observed at a length-scale L , the plasma has no net charge,

$$\sum_{\alpha} q_{\alpha} N_{\alpha} = 0, \quad (2.1)$$

where N_{α} is the number of particles of species α , with charge q_{α} , in the cube with volume $V = L^3$. For a ‘single-species’ plasma¹ with average ionisation state Z , this implies,

$$n_e = Z n_i, \quad (2.2)$$

where n_e and n_i are the number densities of electrons and ions respectively.

Quasineutrality arises because the particles in the plasma are free to move due to forces they experience. Thus, if a local charge imbalance occurs then the electrons, which respond faster than the ion population due to their lower mass, move to rebalance this field and restore quasineutrality. This electron relocation to eliminate local electric fields is known as Debye-screening. The length scale below the electron population cannot effectively screen the charges sets the length scale of quasineutrality, and is known as the Debye-length,

$$\lambda_D^2 = \frac{\epsilon_0 k_B T_e}{n_e e^2}, \quad (2.3)$$

where T_e is the electron temperature, ϵ_0 is the permittivity of free space, k_B is the Boltzmann constant and e is the electron charge. This screening can only occur if there is a large number of particles in the Debye sphere, $n_e \lambda_D^3 \gg 1$.

The timescale of the charge relocation is of particular importance to the interaction of light with plasmas. Light waves are an oscillating electric and magnetic field, which the charged particles in the plasma can respond to. If the particles are able to respond quickly enough, they can influence the propagation of the light and ultimately force the plasma to become opaque to the propagating light wave. The oscillation timescale can be derived by considering a uniform assembly of quasineutral plasma and then displacing the electron population from the ions by a small distance δx along the x -axis. The electric field which develops within the plasma is directed along the x axis with a magnitude found from Gauss’ law,

$$E_x = \frac{n_e e \delta x}{\epsilon_0}. \quad (2.4)$$

¹Single-species here means that there is only a single type of ion.

This field leads to a restoring force $F_x = -eE_x$ on the charged constituents of the plasma. Solving Newton's second Law demonstrates that, when the thermal motion of the electron population is ignored,² oscillations of the electrons occur at the 'plasma frequency',

$$\omega_p^2 = \frac{n_e e^2}{m_e \epsilon_0}. \quad (2.5)$$

If the forcing oscillations occur at a frequency which is lower than ω_p , then the electrons move rapidly enough to nullify the field. Consider light with frequency ω , which is incident normal to a plasma density gradient. Because $\omega_p \propto n_e$, the light is able to propagate until it reaches the density where $\omega = \omega_p$, which is known as the critical density,

$$n_{\text{cr}} = \frac{m_e \epsilon_0 \omega^2}{e^2}. \quad (2.6)$$

After reaching the critical density, the field of the light decays exponentially as an evanescent wave, but cannot propagate.

2.2 The Kinetic and Fluid Formulations of Plasma

Idealised computational modelling of a plasma would solve the long-range electro-magnetic interaction between every pair of particles at all times. However, this approach very rapidly becomes intractable due to the large number N of particles and the $\mathcal{O}(N^2)$ scaling of interactions to solve³. Reduced frameworks must thus be devised to predict the evolution of the plasma. Plasmas are divided into two broad classifications. When in local thermal equilibrium, the plasma is often described as 'thermal' and the fluid formulation is an adequate description. There are numerous situations when this is not true however, for instance if a particular subset of particles is heated at a rate, which is much faster than thermalising collision timescales. In this case, the subset of the system is described as 'non-thermal' or 'kinetic'. The fluid formulation becomes an inadequate description and higher fidelity, kinetic tools must be used to describe the evolution of the system.

When a kinetic description of a plasma is required, the distribution function, $f_\alpha(\mathbf{x}, \mathbf{v}, t)$ is used to describe the state of particle species α . It provides a statistical description of the number density of particles, which inhabit a phase-space, $[\mathbf{x}, \mathbf{v}]$, at time t . If the system macroscopically evolves on a timescale much longer than the collision time and length scales are longer than the collisional mean free-path, then these collisions between particles relax the distribution function toward a Maxwellian,

$$f_{\alpha, \text{Max.}}(v) = n_\alpha \left(\frac{1}{2\pi v_{T\alpha}^2} \right)^{1.5} e^{-[v/(\sqrt{2}v_{T\alpha})]^2}, \quad (2.7)$$

²Inclusion of thermal motion leads to pressure, which acts as a restoring force, and yields the dispersion relation for an Electron Plasma Wave (EPW).

³Note that it is possible to perform N -body plasma simulations by utilising, for example tree-code techniques [79], although this is not a common approach.

where $v_{T\alpha} = \sqrt{k_B T_\alpha / m_\alpha}$ is the thermal speed of the species with temperature and mass T_α and m_α , respectively. The fraction outside the exponential is set such that the integral over velocity space yields the number density of the species. Typically, for bulk of the target configuration throughout a laser-produced ICF implosion, the assumption of a Maxwellian distribution is close to accurate, although there are several notable exceptions. For instance, DT fusion products have energies much higher than thermal energies and are also monoenergetic. LPIs also generate energetic electron populations, which are able to range through the implosion due to their low collisionality [80]. Additionally, laser-heated plasmas typically exhibit steep density and temperature gradients near the ablation surface. Therefore, collisions and related phenomena like the transport of thermal energy, do not act locally [81]. When performing fluid simulations of ICF experiments, an accurate description of these kinetic phenomena requires specific non-local modelling techniques.

2.2.1 The Kinetic Formulation

The evolution of the distribution function for each individual species is described by the Vlasov-Fokker-Planck equation,

$$\frac{\partial f_\alpha}{\partial t} + \mathbf{v} \cdot \nabla_x f_\alpha + \frac{q_\alpha}{m_\alpha} (\mathbf{E} + \mathbf{v} \times \mathbf{B}) \cdot \nabla_v f_\alpha = \left(\frac{\partial f_\alpha}{\partial t} \right)_{\text{collisions}} + \left(\frac{\partial f_\alpha}{\partial t} \right)_{\text{external}}, \quad (2.8)$$

where q_α is the charge of the species, $\mathbf{E}$ and $\mathbf{B}$ are the macroscopic electric and magnetic fields,⁴ respectively, and ∇_x and ∇_v are gradients with respect to position and velocity coordinates, respectively [82]. The equation describes the conservation of number of particles in phase-space. The Vlasov equation, which is obtained from 2.8 by setting the collisional term on the right-hand side to zero, describes collisionless systems. In these collisionless systems, forces from macroscopic fields dominate over short scale length forces, *i.e.* collisions between particles. When collisional timescales are sufficiently small that collisions influence the evolution of the distribution function, then the collisional term cannot be neglected and the full Vlasov-Fokker-Planck equation must be solved. The collision operator on the right-hand side of Eq. 2.8 describes the action of microscopic fields, which arise due to the random motion of the charged plasma particles. Various collision operators can be employed, but the Fokker-Planck collision operator takes the form,

$$\left(\frac{\partial f_\alpha}{\partial t} \right)_{\text{collisions}} = -\nabla_v \cdot (\mathbf{A} f_\alpha) + \nabla_v \cdot (\mathbf{D} \cdot \nabla_v f_\alpha), \quad (2.9)$$

where the first term on the right-hand side represents drag or slowing due to collisions and the second term represents diffusion of particles in velocity space due to small-angle collisions. External forces, for example gravity, are also accounted for in the second right-hand side term.

⁴Macroscopic here means on a scale larger than the Debye length.

Evolution of the macroscopic fields is governed by Maxwell's equations,

$$\nabla \cdot \mathbf{E} = \frac{\rho_{\text{charge}}}{\epsilon_0}, \quad (2.10)$$

$$\nabla \cdot \mathbf{B} = 0, \quad (2.11)$$

$$\nabla \times \mathbf{E} = -\frac{\partial \mathbf{B}}{\partial t}, \quad (2.12)$$

$$\nabla \times \mathbf{B} = \mu_0 \epsilon_0 \frac{\partial \mathbf{E}}{\partial t} + \mu_0 \mathbf{J}, \quad (2.13)$$

where μ_0 is the permeability of free space, ρ_{charge} is the charge density, $\mathbf{J}$ is the current density and from now on, $\nabla \equiv \nabla_x$. In the general case with three spatial dimensions, the phase space of Eq. 2.8 is 6-D and evolves on small spatial and temporal scales. It is therefore expensive to solve in its entirety for the relatively long times and large spatial scales of an entire ICF implosion. Solution methods to the Vlasov equation include Vlasov-Fokker-Planck (VFP) codes. Rather than directly discretising the 6-D phase space, the velocity space of the distribution function can be expanded in a suitable set of basis functions, such as spherical harmonics. This enables a 3-D discretisation to model 2 spatial dimensions [83, 84]. The velocity space can also be directly discretised for lower dimensional problems [85]. Practically, VFP codes often encounter numerical instabilities around the sharp gradients that are encountered in ICF implosions, due to the generation of large magnitude electric fields in these regions [86]. Alternatively, the distribution function can be approximated by Monte Carlo techniques, which is the technique utilised in Particle in Cell (PiC) codes [87].

2.2.2 The Fluid Equations

Due to the expense of the Vlasov equation, the assumption of local thermodynamic equilibrium can be used to derive the fluid equations. By inserting Eq. 2.7 into Eq. 2.8, multiplying each side of the equation by functions of $\mathbf{v}$ and then integrating over velocity space, the fluid equations can be derived [78]. This process is known as taking moments of the VFP equation. A moment of the VFP equation at order n of $\mathbf{v}$, *i.e.* multiplying the VFP equation by a function of $\mathcal{O}(\mathbf{v}^n)$ ⁵ and integrating over $d\mathbf{v}$, yields an equation which depends upon the $n+1$ moment. A solvable system of equations thus requires an external closure. The inviscid hydrodynamic equations⁶ are obtained from $n = [0, 1, 2]$ moments,

$$\left[\frac{\partial}{\partial t} + \mathbf{u}_\alpha \cdot \nabla \right] \rho_\alpha + \rho_\alpha \nabla \cdot \mathbf{u}_\alpha = 0, \quad (2.14)$$

$$\rho_\alpha \left[\frac{\partial}{\partial t} + \mathbf{u}_\alpha \cdot \nabla \right] \mathbf{u}_\alpha = -\nabla P_\alpha + \mathbf{F}_{\alpha,L} + \mathbf{F}_{\alpha,\text{ext.}} + \mathbf{F}_{\alpha,\text{col.}}, \quad (2.15)$$

$$\left[\frac{\partial}{\partial t} + \mathbf{u}_\alpha \cdot \nabla \right] U_\alpha + (U_\alpha + P_\alpha) \nabla \cdot \mathbf{u}_\alpha = -\nabla \cdot \mathbf{q}_\alpha + Q_{\alpha,\text{ext.}} + Q_{\alpha,\text{col.}}, \quad (2.16)$$

⁵Note that either an inner- or outer-product formulation of the velocity vectors may be employed. Therefore, generically, the $\mathcal{O}(\mathbf{v}^n)$ moment yields a tensor equation of rank n .

⁶Note that since the velocity dependence has been integrated out, all variables are now only functions of $\mathbf{x}$.

where ρ_α is the mass density of species α , $\mathbf{u}_\alpha$ is the fluid velocity, U_α is the internal energy, P_α is the isotropic pressure, $\mathbf{F}_{\alpha,L} = n_\alpha q_\alpha (\mathbf{E} + \mathbf{u}_\alpha \times \mathbf{B})$ is the Lorentz, $\mathbf{F}_{\alpha,\text{ext.}}$ is external forcing, $\mathbf{F}_{\alpha,\text{col.}}$ is collisional forcing, $\mathbf{q}_\alpha$ is the heat flux, $Q_{\alpha,\text{ext.}}$ is external heating or cooling and $Q_{\alpha,\text{col.}}$ is collisional heating or cooling [88]. In its most general form, Eq. 2.16 is obtained by taking the outer product, $\mathbf{v}\mathbf{v}$ moment, which yields an anisotropic pressure tensor. Typically, for unmagnetised plasmas the pressure can be treated as isotropic, however. A closure relation for the plasma pressure can be obtained from an Equation of State (EoS), such as the ideal gas law, which relates pressure to density, temperature and energy density. The heat flux can be obtained from Fourier's law,

$$\mathbf{q}_\alpha = -\kappa \nabla T_\alpha, \quad (2.17)$$

where the thermal conductivity, κ , is obtained from local transport theory [89, 90].

The single-fluid, two temperature hydrodynamic equations can be obtained by assuming quasineutrality ($Zn_i = n_e = n$), that the electron mass is negligible to the ion mass ($m_i \gg m_e$) and modelling only a single ion-species.⁷ The populations will therefore be co-moving, and it follows that,

$$\rho \equiv m_i n_i + m_e n_e \approx m_i n, \quad (2.18)$$

$$\mathbf{u} \equiv \frac{(\rho_i \mathbf{u}_i + \rho_e \mathbf{u}_e)}{\rho} \approx \mathbf{u}_i, \quad (2.19)$$

where the fluid density, ρ , and velocity, $\mathbf{u}$, have been defined. Noting also that collisions within a fluid will not affect its net momentum and that collisional forces between the electrons and ions will cancel, the fluid equations for electrons and ions can be added to obtain the single-fluid, two temperature equations,

$$\left[\frac{\partial}{\partial t} + \mathbf{u} \cdot \nabla \right] \rho + \rho \nabla \cdot \mathbf{u} = 0, \quad (2.20)$$

$$\rho \left[\frac{\partial}{\partial t} + \mathbf{u} \cdot \nabla \right] \mathbf{u} = -\nabla P + \mathbf{F}_L, \quad (2.21)$$

$$\left[\frac{\partial}{\partial t} + \mathbf{u} \cdot \nabla \right] U_e + (U_e + P_e) \nabla \cdot \mathbf{u} = -\nabla \cdot \mathbf{q}_e + Q_{e,\text{ext.}} + Q_{e,\text{col.}}, \quad (2.22)$$

$$\left[\frac{\partial}{\partial t} + \mathbf{u} \cdot \nabla \right] U_i + (U_i + P_i) \nabla \cdot \mathbf{u} = -\nabla \cdot \mathbf{q}_i + Q_{i,\text{ext.}} + Q_{i,\text{col.}}, \quad (2.23)$$

where the total pressure is defined, $P = P_e + P_i$, the external forcing has been dropped, and equilibration between the electron and ion temperatures occurs via the collisional heating terms $Q_{\alpha,\text{col.}}$. By defining the current from electron motion in the frame of the plasma, $\mathbf{J} = -n_e e (\mathbf{u}_e - \mathbf{u})$ and utilising quasineutrality, the Lorentz force term in the single fluid approximation becomes,

$$\mathbf{F}_L = \mathbf{J} \times \mathbf{B}. \quad (2.24)$$

The CHIMERA code solves the two-temperature, single fluid equations. In the absence of

⁷When multiple ion-species compose the plasma, *e.g.* for CH ablator or DT fuel, then the average ion mass (by number density) must be used. Frictional forces and temperature separation between the ion species are also ignored.

magnetic fields, the Lorentz force term goes to zero. Macroscopic magnetic fields can be included under the extended-Magnetohydrodynamics (MHD) framework. The evolution of the fields and their effect on plasma transport and dynamics under this framework is discussed in Sec. 2.2.4. Further details of the numerical solution in CHIMERA to the fluid equations can be found in Refs. [91, 92, 16, 93, 94, 95]. Additional physics packages are included by calculating contributions to the forcing and heating. For example, the non-thermal fusion products are modelled by a Monte Carlo treatment. These energetic particles collide with the thermal plasma constituents and thus deposit energy, which contributes to the electron and ion $Q_{\alpha, \text{ext}}$. [96]. The additional physics packages utilised in the subsequent work presented in this manuscript are discussed briefly below.

2.2.3 Radiation Transport

In High Energy Density Physics (HEDP), a distinction between coherent, long wavelength ($\omega \lesssim \omega_p$) laser radiation, and higher frequency ($\omega \gg \omega_p$) x-ray radiation is drawn. Modelling the former is the main focus of Chap. 3. The long wavelength visible radiation refracts significantly in coronal plasma density gradients, and radiation at this wavelength is not significantly re-emitted by the thermal plasma. In contrast, the x-ray radiation does not significantly refract and is re-emitted in ICF plasma conditions. This x-ray radiation is accounted for numerically in CHIMERA by a radiation transport algorithm, which is cursively described here.

In addition to driving the implosion of indirect-drive ICF experiments, x-ray radiation is also significant in a wide array of laser-driven HEDP physics experiments. For example, the hot coronal plasma in direct-drive radiates a significant amount of energy as thermal Bremsstrahlung emission. This both lowers the coronal temperatures, which reduces the thermal conduction drive efficiency, and also preheats the fuel making it harder to compress. Radiation acts as both a source and sink of energy, as it is radiated and emitted by the plasma. Thermal emission is wavelength-dependent and atomic transitions create sharp resonances of emissivity and opacity in wavelength space. Therefore, the radiation-transport algorithm must be discretised in wavelength, depending on the properties of the material and the plasma conditions.

The radiative transfer equation describes the propagation, absorption, emission and scattering of photons with matter,

$$\frac{\partial I_v}{\partial t} + c\hat{\Omega} \cdot \nabla I_v = \left(\frac{\partial I_v}{\partial t} \right)_{\text{collisions}} + \left(\frac{\partial I_v}{\partial t} \right)_{\text{source}}, \quad (2.25)$$

where c is the speed of light in vacuum, I_v is the spectrally resolved radiation intensity and $\hat{\Omega}$ is the photon direction of travel [88]. The left-hand side describes the advection of radiation at speed c and the first and second terms on the right-hand side describe collisions between matter and photons, and emission or absorption of photons by the matter respectively. Similarly to the Vlasov equation, Eq. 2.25 is 7-D, and therefore is typically highly expensive to solve, so approximations are often employed to make the solution more tractable.

In analogy to derivation of the fluid equations from the Vlasov equation, angular moments of Eq. 2.25 can be taken to reduce the dimensionality of the problem. This also leads to the requirement for a closure relation, and the approach taken in CHIMERA is to use the $P_{1/3}$ closure, which works well for highly isotropic radiation fields [97]. More detail of the CHIMERA implementation is provided in Refs. [98, 99].

2.2.4 Magnetohydrodynamics

The method used to evolve macroscopic magnetic fields and their impact upon the plasma dynamics and transport is discussed in this section. Magnetic fields are able to alter the evolution of the plasma state via the Lorentz force in Eq. 2.21 and by altering the conductivities in Eq. 2.17. The relative importance of magnetisation can typically be broadly assessed by calculating dimensionless numbers. For instance, the plasma β describes the ratio of thermal pressure to the magnetic pressure,⁸

$$\beta = \frac{2\mu_0 P}{|\mathbf{B}|^2}. \quad (2.26)$$

This gives an order of magnitude estimate for when it is necessary to include the Lorentz force in the momentum equation. For example, in high beta plasmas, thermal pressure dominates over magnetic pressure and the Lorentz force is thus unimportant. This is often the case when magnetising a laser-driven ICF implosion, as shall be demonstrated in Chap. 5. Similarly, the Hall parameter, $\omega_\alpha \tau_\alpha$, describes the ratio of the gyrofrequency to the collision frequency for a species, α . Here, ω_α is the gyrofrequency of the species and τ_α is specifically its momentum loss collision time, which from now shall simply be referred to as the collision time. For example, the electron Hall parameter,

$$\omega_e \tau_e = \frac{e|\mathbf{B}|}{m_e} \frac{3\sqrt{m_e}(k_B T_e)^{1.5}}{4\sqrt{2\pi}e^4 Z^2 n_i \ln \Lambda}, \quad (2.27)$$

where $\ln \Lambda$ is the Coulomb logarithm, which is an important parameter for collisional phenomena that is related to the impact parameter of collisions [100, 101, 102]. When $\omega_\alpha \tau_\alpha \sim 1$, then collisional phenomena, such as thermal conduction, are significantly affected by the magnetic field.

The evolution of the magnetic field is governed by Faraday's law (Eq. 2.12) and the electric field by Ampère's law (Eq. 2.13). This does not form a closed system of equations, however, due to the appearance of $\mathbf{J}$ in Ampère's law. Therefore, an additional equation is required to evolve the fields. This can be derived from the momentum equations for ions and electrons, Eq. 2.15 [93]. In the limit of no electron inertia, the left-hand side of Eq. 2.15 is zero, leading to the 'generalised Ohm's law',

$$\mathbf{E} = -\mathbf{u} \times \mathbf{B} + \frac{\mathbf{J} \times \mathbf{B}}{n_e e} - \frac{\nabla P_e}{n_e e} + \frac{1}{n_e e} \left(\frac{\bar{\alpha} \cdot \mathbf{J}}{n_e e} - n_e \bar{\beta} \cdot \nabla T_e \right), \quad (2.28)$$

⁸Using Ampère's law with $(\partial \mathbf{E} / \partial t) \ll c^2$, the Lorentz force in Eq. 2.24 can be recast as the sum of magnetic pressure and tension [94].

where $\bar{\alpha}$ and $\bar{\beta}$ are transport coefficient tensors, the precise form of which is described in, for example, Ref. [94]. Taking the curl of Eq. 2.28 and combining with Faraday's law yields the magnetic induction equation,

$$\frac{\partial \mathbf{B}}{\partial t} = \nabla \times \left(\mathbf{u} \times \mathbf{B} - \frac{\mathbf{J} \times \mathbf{B}}{n_e e} + \frac{\nabla P_e}{n_e e} - \frac{\bar{\alpha} \cdot \mathbf{J}}{n_e^2 e^2} + \frac{\bar{\beta} \cdot \nabla T_e}{e} \right). \quad (2.29)$$

Each term on the right-hand side of the Eq. 2.29 is ascribed a physical affect. The first term describes advection of the field with the movement of the plasma. When this term dominates, the 'ideal MHD' framework is obtained, in which the field is said to be 'frozen-in' to the flow. Ideal MHD can often be applied in fields such as space physics [103, 104] and Magnetic Confinement Fusion (MCF) [105], where many phenomena featuring highly conductive plasmas are well described by this limit. In order from the second term on the right-hand side, the subsequent terms are related to, the Hall effect (collisionless advection of field with current), the Biermann effect (generation of field in plasma gradients), resistive phenomena and thermoelectric phenomena. The non-ideal terms of relevance to the work conducted in this thesis are discussed below.

2.2.4.1 Resistive MHD

When only the first term and the resistive term from Eq. 2.29 are significant (and using a simple form for the transport coefficient, $\bar{\alpha}$), then the induction equation can be written,

$$\frac{\partial \mathbf{B}}{\partial t} = \nabla \times (\mathbf{u} \times \mathbf{B}) - \nabla \times \left(\frac{\eta}{\mu_0} \nabla \times \mathbf{B} \right), \quad (2.30)$$

where η is the resistivity of the plasma [93]. As already discussed, the first term describes advection of the field with the fluid flow, whereas the second term describes resistive diffusion of the plasma. This effect moves $\mathbf{B}$ from regions of high- to low-field, via the damping of current due to collisions between electrons and ions. Resistive diffusion can lead to phenomena such as magnetic reconnection [106, 107]. The relative importance of these two terms is described by the magnetic Reynolds number, which is the ratio of field advection to diffusion,

$$R_m = \frac{\mu_0 |\mathbf{u}| L}{\eta}, \quad (2.31)$$

where L is the characteristic length scale of the gradients.

2.2.4.2 The Nernst Effect

The Nernst effect arises from the thermoelectric term in Eq. 2.29, and makes a contribution to the field induction,

$$\left(\frac{\partial \mathbf{B}}{\partial t} \right)_{\text{Nernst}} = \nabla \times \left(- \frac{\beta_{\perp}}{e |\mathbf{B}|} \nabla T_e \times \mathbf{B} \right), \quad (2.32)$$

where $\beta_\perp$ is a transport coefficient. This contribution looks like an advection of the field in the direction of $-\nabla T_e$ with a velocity,

$$\mathbf{v}_N = -\frac{\beta_\perp}{e|\mathbf{B}|}\nabla T_e. \quad (2.33)$$

Nernst-advection of magnetic field is often of particular importance to magnetised HEDP experiments, acting to move field down temperature gradients, even when the plasma is highly conductive [108, 92]. As this is a collisional phenomenon, it is most important in regions where the electron Hall parameter is low, $\omega_e \tau_e \ll 1$. Explicitly, the transport coefficient $\beta_\perp$ is a function of the Hall parameter, and is large at low values of $\omega_e \tau_e$.

2.2.4.3 Magnetised Thermal Conduction

Magnetic fields can affect collisional transport anisotropically by forcing particles to gyrate around field lines, which changes their collisional behaviour anisotropically. Because the Lorentz force does not affect particle motion in the direction parallel to field lines, transport along the direction, $\hat{\mathbf{b}} = \mathbf{B}/|\mathbf{B}|$, is unaffected by magnetisation. However, the gyration around field lines does restrict collisional transport in the perpendicular direction. Local transport analysis of the heat flux demonstrates that the electron thermal conduction heatflow can be expanded as,

$$\mathbf{q}_e = -\kappa_\parallel \hat{\mathbf{b}} (\hat{\mathbf{b}} \cdot \nabla T_e) - \kappa_\perp \hat{\mathbf{b}} \times (\nabla T_e \times \hat{\mathbf{b}}) - \kappa_\wedge \hat{\mathbf{b}} \times \nabla T_e, \quad (2.34)$$

where the $\kappa_\parallel$ term describes thermal conduction parallel to the field, $\kappa_\perp$ perpendicular to the field and $\kappa_\wedge$ perpendicular to the field and the temperature gradient [90]. The wedge term, $\kappa_\wedge$, describes ‘Righi-Leduc’ heatflow, which occurs along isotherms. Considering the first two terms, when unmagnetised, $\kappa_\parallel = \kappa_\perp$ and conduction is isotropic. As the magnetisation increases, however, $\kappa_\perp$ decreases with respect to $\kappa_\parallel$, for example, $\kappa_\perp/\kappa_\parallel \sim 1/3$ at $\omega_e \tau_e = 1$. This anisotropises the thermal conduction with respect to the orientation of $\mathbf{B}$.

2.2.5 Kinetic Heatflow

While the CHIMERA implementation of thermal transport is simply to use the local limit described in Eq. 2.17, this is not always a valid approximation in laser produced plasmas. When laser-heating is applied to a plasma, laser energy is transferred to the electron population mostly by Inv-Brem, described in Sec. 2.4.3. This process heats the electron fluid to significant temperatures, and energy is transferred to the ions via collisions on a slower timescale. For laser-solid interactions, this often results in sharp electron temperature and density gradients in the conduction zone, where the energy of the hot corona is transported via thermal conduction to an ablation surface. To assess whether local transport theory is valid, the Knudsen number is a convenient dimensionless parameter, which compares the mean free path of electrons λ_{mfp} to the length scale of the gradient, L . Specifically, the Knudsen number is defined,

$$\text{Kn} = \frac{\lambda_{\text{mfp}}}{L} = \frac{3(k_B^2 T_e^2)}{4\sqrt{2\pi} e^4 Z^2 n_i \ln \Lambda L}, \quad (2.35)$$

and transport effects become significantly non-local when $\text{Kn} \sim 0.1$. Knudsen numbers of this magnitude are often observed for high-power laser solid conduction zones [109]. Physically, more energetic particles within a plasma are less collisional and also play a more significant role in heat flux. This is because the heat flux is obtained from a higher order moment than density or fluid velocity, so it is more sensitive to the particles in the high energy tail of the distribution. Therefore, the fast heat carrying particles are able to range through longer length scales and preheat the colder matter more than a local treatment would predict. Although not included within CHIMERA, models for non-local transport exist, which can be included in fluid codes. These include SNB [110, 111, 112, 113], Fast-VFP [114] and RKM [115]. They obtain an improved heat flux estimate, which accounts for non-local conduction, and can be included in the fluid framework as a closure on the energy equations. Implementation of one of these models into CHIMERA could significantly improve the predictive capability for simulations involving laser-solid interactions, such as direct-drive calculations.

2.3 Fluid Description of Waves in Plasma

LPIs are a class of multi-wave coupling phenomena that occur when a light wave excites additional plasma waves. Therefore, understanding LPIs requires some background theoretical knowledge of the relevant waves. In the absence of a macroscopic magnetic field, three classes of wave can propagate in a plasma: the Ion Acoustic Wave (IAW), the EPW and the Electro-Magnetic Wave (EMW), *i.e.* a light wave. This section shall describe the basic theory of these waves and provide dispersion relations for their propagation. The treatment provided in this section assumes that the plasma acts as a fluid. For a full kinetic treatment, the reader is referred to the work of Michel in Ref. [116], from which many of the following relations are also obtained.

2.3.1 Plasma as a Dielectric Medium

Freely moving particles in a plasma are able to respond to external electric fields by relocating, which reduces the size of the field that the plasma sees. In other words plasma is a dielectric medium, and it is therefore convenient to define the permittivity of the plasma in analogy to the solution of Maxwell's equations in matter. The dielectric displacement vector is defined,

$$\mathbf{D} = \epsilon_0 \epsilon \mathbf{E}, \quad (2.36)$$

where ϵ is the permittivity of the plasma.⁹ Defining the electrical conductivity, σ , such that $\mathbf{J} = \sigma \mathbf{E}$ it can be shown that the permittivity takes the form,

$$\epsilon = 1 + i \frac{\sigma}{\epsilon_0 \omega} \equiv 1 + \chi, \quad (2.37)$$

⁹Outside of plasma physics, the displacement vector is usually defined $\mathbf{D} = \epsilon \mathbf{E} = \epsilon_0 \epsilon_r \mathbf{E}$ and ϵ_r is the 'relative' permittivity.




Figure 2.1: Particle motion in the electric fields of a) transverse and b) longitudinal plasma waves.

where the dielectric susceptibility, χ , has been introduced [116]. The permittivity and susceptibility represent the plasma response to applied fields. In general, each species within a plasma contributes separately to the susceptibility, such that $\chi = \sum_{\alpha} \chi_{\alpha}$, where χ_{α} are the contributions from electrons and the individual ion species.

In order to describe waves, the response of the plasma must be described in observable variables. Here, the plasma is treated as a fluid, and is allowed to respond linearly to the applied fields. An initially uniform, stationary plasma in equilibrium is considered with a number density for each species, $n_{\alpha 0}$. The species is assumed to be isothermal or adiabatic, and therefore solution of the energy equation, Eq. 2.16 is not required. An ideal gas equation of state is used to obtain a closure for the pressure in the momentum equation, Eq. 2.15,

$$P_{\alpha} = n_{\alpha} T_{\alpha} = C n_{\alpha}^{\gamma_{\alpha}}, \quad (2.38)$$

where γ_{α} is the adiabatic index of species α and C is a constant. In analogy to ideal gases, γ_{α} is related to the number of degrees of freedom of the fluid, for example, $\gamma_{\alpha} = 1$ for an isothermal fluid and $\gamma_{\alpha} = 3$ for an adiabatic fluid. In this framework, the plasma response is completely described by a velocity ($\delta \mathbf{v}_{\alpha}$) and density (δn_{α}) perturbation. Linearising, we seek a solution to the following plasma response,

$$n_{\alpha}(\mathbf{x}, t) = n_{\alpha 0} + \delta n_{\alpha}(\mathbf{x}, t), \quad (2.39)$$

$$\mathbf{v}_{\alpha}(\mathbf{x}, t) = \mathbf{v}_{\alpha 0} + \delta \mathbf{v}_{\alpha}(\mathbf{x}, t) = 0 + \delta \mathbf{v}_{\alpha}(\mathbf{x}, t), \quad (2.40)$$

$$\mathbf{E}(\mathbf{x}, t) = \mathbf{E}_0 + \delta \mathbf{E}(\mathbf{x}, t) = 0 + \delta \mathbf{E}(\mathbf{x}, t), \quad (2.41)$$

$$\mathbf{B}(\mathbf{x}, t) = \mathbf{B}_0 + \delta \mathbf{B}(\mathbf{x}, t) = 0 + \delta \mathbf{B}(\mathbf{x}, t), \quad (2.42)$$

where the subscript 0 quantities are the (constant) equilibrium values and $\delta \mathbf{E} = \mathbf{E}$ and $\delta \mathbf{B} = \mathbf{B}$ are the applied fields. Inserting the linearised equations for each species into Eq. 2.14 (continuity equation) and Eq. 2.15 (momentum equation) it can be shown that in the absence of

a magnetic field,

$$\frac{\delta n_\alpha}{n_{\alpha 0}} = \frac{\mathbf{k} \cdot \delta \mathbf{v}_\alpha}{\omega}, \quad (2.43)$$

$$\omega \delta \mathbf{v}_\alpha = \gamma_\alpha v_{T\alpha}^2 \frac{\delta n_\alpha}{n_{\alpha 0}} \mathbf{k} + i \frac{q_\alpha}{m_\alpha} \mathbf{E}, \quad (2.44)$$

where ω is the conjugate variable of time, t , from a Laplace transform and $\mathbf{k}$ is the conjugate variable of space, $\mathbf{x}$, from a Fourier transform. The same equations can be obtained with a less formal treatment, by assuming a plane wave Ansatz for the perturbed quantities, $\delta \sim \exp[i(\mathbf{k} \cdot \mathbf{x} - \omega t)]$. These equations can then be combined to relate δn_α and $\delta \mathbf{v}_\alpha$ to $\mathbf{E}$,

$$[\omega^2 - \gamma_\alpha v_{T\alpha}^2 k^2] \frac{\delta n_\alpha}{n_{\alpha 0}} = i \frac{q_\alpha}{m_\alpha} \mathbf{k} \cdot \mathbf{E}, \quad (2.45)$$

$$[\omega^2 - \gamma_\alpha v_{T\alpha}^2 \mathbf{k}(\mathbf{k})] \delta \mathbf{v}_\alpha = i \frac{q_\alpha \omega}{m_\alpha} \mathbf{E}. \quad (2.46)$$

Assuming a plane wave of the form $\mathbf{E} \propto \exp[i(\mathbf{k} \cdot \mathbf{x} - \omega t)]$, we identify $\mathbf{k}$ as the wavenumber and ω as the angular frequency. Using the Ansatz that waves can either be transverse ($\mathbf{k} \cdot \mathbf{E} = 0$) or longitudinal ($\mathbf{k} \times \mathbf{E} = 0$), dispersion relations in the fluid regime can be obtained for admissible solutions using these equations. Fig. 2.1.a and Fig. 2.1.b plot the motion of particles under the influence of the electric field from each of these wave varieties.

2.3.2 Electromagnetic Waves

Firstly, the transverse wave solution shall be obtained, which is an EMW propagating in the plasma. Inserting the transverse wave Ansatz ($\mathbf{k} \cdot \mathbf{E} = 0$) into Eq. 2.46 demonstrates that $\mathbf{k} \cdot \delta \mathbf{v}_\alpha = 0$, and therefore $\delta \mathbf{v}_\alpha = \delta \mathbf{v}_{\alpha \perp}$. The transformed equation for the conservation of momentum, Eq. 2.44 yields,

$$\delta \mathbf{v}_{\alpha \perp} = i \frac{q_\alpha}{m_\alpha \omega} \mathbf{E}, \quad (2.47)$$

which demonstrates that under the Lorentz force of the field, the particles undergo oscillatory ‘quiver’ motion. The quiver velocity of the ions is ignored because $m_i \gg m_e$ and therefore the ions are considered as a static background compared to the electron oscillations. Therefore, the linearised current is only from the electron motion, $\sum_\alpha \delta \mathbf{J}_\alpha \sim \delta \mathbf{J}_e = -en_{e0} \delta \mathbf{v}_e = \sigma \mathbf{E}$. Utilising the definition of the permittivity from Eq. 2.37, the plasma response to an EMW is obtained,

$$\varepsilon = 1 - \frac{\omega_p^2}{\omega^2} = 1 - \frac{n_{e0}}{n_{cr}}, \quad (2.48)$$

$$\chi = -\frac{\omega_p^2}{\omega^2}, \quad (2.49)$$

where the plasma frequency, defined previously in Eq. 2.5, has emerged.

A dispersion relation can also be obtained by utilising Maxwell’s equations. Taking the curl of Faraday’s law (Eq. 2.12) and combining this with Ampère’s law (Eq. 2.13) to eliminate




Figure 2.2: The dispersion relation for an EMW propagating in a plasma. For $\omega < \omega_p$, the oscillatory plane wave Ansatz turns into a decaying evanescent wave, thus the wave does not propagate.

B yields,

$$[\partial_t^2 - c^2 \nabla^2 + c^2 \nabla(\nabla \cdot)] \mathbf{E} = -\frac{1}{\epsilon_0} \partial_t \mathbf{J}. \quad (2.50)$$

Assuming a plane wave Ansatz for $\mathbf{E}$ and inserting Eq. 2.48 to eliminate $\mathbf{J}$, the dispersion relation and wave equation for an EMW are obtained,

$$\omega^2 - \omega_p^2 - c^2 k^2 = 0, \quad (2.51)$$

$$(\partial_t^2 + \omega_{pe}^2 - c^2 \nabla^2) \mathbf{E}(\mathbf{x}, t) = 0. \quad (2.52)$$

The dispersion relation, Eq. 2.51, is plotted in Fig. 2.2. For $\omega < \omega_p$, k becomes imaginary and therefore the plane wave Ansatz turns into a non-propagating, exponentially decaying evanescent wave. The phase v_ϕ and group v_g velocities of the light are also obtained from the dispersion relation,

$$v_\phi = \frac{\omega}{k} = \frac{c}{\sqrt{\epsilon}}, \quad (2.53)$$

$$v_g = \frac{\partial \omega}{\partial k} = \sqrt{\epsilon} c, \quad (2.54)$$

which defines the refractive index of plasma, $n_{\text{ref}} = \sqrt{\epsilon} = \sqrt{1 - n_e/n_{\text{cr}}}$. Note that $\sqrt{\epsilon} < 1$ so $v_\phi > c$. However, the group velocity, which is the speed that information travels, remains subluminal.

2.3.3 Plasma Waves

Two propagating longitudinal waves can exist in an unmagnetised plasma. One is an EPW, which is an oscillation of the electron fluid on a timescale which is too fast for the ions to respond. The other is an IAW, which is a slower, joint oscillation of the electron and ion fluid. The inertia of the EPW and IAW are set by the electron and ion masses, respectively. Electron




Figure 2.3: Illustration of electron (small blue circle) and ion (large red circle) motion in a) an IAW and b) an EPW.

pressure drives the EPW, whereas both fluids contribute to the IAW pressure, depending on the ionisation level of the plasma. Fig. 2.3 illustrates particle motion in each wave. Dispersion relations can be obtained similarly to the EMW method from Sec. 2.3.2. By first inserting the longitudinal Ansatz ($\mathbf{k} \times \mathbf{E} = 0$) into Eq. 2.46 to obtain the particle motion in a longitudinal wave,

$$\delta \mathbf{v}_\alpha = i \frac{q_\alpha \omega}{m_\alpha} \frac{\mathbf{E}}{\omega^2 - \gamma_\alpha v_{T\alpha}^2 k^2}. \quad (2.55)$$

Unlike the EMW, we make no assumptions yet about which species contribute to the current, and thus the general form of the conductivity for longitudinal waves includes contributions from all species,

$$\sigma = i\omega\epsilon_0 \sum_\alpha \frac{\omega_{p\alpha}^2}{\omega^2 - \gamma_\alpha v_{T\alpha}^2 k^2}, \quad (2.56)$$

which yields the longitudinal wave permittivity and equivalently susceptibility,

$$\epsilon = 1 + \sum_\alpha \chi_\alpha, \quad (2.57)$$

$$\chi_\alpha = -\frac{\omega_{p\alpha}^2}{\omega^2 - \gamma_\alpha v_{T\alpha}^2 k^2}. \quad (2.58)$$

Inserting the plane wave Ansatz into 2.50 (combined Faraday-Maxwell laws) and using Eq. 2.57 to eliminate $\mathbf{J}$ yields the generalised dispersion relation for longitudinal waves in a plasma,

$$\epsilon \mathbf{E} = 0. \quad (2.59)$$

Further use of an Ansatz can be utilised to obtain a specific description of a particular plasma wave. Note that unlike transverse waves, where $\delta n_\alpha = 0$, which is seen by inserting $\mathbf{k} \cdot \mathbf{E} = 0$ into Eq. 2.45 (particle density under the electric field of a wave), longitudinal waves do give rise to a density perturbation. This can be heuristically understood by looking at Fig. 2.1, which illustrates particle bunching in an oscillatory, longitudinal electric field. The charge separation from these longitudinal oscillations leads to an additional electric field component within the plasma, which is important for wave damping.

The longitudinal susceptibility (Eq. 2.58) has a pole at $\omega^2 = \gamma_\alpha v_{T\alpha}^2 k^2$, suggesting unbounded growth of the wave when its phase velocity is close to the thermal velocity of the species, α . This does not occur in actuality however, due to a process known as Landau damping, which is a resonant exchange of energy from the wave to particles. When the phase velocity of the wave approaches this resonance, particles with velocity just below the phase velocity, $|\mathbf{v}| \lesssim |\mathbf{v}_\phi|$, are trapped in the potential created by the charge separation and accelerated. This transfers energy from the wave to the particles and thus prevents unbounded growth of the density perturbation. Particles with velocity $|\mathbf{v}| \gtrsim |\mathbf{v}_\phi|$ are equivalently trapped and decelerated, leading to energy gain of the wave. However, phase velocities of the plasma waves are greater than the thermal velocity where the fluid, Maxwellian distribution function peaks, and thus more particles are accelerated by this effect. This energy gain of particles in longitudinal plasma waves is the basis for laser-plasma accelerators [117, 118].

2.3.3.1 Electron Plasma Waves

Firstly, the EPW shall be considered. These are fast oscillations of the electron fluid such that the heavier ions do not respond on the oscillation timescale. To arrive at the dispersion relation, the assumption that the phase velocity is much faster than the electron thermal velocity is made, such that the wave can freely propagate without experiencing strong Landau damping, $\omega/k \gg v_{Te}$. This assumption also yields $\lambda \gg v_{Te}/\omega$, which states that electrons travelling at the thermal velocity will not move a significant portion of the wavelength, λ , in one oscillation. Temperature equilibration is thus limited throughout the wave, which means that there is little heat flux. Therefore, the electron fluid is assumed to be adiabatic. An adiabatic fluid has $\gamma_e = 3$, which is equivalent to a single degree of freedom in the motion. Finally, the assumption is made that the ions are slow due to their large mass, and thus do not contribute to the permittivity.¹⁰ By using these assumptions in the generalised dispersion relation for longitudinal plasma waves, Eq. 2.59, the dispersion relation for the EPW is obtained,

$$\omega^2 = \omega_{pe}^2 + 3k^2 v_{Te}^2. \quad (2.60)$$

¹⁰Explicitly for the EPW, $\omega/k \gg v_{Te} \gg v_{Ti}$ and thus $|\chi_e| \gg |\chi_i|$

Just as for light waves, EPW cannot propagate at $\omega \geq \omega_{pe}$. The phase and group velocity of the wave are obtained from this dispersion relation,

$$v_\phi = \frac{\omega}{k} = v_{Te} \frac{\sqrt{1 + 3k^2 \lambda_D^2}}{k \lambda_D}, \quad (2.61)$$

$$v_g = \frac{\partial \omega}{\partial k} = 3 v_{Te} \frac{k \lambda_D}{\sqrt{1 + 3k^2 \lambda_D^2}}. \quad (2.62)$$

2.3.3.2 Ion Acoustic Waves

The IAW is of particular importance to the work conducted in this thesis, because it is the wave that mediates the energy exchange between light waves in CBET. To obtain the dispersion relation for this wave, the phase velocity of the wave is assumed to be much greater than the thermal velocity of the ions, but much less than the electron thermal velocity.¹¹ Using the same logic as the electrons in the EPW, the ion fluid is assumed to be adiabatic, so $\gamma_i = 3$. The electrons, however, move much faster and can thus propagate over many wavelengths. Therefore, they are able to equilibrate their temperature across the wave in a single oscillation, and are taken to be isothermal ($\gamma_e = 1$). When inserted into Eq. 2.59, these assumptions yield the dispersion relation,

$$\omega^2 = \frac{k^2}{1 + k^2 \lambda_D^2} \frac{Z T_e}{m_i} + 3k^2 v_{Ti}^2. \quad (2.63)$$

For oscillation wavelengths much larger than the Debye length, $k \lambda_D \ll 1$, this simplifies to,

$$\omega \approx c_s k, \quad (2.64)$$

where the sound speed is defined,

$$c_s = \sqrt{\frac{Z k_B T_e + 3 k_B T_i}{m_i}}. \quad (2.65)$$

The phase and group velocity of the light are simply c_s in this limit, so the wave is non-dispersive.

Commonly in ICF, plasmas are composed of multiple ion species, for instance the ablator of OMEGA targets is commonly composed of a CH plastic or SiO₂ glass shell. The IAW properties are typically greatly altered due to the additional susceptibility contributions from each species. This leads to multiple IAW modes with distinct speeds and damping. Denoting the ‘heavy’ and ‘light’ ion species as h and l respectively, a fast ion mode exists when the phase velocity is greater than both ion thermal velocities,

$$v_{Th} \leq v_{Tl} \ll v_{\phi, \text{fast}} \ll v_{Te}. \quad (2.66)$$

A fast mode is also observed when only the heavy ion species is slower than the phase veloc-

¹¹For the IAW, $v_{Ti} \ll \omega/k \ll v_{Te}$.

ity,

$$\nu_{Th} \ll \nu_{\phi, \text{slow}} \ll \nu_{Tl} \ll \nu_{Te}. \quad (2.67)$$

This is of particular relevance to the work in this manuscript, since CBET in direct-drive typically occurs in multi-ion ablaters, and it is moderated by an energy exchange via an IAW. Therefore, for example, different wave phenomena can affect the growth and saturation of energy transfer and must be accurately accounted for in computational modelling. This is accounted for by the fluid and kinetic CBET treatments, described in Sec. 3.5.1.

2.4 Propagation and Absorption of Light in a Plasma

This section shall describe the propagation and absorption of light in plasma, in the regime of typical ICF conditions. Typically, sub-critical laser-driven ICF plasmas can be treated as a linear medium and intensity of the light is sufficiently modest that it does not directly alter the refractive index of the plasma. By making the WKB approximation, which is that over a wavelength the plasma may be treated as uniform, the equations of ray-tracing are derived. Ray-tracing is the tool used to model lasers in the SOLAS module, the development of which is described in Chap. 3. Absorption processes, which are important to energy deposition in laser-drive ICF are then introduced, particularly collisional absorption, otherwise known as Inv-Brem. An absorption per-unit-length for Inv-Brem is provided, which is suitable for integration along ray trajectories.

2.4.1 The WKB Approximation

Laser-driven plasmas are heated and expand, leading to density gradients in the region where the light propagates. The propagation path of the light through this plasma is governed by the refractive index, or equivalently permittivity, $\epsilon = n_{\text{ref}}^2 = 1 - n_e/n_{\text{cr}}$. Thus, light refracts in plasmas with a density gradient. Here, a description of light propagation through a non-uniform plasma (*i.e.* $n_e = n_e(\mathbf{x})$) is obtained, in the limit that the plasma density does not vary significantly over the wavelength of the oscillation. It is convenient to work with vector potentials, $\mathbf{B} = \nabla \times \mathbf{A}$ and $\mathbf{E} = -\partial_t \mathbf{A}$, rather than electric field for laser-plasma waves.¹² This is because electric fields can describe both longitudinal and transverse waves, whereas the vector potential is purely for transverse oscillations and as such cannot describe plasma waves. Using the vector potential, the EMW wave equation, Eq. 2.52, can be rederived from Maxwell's equations,

$$\left(\partial_t^2 + \omega_{pe}^2 - c^2 \nabla^2 \right) \mathbf{a} = 0, \quad (2.68)$$

where the normalised vector potential is defined,

$$\mathbf{a} = \frac{e}{\omega m_e c} \mathbf{A}, \quad (2.69)$$

¹²Note that the Coulomb gauge has been used here, $\nabla \cdot \mathbf{A} = 0$, otherwise $\mathbf{E} = -\partial_t \mathbf{A} - \nabla \phi$, where ϕ is the scalar potential.

where ω is the angular frequency of the light wave. The light is said to be relativistic when $|\mathbf{a}| \sim 1$, because $|\mathbf{a}|$ is equivalent to the maximum quiver velocity of an electron divided by the speed of light. If the wave propagates normally up a plasma density gradient (aligned with z), polarisation of the wave can be ignored by rotational symmetry and a solution can be sought of the form,

$$a(z, t) = \frac{1}{2} f(z) a_0 e^{-i\omega t} + c.c., \quad (2.70)$$

where a_0 is the oscillation amplitude, $c.c.$ is the complex conjugate and $f(z)$ is the spatial dependence of the oscillation. In this limit, the WKB (for Wentzel, Kramers, Brillouin) approximation can be used to find a solution for $f(z)$.

Inserting Eq. 2.70 into the wave equation, Eq. 2.68, yields the 1-D Helmholtz equation,

$$[\partial_z^2 + \varepsilon(z) k_0^2] f(z) = 0, \quad (2.71)$$

where $k_0 = \omega/c$ is the vacuum wavenumber. The Helmholtz equation describes the steady state propagation of light through a plasma, which is defined by the permittivity, ε . An Ansatz is used [116],

$$f(z) = e^{ik_0\varphi(z)}, \quad (2.72)$$

where φ is a series expansion [119],

$$\varphi(z) = \left[S_0(z) + \frac{S_1(z)}{ik_0} + \frac{S_2(z)}{(ik_0)^2} + \dots \right], \quad (2.73)$$

which is valid so long as the expansion parameter is large, $k_0 \gg |\nabla\varepsilon|/\varepsilon$, which is shown explicitly at the end of the section.

By combining Eqs. 2.71, 2.72 and 2.73 and keeping only the dominant, $\mathcal{O}(k_0^{n=0})$ terms, the expression for the first order expansion term is obtained,

$$S_0(z) = \pm \int_0^z \sqrt{\varepsilon(z)} dz. \quad (2.74)$$

If only this term is kept and there is a uniform plasma (therefore ε is constant), then a plane wave propagating solution is obtained for the field,

$$a(z, t) = \frac{a_0}{2} \exp[i(\pm k_0 z - \omega t)]. \quad (2.75)$$

If the next order terms, $\mathcal{O}(k_0^{n=-1})$, are collected and equated, then the following expression is retrieved,

$$S_1(z) = i \ln(\varepsilon^{1/4}) + \varphi_0, \quad (2.76)$$

where φ_0 is an integration constant, identified as the initial phase of the light. Note that in this work, the *dimensional* phase is used, such that the phase terms φ , have spatial units. This is often convenient to work with, because the parametric variables of phase, path and arc length all share the same dimension. The *dimensionless* phase is obtained by multiplication with the vacuum wavenumber, $\psi \equiv k_0\varphi$. An expression for the light field is obtained by

keeping the next order terms [116],

$$a(z, t) = \frac{1}{2} \frac{a_0}{\varepsilon^{1/4}} \exp \left[\pm i \int_0^z k(z) dz - i\omega t + i k_0 \varphi_0 \right] + c.c., \quad (2.77)$$

where $k(z) = \sqrt{\varepsilon(z)} k_0$. The $\varepsilon^{1/4}$ term outside the exponential leads to higher field in regions of the plasma with higher densities and is known as ‘field-swelling’. It occurs due to the reduced group velocity of light at lower ε values, resulting in a pile up of field as light propagates up a density gradient. The validity condition of the WKB solution in Eq. 2.77 can be demonstrated by mandating that the $\mathcal{O}(k_0^{n=-2})$ terms from the expansion are much smaller than the $\mathcal{O}(k_0^{n=-1})$ terms. This leads to the validity domain (for the normally incident light considered here),

$$\frac{|\nabla \varepsilon|}{\varepsilon} \ll k, \quad (2.78)$$

where, stated again for emphasis, $k = \sqrt{\varepsilon} k_0$. In other words, the length scale L associated with $\nabla \varepsilon$ must be far larger than the oscillation wavelength $\lambda = 2\pi/(k_0 \sqrt{\varepsilon})$. This condition is violated either for plasmas with extreme density gradients, or in a narrow region near the critical surface, where the oscillation wavelength grows to infinity. When calculating field amplitudes under the WKB approximation therefore, specialised treatments are required in the region close to the critical surface (or turning point of the light when it is not normally incident), as is shown explicitly in Sec. 3.4.1.

2.4.2 Ray-Tracing

The WKB approximation correctly captures the behaviour of light, apart from a narrow region near the wave turning point. Ray-tracing is a technique which provides a framework to integrate the WKB solution along the trajectory of the light, at discrete points on an initial phase front. This turns out to be a highly useful tool for numerical implementation of laser modelling. Starting from the Helmholtz equation, now generalised to multiple dimensions [119],

$$[\nabla^2 + \varepsilon(\mathbf{x}) k_0^2] a(\mathbf{x}) = 0, \quad (2.79)$$

a solution is sought of the form,

$$a(\mathbf{x}) = A(\mathbf{x}) e^{i k_0 \varphi(\mathbf{x})} \equiv A(\mathbf{x}) e^{i \psi(\mathbf{x})}, \quad (2.80)$$

where going from 1-D→3-D the change $f(z) \rightarrow A(\mathbf{x})$ is made, and $A(\mathbf{x})$ is called the ‘ray-amplitude’. Similarly to Sec. 2.4.1, $\varphi(\mathbf{x})$ is expanded, and the Ansatz is inserted into the Helmholtz equation, which returns the following equations:

$$[\nabla \varphi(\mathbf{x})]^2 - \varepsilon(\mathbf{x}) = 0, \quad (2.81)$$

$$2 [\nabla A(\mathbf{x}) \cdot \nabla \varphi(\mathbf{x})] + A(\mathbf{x}) \nabla^2 \varphi(\mathbf{x}) = 0, \quad (2.82)$$

where Eq. 2.81 is known as the Eikonal equation, and Eq. 2.82 is known as the transport equation. So far, this is simply a multidimensional corollary of the same procedure that was

followed in Sec. 2.4.1.

The ray-tracing equations are obtained by noticing that Eq. 2.81 describes the Hamiltonian of a system,

$$\mathcal{H} = \frac{1}{2} [\mathbf{p}^2 - \varepsilon(\mathbf{x})], \quad (2.83)$$

which has a momentum $\mathbf{p} = \nabla\varphi(\mathbf{x})$, and a potential $-\varepsilon(\mathbf{x})$. Thus, using the characteristic technique, the following equations of motion are obtained,

$$\frac{d\mathbf{x}}{d\tau} = \mathbf{p}, \quad (2.84)$$

$$\frac{d\mathbf{p}}{d\tau} = \frac{1}{2} \nabla \varepsilon(\mathbf{x}), \quad (2.85)$$

which describe the evolution of a wave packet (typically called a ray) along its optical path length, τ , through a medium defined by $\varepsilon(\mathbf{x})$ [119]. The wavevector is simply related to this momentum $\mathbf{k} = k_0 \mathbf{p}$. Similarly to the WKB solution, it is valid as long as the evolution of the envelope function, *i.e.* the ray amplitude $A(\mathbf{x})$, are slow compared to the spatial oscillations described by $|\mathbf{k}|$. These equations can be readily integrated along τ to describe the trajectory of the field, from a discrete initial point on its initial surface.

The evolution of the ray amplitude is obtained by integrating the transport equation, Eq. 2.82,

$$A(\tau) = A(0) \left| \frac{D(0)}{D(\tau)} \right|^{1/2}, \quad (2.86)$$

$$D(\tau) = \begin{bmatrix} \frac{\partial x}{\partial \zeta_1} & \frac{\partial x}{\partial \zeta_2} & \frac{\partial x}{\partial \tau} \\ \frac{\partial y}{\partial \zeta_1} & \frac{\partial y}{\partial \zeta_2} & \frac{\partial y}{\partial \tau} \\ \frac{\partial z}{\partial \zeta_1} & \frac{\partial z}{\partial \zeta_2} & \frac{\partial z}{\partial \tau} \end{bmatrix}, \quad (2.87)$$

where $[x, y, z]$ and $[\zeta_1, \zeta_2, \tau]$ are the ray real-space and phase-space coordinates, respectively and $D(\tau)$ is the Jacobian for the coordinate transform from phase-space to real-space [67]. Physically, the phase space coordinate τ represents the distance along a wave packet trajectory, and $[\zeta_1, \zeta_2]$ are the initial position of the wave packets on the 2-D beam port ($\tau = 0$). This represents the conservation of energy along a tube of the light, defined by an infinitesimally small separation of rays. Here, infinitesimally small means that it is small compared to the spatial variations of the medium, $L \sim |\varepsilon/\nabla\varepsilon|$. Unlike Eqs. 2.84 and 2.85, it is not immediately obvious how to integrate this equation along a ray path. The solution method taken in this manuscript is to trace a small bundle of rays and calculate their separation at discrete points along the path, which is described in more detail in Sec. 3.4.1.

Temporal variation of the medium may also affect the evolution of the field. For most HEDP experiments, the time taken for light to traverse the region of interest is small compared to the temporal evolution of ε , therefore it is usually safe to disregard its effect. One exception to this is LPIs. The mechanism for these instabilities is described in more detail in Sec. 2.5, but they can rely on a precisely matched EMW frequency condition to excite. Therefore, even small differences to wave frequencies can lead to errors in modelling. The




Figure 2.4: An electron-ion collision in the presence of an oscillating electric field. The collision time is much shorter than the oscillation time. The magnitude of the electron velocity just before ($\mathbf{w}$) and just after ($\mathbf{w}'$) the collision is the same, since the collision is elastic. However, the electron continues its quiver motion ($\mathbf{u}$) just after the collision and thus the time-averaged velocities before ($\mathbf{v}$) and after ($\mathbf{v}'$) are different. The figure has been reproduced with permission from Ref. [116].

evolution of the frequency of the light is obtained from the dispersion relation [120],

$$\frac{d\omega^2}{dt} = \frac{d\omega_p^2}{dt}. \quad (2.88)$$

This leads to an equation for the integration along ray trajectories,

$$\frac{d\omega}{d\tau} = \frac{\omega}{2c} \frac{\partial(n_e/n_{cr})}{\partial t}, \quad (2.89)$$

where $dt = d\tau/c$. Eq. 2.89 describes the temporal bunching or rarefaction of successive wavefronts propagating at a discrete point in space, due to a temporally increasing or decreasing refractive index. The evolution timescale is typically long for ICF compared $1/\omega$ and therefore the percentage change of ω is small. Eq. 2.89 can thus be evolved independently of Eqs. 2.84 and 2.85. In direct-drive ICF experiments, evolution of the coronal plasma can be diagnosed by measuring this frequency shift [121].

2.4.3 Inverse-Bremsstrahlung

Lasers are used in HEDP experiments as an energy source, which drives the plasma dynamics. The energy from the laser must therefore be absorbed into the plasma to drive this motion. Inv-Brem, otherwise known as collisional absorption, is the dominant heating mechanism for modern direct-drive experiments. This is the absorption of wave energy by the electron population of the plasma as they undergo collisions with the ions. Collisions with the ions are required for energy transfer from the wave to the electrons, as without collision, the electrons would simply gain and lose energy cyclically as they oscillated in the laser elec-

tric field. When collisions with ions are included, electrons are deflected, disrupting their cyclical quiver energy exchange. Therefore, the full portion of kinetic energy that an electron gains from a quiver oscillation, is not restored to the field, resulting in damping of the wave and net energy gain of the electron population.

The physical picture of the absorption process is outlined in Fig. 2.4. In this figure, an electron undergoes a collision with an ion, while oscillating in an electric field. The total motion of the electron, $\mathbf{w}$, is described by a quiver component, $\mathbf{u}$, and a guiding centre component, $\mathbf{v}$, which is averaged over the oscillatory motion,

$$\mathbf{w} = \mathbf{v} + \mathbf{u}, \quad (2.90)$$

where $|\mathbf{u}| = e|\mathbf{E}|/m_e\omega$ is the quiver velocity of the electrons in the field, defined in Eq. 2.47. The thermal energy of the population is effectively the randomly distributed velocities, $\mathbf{v}$, which is a Maxwellian if the plasma is in equilibrium. The quiver velocity of the electrons is thus effectively a guiding centre drift of a Maxwellian compared to a stationary ion population, which relaxes to a Maxwellian by collisions with the ions. An equation of motion for the quiver component can therefore be written,¹³

$$\frac{\partial \mathbf{u}}{\partial t} = -\frac{e}{m_e} \mathbf{E} - \nu_{ei} \mathbf{u}, \quad (2.91)$$

where ν_{ei} is the electron-ion collision rate [116],

$$\nu_{ei} = \frac{2n_i}{\sqrt{2\pi}m_e^2 v_{Te}^3} \left(\frac{Ze^2}{4\pi\epsilon_0} \right)^2 \frac{4\pi}{3} \ln(\Lambda). \quad (2.92)$$

The Coulomb logarithm, $\ln(\Lambda)$, is the logarithm of the ratio of minimum and maximum impact parameters for a collision. Its expression takes different forms for different physical effects. When considering a quivering electron, oscillating in an electric field, the appropriate form for Λ is,

$$\Lambda = \frac{v_{Te}}{V}, \quad (2.93)$$

$$V = \max(\omega, \omega_p) \times \max\left(\frac{Ze^2}{k_B T_e}, \frac{\hbar}{\sqrt{m_e k_B T_e}} \right), \quad (2.94)$$

where $\hbar$ is the reduced Planck constant [66].

Assuming a plane wave solution, $(\mathbf{E}, \mathbf{v} \propto \exp(i\omega t))$, an expression for the quiver velocity is found from Eq. 2.91,

$$\mathbf{u} = \frac{-ie}{m(\omega + i\nu_{ie})} \mathbf{E}. \quad (2.95)$$

Recalling the definition of the current, $\mathbf{J}_e = -en_e \mathbf{u} = -i\omega\epsilon_0\chi_e$, the permittivity ($\epsilon = 1 - \chi_e$) is

¹³Note that only the dominant terms in an expansion around the ‘small’ parameter, $|\mathbf{u}|/|\mathbf{v}| \equiv \nu_{\text{quiv.}}/\nu_{Te}$ have been kept. It has been shown that the following Inv-Brem absorption rate is valid for $\nu_{\text{quiv.}}/\nu_{Te} \leq 1$ [122].

found to be,

$$\varepsilon = 1 - \frac{\omega_p^2}{\omega^2 + i\omega\nu_{ei}}. \quad (2.96)$$

In the small absorption limit, where $\nu_{ei} \ll \omega$, the wavevector, $k = \sqrt{\varepsilon}\omega/c$ is complex,

$$k = \frac{\omega}{c}\sqrt{\varepsilon} + i\frac{1}{2}\frac{\nu_{ei}}{c\sqrt{\varepsilon}}\frac{n_e}{n_{cr}}. \quad (2.97)$$

The small absorption limit is generally valid for the hot coronal plasmas of direct-drive ICF experiments, apart from in a thin region close to the critical surface [123]. The imaginary component describes a damping of the wave amplitude as it propagates in space, with an absorption rate per $d\tau$,

$$\kappa_{IB} = 2\frac{\nu_{ei}}{c}\frac{n_e}{n_{cr}}, \quad (2.98)$$

where the $\sqrt{\varepsilon}$ drops out when converting from unit arc length to path length. The factor of 2 in Eqs. 2.97 and 2.98, arises from a Taylor expansion of $\sqrt{\varepsilon}$ around the small parameter ν_{ei}/ω .

A further improvement can be made to Eq. 2.98, by noting that Inv-Brem preferentially heats the colder population of electrons, because the collision frequency $\nu_{ei} \propto v^{-3}$. If the heating rate is fast compared to the collisional rate which returns the distribution function to Maxwellian, then this can lead to super-Gaussian distribution functions. These distributions have a lower population of the cold electrons, therefore the overall absorption rate is decreased. This process is known as the Langdon effect [124] and is parameterised by the ratio of the heating rate to the thermalising, electron-electron collision rate,

$$\alpha = Z^* \frac{\nu_{quiv}^2}{\nu_{Te}^2}, \quad (2.99)$$

where $Z^* = \langle Z^2 \rangle / \langle Z \rangle$ (the average is taken over multiple ion species if present) and $\nu_{quiv} = e|\mathbf{E}|/m_e\omega$ is the quiver velocity. The modified absorption kernel is then the value from Eq. 2.98, multiplied by the factor [67],

$$f_L = [1 - 0.553 / (1 + [0.27/\alpha]^{0.75})]. \quad (2.100)$$

One issue with accurately implementing this effect in a ray-tracing code is that the field along each ray must be known to calculate ν_{quiv} , and therefore the amplitude of the ray must be integrated, which is not trivial. Non-Maxwellian electron distributions also lead to modified transport properties of the plasma [125]. These modifications to thermal conduction are important to consider for laser direct-drive implosions, because the laser energy is transported from the absorption region to the ablation surface via thermal conduction. Recent work has also been conducted, using laser heated gas jets to explore the validity of Eq. 2.98 for describing the Inv-Brem absorption rate [126]. This work demonstrated that both the Langdon effect and modifications due to ion-screening are required to best reproduce the measured laser absorption for direct-drive relevant conditions.




Figure 2.5: Illustration of resonance absorption for a non-normally incident, p-polarised light wave, propagating up a density gradient. At the turning point of the light, if sufficiently close to the critical density, then the evanescent light can tunnel through to the critical surface and resonantly excite an EPW. The EPW propagates down the density ramp and transfers energy to the electrons through damping. The figure has been reproduced with permission from Ref. [116].

2.4.4 Resonance Absorption

An additional absorption mechanism is also important for many HEDP experiments, known as resonance absorption. This mechanism is far less important for direct-drive experiments on OMEGA and the National Ignition Facility (NIF), due to the use of frequency tripled light and long plasma scale lengths, however a short discussion is included for completeness. An illustration of the procedure is shown in Fig. 2.5. Light propagating up a density gradient, reaches a maximum plasma density of $n_e = n_{cr} \cos \theta$, where θ is the angle of the light to the density gradient. If the light is p-polarised, then the field oscillates in the plane of the gradient, therefore a finite portion of the evanescent field is able to tunnel through to the critical density. This field will resonantly excite an EPW at the frequency of the light wave, which propagates down the density gradient. As it propagates down, the EPW will irreversibly lose energy to the plasma via Landau damping or collisional damping, resulting in laser energy absorption by the plasma.

If the light is normally incident to the density ramp, then the field is purely s-polarised, thus no resonance absorption can occur. However, the light must propagate sufficiently close to normal incidence in order to get close to the critical surface. This places strict requirements on plasma conditions and laser propagation for resonance absorption to be significant compared to Inv-Brem. In long density length scale plasmas, only a very small portion of the incident laser light gets sufficiently close to the critical surface for large portions of resonance absorption to occur. Additionally, Inv-Brem often dominates apart from when it becomes ineffective at high plasma temperatures. It can however be a significant absorption mechanism for short length and timescale, high intensity laser-solid interactions, before the target has significantly heated and expanded. Specifically for ICF however, resonance absorption is thought to be mostly insignificant on OMEGA since the move to frequency tripled lasers [127].

2.5 Laser Plasma instabilities

This section shall provide some introductory theory which is required to understand LPIs. The ponderomotive force is introduced, which is the mechanism by which crossing light waves are able to excite longitudinal plasma waves and thus lead to LPIs. A quantitative picture of these three-wave coupling processes is provided, which describes how the plasma waves are produced and act to scatter the light waves. The specific mechanism for CBET is described and the gain, which is the rate at which each EMW gains or loses energy, is provided in the limit of a fluid description of the plasma. The effect of CBET in ICF experiments is discussed and some further modifications to the simple picture provided are given, such as the Langdon effect of Inv-Brem heating on the CBET gain.

2.5.1 The Ponderomotive Force

When interacting with a wave that has a non-uniform envelope, plasma particles undergo a drift from the regions of high field amplitude to low amplitude, on top of their fast quiver motion. This effect is known as the ponderomotive force and is crucial to the development of LPIs. Although the ponderomotive force acts on particles in both longitudinal and transverse waves in plasmas, this section shall focus only on the effect due to an EMW. Considering an EMW with a non-uniform envelope, polarised along the $\pm \hat{\mathbf{x}}$ axis and propagating along $+\hat{\mathbf{z}}$, the electric field of the wave can be written,

$$\mathbf{E}(\mathbf{x}, t) = E_0(\mathbf{x}, t) \cos(\psi) \hat{\mathbf{x}}, \quad (2.101)$$

where $\psi = kz - \omega t$ and E_0 is the slowly varying envelope. As previously discussed, the particle will undergo fast quiver oscillations due to the $\cos(\psi)$ term. The non-uniform component of the field shall introduce a slow drift motion of the particle if the envelope is slowly varying, which are analogous to single-particle motion ' $E \times B$ ' and ' ∇B ' drifts [78]. Thus, an Ansatz to the particle motion is sought of the form,

$$\mathbf{x}(t) = \mathbf{x}_0(t) + \mathbf{x}_s(t), \quad (2.102)$$

where $\mathbf{x}_0(t)$ is the quiver motion and $\mathbf{x}_s(t) = \langle \mathbf{x}(t) \rangle$ is the slow drift motion of the particle. The brackets here refer to an average over a single oscillation period.

The field amplitude is assumed to be slowly varying and therefore does not significantly change over the scale of a particle quiver oscillation. A Taylor expansion can thus be performed for the field about the cycle-averaged $\mathbf{x}_s$ for small quiver movements, $\mathbf{x}_0$,

$$\mathbf{E}[\mathbf{x}(t), t] = \mathbf{E}[\mathbf{x}_s(t) + \mathbf{x}_0(t), t] \approx \mathbf{E}_s + [\mathbf{x}_0(t) \cdot \nabla] \mathbf{E}_s, \quad (2.103)$$

where,

$$\mathbf{E}_s = \mathbf{E}[\mathbf{x}_s(t), t] = E_{0s} \cos(\psi_s) \hat{\mathbf{x}}, \quad (2.104)$$

and $\psi_s = kz_s - \omega t$. For simplicity, the magnetic field is simply taken to be the equal to the

average field over a cycle, *i.e.* $\mathbf{B} \approx \mathbf{B}_s$. A Taylor expansion of the magnetic field around $\mathbf{x}_s$ yields second order corrections to the final expression. An expression for the cycle-averaged magnetic field is found by integrating Faraday's law, Eq. 2.12,

$$\mathbf{B}_s = \frac{1}{\omega} \nabla \times [E_{0s} \sin(\psi) \hat{\mathbf{x}}]. \quad (2.105)$$

The equation of motion, averaged over the quiver motion for a charged particle, interacting via the Lorentz force with these fields can thus be written,

$$m \frac{d^2 \mathbf{x}_s}{dt^2} \approx q \langle (\mathbf{x}_0 \cdot \nabla) \mathbf{E}_s \rangle + q \langle \mathbf{v}_0 \times \mathbf{B}_s(\mathbf{x}, t) \rangle, \quad (2.106)$$

where the fact that $\langle \mathbf{E}_s \rangle = 0$, as it is averaged over a cycle, has been used to simplify the first term on the left. The velocity of the particle $\mathbf{v} = \mathbf{v}_0 + \mathbf{v}_s \approx \mathbf{v}_0$, because it was assumed that the quiver motion is much faster the drift in the Ansatz, Eq. 2.102. For sufficiently low intensities, the particle velocities will be non-relativistic and therefore the magnetic field will not influence the leading order quiver motion, $\mathbf{x}_0(t)$. The quiver motion is therefore a simple oscillation in the electric field,

$$\mathbf{x}_0(t) = \frac{-q}{m\omega^2} E_{0s} \cos(\psi) \hat{\mathbf{x}}, \quad (2.107)$$

$$\mathbf{v}_0(t) = \frac{-q}{m\omega} E_{0s} \sin(\psi) \hat{\mathbf{x}}, \quad (2.108)$$

where $\mathbf{v}_0$ is the quiver velocity. Substituting in Eqs. 2.105, 2.107 and 2.108 into the equation of motion Eq. 2.105 yields a drift force term [116],

$$m \frac{d^2 \mathbf{x}_s}{dt^2} = \mathbf{F}_p, \quad (2.109)$$

where the ponderomotive force is identified,

$$\mathbf{F}_p = \frac{-q^2}{2m\omega^2} \nabla \langle \mathbf{E}_s^2 \rangle. \quad (2.110)$$

The force term in Eq. 2.109 acts to push charged particle undergoing quiver motion down the gradient of the slowly varying field envelope. One important difference between the ponderomotive force for longitudinal plasma waves and transverse light waves, is that the gradient term in the direction $\mathbf{F}_p \perp \mathbf{E}$ originates from the magnetic field component of the Lorentz force. Longitudinal plasma waves, *i.e.* the IAW and the EPW, do not have a magnetic component and therefore the force is only $\mathbf{F}_p \parallel \mathbf{E}$. Although the force acts on all charged particles, it pushes electrons much more easily than ions as $\mathbf{F}_p \propto 1/m$. This electron movement sets up a charge displacement, which then pulls the ions along with the electron population. Eq. 2.110 was derived by assuming a Taylor expansion of $\mathbf{E}$ in the oscillation displacement from the cycle average, $\mathbf{x}_0$. This is valid if the displacement is not too large, which is true for non-relativistic intensities, *i.e.* the normalised vector potential amplitude $a_0 \ll 1$ [116].




Figure 2.6: Two plane light waves crossing each other and leading to a beat pattern in the summed field. When in a plasma background, the charged particles will experience a drift motion down the beat pattern envelope, known as the ponderomotive force.

2.5.2 Three-Wave Coupling with Two Light Waves

An important example of the ponderomotive force in action for LPIs is the case of two crossing light waves in a plasma, with slightly different wave numbers ($\mathbf{k}_0 - \mathbf{k}_1 = \mathbf{k}_\Delta$) and frequencies ($\omega_0 - \omega_1 = \omega_\Delta$). The waves will thus create a beat pattern which has an envelope with wavevector $\mathbf{k}_\Delta$ and oscillation frequency ω_Δ . Two light waves crossing and creating a beat pattern is plotted in Fig. 2.6. Only the spatial variation is plotted, but the beat pattern also oscillates in time. This beat profile leads to a ponderomotive force term [116],

$$\mathbf{F}_p = \frac{q^2}{4m\bar{\omega}^2} |E_0||E_1| \hat{\mathbf{e}}_0 \cdot \hat{\mathbf{e}}_1 \sin \psi_\Delta \mathbf{k}_\Delta, \quad (2.111)$$

where $\hat{\mathbf{e}}_i$ and $|E_i|$ are the polarisation unit vector and field amplitude of wave i , respectively, and $\bar{\omega} = (\omega_0 + \omega_1)/2$. The phase term is $\psi_\Delta = |\mathbf{k}_\Delta|z - \omega_\Delta t$. The force is thus modulated at the oscillation wavelength and frequency of the beat envelope. This acts to push particles into the beat region with the lowest envelope at a given time and if ω_Δ and $\mathbf{k}_\Delta$ match the dispersion relation for a plasma wave, an EPW or IAW can be resonantly excited. Electron evacuation from the ponderomotive force sets up a modulation in n_e and thus the plasma refractive index. This acts as a grating, which can scatter light from one wave to the other. An unstable feedback loop can thus occur, whereby light is scattered from one wave to another, increasing the magnitude of $|E_0||E_1|$. A stronger ponderomotive force thus increases the amplitude of the grating, scattering yet more light between the waves. This continues until some process saturates the interaction, for example ‘pump depletion’ saturation occurs when the beam transferring energy to the other waves sufficiently depletes its energy that the rate of energy transfer no longer increases. The instabilities CBET, Stimulated Raman Scattering (SRS) and Stimulated Brillouin Scattering (SBS) are all excited via this ponderomotive process.




Figure 2.7: The frequency and wave number matching conditions for LPIs, whereby a pump EMW, labelled 0, loses energy to two daughter waves. Here the daughter waves are another EMW with lower frequency ($\omega_1 < \omega_0$) and an IAW, which are labelled 1 and 2, respectively. EMW $\rightarrow$ EMW + IAW identifies the process as CBET or SBS.

The frequency and wave number matching conditions for a resonant excitation can be expressed in the form,

$$\omega_2 = \omega_0 - \omega_1, \quad (2.112)$$

$$\mathbf{k}_2 = \mathbf{k}_0 - \mathbf{k}_1, \quad (2.113)$$

where ω_2 and $\mathbf{k}_2$ match the dispersion relation for either an EPW or IAW. Wave '0' is often referred to as the 'pump' wave and the other waves are known as the 'daughter' waves. Eqs. 2.112 and 2.113 can heuristically be thought of as conservation of wave quanta energy and momentum, respectively. This is seen more clearly by multiplying both sides by $\hbar$ and noting that the energy and momentum of a wave quanta are $\hbar\omega$ and $\hbar\mathbf{k}$, respectively. This conservation is illustrated in Fig. 2.7, which plots the dispersion relations for CBET or SBS, whereby a pump EMW loses energy to a daughter EMW and IAW.

2.5.3 Cross-Beam Energy Transfer

As stated, CBET is one of the LPIs, which involves the decay of an EMW into another EMW and an IAW. This is almost the same as SBS, but distinct. For SBS, the daughter EMW wave occurs due to light scattering from the plasma, and therefore grows from noise. The scattered light wave is low amplitude, therefore its field can be amplified manyfold via the process. In CBET however, the daughter EMW is another laser beam, *i.e.* two laser beams cross leading to a beat pattern and subsequent energy exchange from one beam to another. This procedure is illustrated in Fig. 2.8 Both light waves are relatively large in field amplitude and therefore a significant amount of energy can be transferred, even when the waves do not quite resonantly drive an IAW [116].

A fluid description of the plasma may be used as a useful starting point to introduce the key physics of CBET and obtain a rate of energy transfer between the light waves, known




Figure 2.8: A 1-D Illustration of the mechanism by which CBET grows. Panel a) plots the incident electric fields from two lasers, overlapping in a plasma background. This creates a beat pattern which ponderomotively pushes charges within the plasma to regions where the beat envelope is small, as is shown in panel b). If the spatial and temporal frequency of the beat oscillation matches the IAW dispersion relation, then an IAW grows, which is shown in panel c). The density modulation from the IAW acts as a grating which transfers energy from one beam to the other.

as the gain. A kinetic plasma description is often required to accurately model the interaction, especially when occurring in multi-species plasmas [128]. While not provided here, the kinetic gain (derived in, for example, Ref. [129]) is presented in Sec. 3.5.1, along with a description of its benefits over the fluid gain. The fluid gain derivation was first calculated in Ref. [130].

The system of equations to solve to calculate the gain rate of CBET may be written as,

$$(\partial_t^2 + \omega_p^2 - c^2 \nabla^2) \mathbf{a}_0 = -\omega_p^2 \frac{\delta n_e}{n_e} \mathbf{a}_1, \quad (2.114)$$

$$(\partial_t^2 + \omega_p^2 - c^2 \nabla^2) \mathbf{a}_1 = -\omega_p^2 \frac{\delta n_e}{n_e} \mathbf{a}_0, \quad (2.115)$$

$$[(\partial_t + \nu + \mathbf{u} \cdot \nabla)^2 - c_s^2 \nabla^2] \frac{\delta n_e}{n_e} = \frac{Z m_e c^2}{m_i} \nabla^2 (\mathbf{a}_0 \cdot \mathbf{a}_1). \quad (2.116)$$

In Eqs. 2.114 and 2.115, $\mathbf{a}_i$ are the normalised vector potentials of the light waves, which undergo an interaction proportional to the density modulation generated by an IAW, $\delta n_e / n_e$. These equations describe the general coupling of two light waves via an IAW. Eq. 2.116 is the driven wave equation for the IAW in a plasma with flow velocity $\mathbf{u}$ and ionisation Z , which




Figure 2.9: The two CBET modes which are identified from the fluid CBET gain in Eq. 2.117, in the context of direct-drive ICF. The backscatter mode (left) dominates for direct-drive ICF. In this regime, the light waves have approximately equal frequencies and therefore the required frequency difference in the plasma frame to excite an IAW is a result of the Doppler shift. In the sidescatter mode (right) which is employed in indirect-drive to tune symmetry [44], the flow is orthogonal to the wave vector and therefore CBET only occurs due to a frequency difference between the waves. This figure has been reproduced with permission from Ref. [131].

experiences damping, ν . Much more detail on the derivation and solution of this system of equations is provided in Ref. [116], however, a final amplification rate is obtained,

$$\gamma_{01} = \frac{n_e e}{4m_e c \omega_0} \frac{1}{T_e (1 + 3T_i / Z T_e)} \frac{R(\eta_{01})}{\nu}, \quad (2.117)$$

$$R(\eta_{01}) = \frac{(\nu / \omega_s)^2 \eta_{01}}{(\eta_{01}^2 - 1)^2 + (\nu / \omega_s)^2 \eta_{01}^2}, \quad (2.118)$$

$$\eta_{01} = \frac{\omega_\Delta - \mathbf{k}_\Delta \cdot \mathbf{u}}{\omega_s}, \quad (2.119)$$

where T_e and T_i are the electron and ion temperatures in eV respectively and $\omega_s = |\mathbf{k}_\Delta| c_s$ is the IAW frequency. The gain, γ_{01} describes the spatial amplification rate per unit path length of beam 0 interacting with beam 1, such that $|a_i|^2 \sim |a_{i0}|^2 \exp(\gamma_{ij} |a_{j0}|^2 d\tau)$. The key assumptions in the derivation of Eq. 2.117 include assuming a linear plasma response, such that the effect of the ponderomotive force can be treated perturbatively. This proves to be a good assumption for the intensities used in direct-drive ICF, which are specifically limited to avoid the non-linear regime where LPs can dominate. Another assumption is that the plasma and pump field are homogeneous over the interaction length. For implementation in a ray-tracing code, this is a good assumption because the gain is applied over ray steps, which are by necessity small compared to gradient length scales of the plasma. Additionally, the interaction is assumed to be in a steady state. This is satisfied for direct-drive ICF on the OMEGA facility because the CBET saturation timescale of ~ 10 ps is shorter than the hydrodynamic time scale of ~ 100 ps. A final assumption was used, that the fields were had parallel polarisation, which shall be discussed in Sec. 2.5.3.1.

Maximal scattering occurs when the resonance parameter, $\eta_{01} = 1$. Eq. 2.119 therefore

identifies two distinct regimes when a resonance can be achieved, which are shown in Fig. 2.9. In a stationary plasma (or when $\mathbf{k}_\Delta \perp \mathbf{u}$), scattering can only occur when the beams have a frequency difference, which is known as the sidescatter mode. Alternatively, if the beams have equal frequency, the plasma flow term in the numerator of Eq. 2.119 ($\mathbf{k}_\Delta \cdot \mathbf{u}$), acts to Doppler shift the frequencies of the laser beams in the frame of the plasma. This results in a frequency difference which can drive an IAW if $|\mathbf{u}| \approx c_s$. Direct-drive coronal plasma flow is supersonic and therefore backscatter CBET occurs close to the Mach-1 surface. The plasma flow in the corona is approximately radially outward, and thus the Doppler shift acts to increase the frequency of the radially inward travelling beam in the plasma frame. From the energy conservation condition in Eq. 2.112, energy is therefore depleted from the inbound laser beam and transferred to the outbound beam. The Mach-1 surface is significantly separated from the region close to critical, where Inv-Brem deposition is important and therefore energy is effectively reflected in direct-drive implosions by backscatter CBET before it is deposited within the plasma. Although all lasers on the OMEGA laser facility have the same frequency, the frequency shift term from Eq. 2.89, means that lasers which have propagated through a different plasma profile have slightly different frequencies. Although small, this shift is often of the order of the numerator terms in Eq. 2.119 and thus must be included when modelling direct-drive CBET.¹⁴ For indirect-drive implosions, the Laser Entrance Hole (LEH) where the lasers cross does not have significant plasma flow. Therefore, in order to tune the symmetry of capsules implosions with CBET, the sidescatter mechanism is utilised, by introducing a finite frequency difference between the inner and outer cones of beams. The frequency difference is experimentally tuned to obtain a round implosion morphology [44].

2.5.3.1 Effect of Polarisation

CBET is affected by the polarisation of the light waves which seed it. Only the parallel polarisation components of the waves create the ponderomotive beat field which sets up the density modulation. In the case of mixed polarisation, the gain rate can simply be reduced by multiplying by a random polarisation factor, which is provided in Eq. 3.27. Polarisation smoothing optics are typically employed on direct-drive facilities which create a mixed polarisation beam, as described in Sec. 1.3.2.1. In reality, however, the Distributed Polarization Rotators (DPRs) on OMEGA do not actually create a single, mixed polarisation laser spot, but two slightly offset sub-beams on the target. CBET models which track polarisation or rays have been developed, which account for the true polarisation angle between two beams [132, 133, 134, 135]. These models have demonstrated that polarised CBET on OMEGA creates a mode-1 which is always in the same direction. This persistent mode-1 was experimentally observed from neutron time of flight diagnostics before its origin was understood [136].

¹⁴Typically, the frequency difference for light which has propagated through a steady state direct-drive corona is on the order of $\Delta\omega/\omega \sim \mathcal{O}(0.1\%)$.

2.5.3.2 Langdon Effect on CBET

As was described in Sec. 2.4.3, Inv-Brem preferentially heats the colder electron population, leading to a super-Gaussian electron distribution function. A non-Maxwellian distribution function has an altered susceptibility to the IAW [137], and therefore Inv-Brem heated plasmas have a modified resonance condition for CBET. Dedicated experiments have been performed which used Thompson scattering to measure the electron distribution function of an Inv-Brem heated, gas jet plasma in order to study how CBET is modified [138]. These experiments demonstrated that the CBET gain was systematically reduced by an altered dispersion relation,

$$\omega = kc_s \left[\frac{3\Gamma^2(3/m)}{\Gamma(1/m)\Gamma(5/m)} \right]^{1/2}, \quad (2.120)$$

where Γ is the Gamma function and,

$$m(\alpha) = 2 + \frac{3}{1 + 1.66/\alpha^{0.724}}, \quad (2.121)$$

and α is defined in Eq. 2.99. This is a small effect in direct-drive experiments on OMEGA due to the relatively low amplitude overlapped field strengths in a direct-drive geometry [67]. However, reduction of the gain via Langdon disturbing the electron distribution function could be one possible explanation for why CBET calculations on the NIF are not predictive and require an artificial clamp on the IAW amplitude [139, 46].

2.5.4 Mitigation of Cross-Beam Energy Transfer

LPIs are significantly detrimental to ICF experiments. For example, direct-drive experiments on OMEGA experience a 50% deposition reduction instantaneously, which particularly occurs late in the implosion at peak incident power [67]. This drastically reduces the coupled energy ablation and ablation pressures which can be achieved. If CBET were fully mitigated, scaling relations demonstrate that the hotspot pressure required for ignition would be substantially reduced [140, 141]. This would enable direct-drive ignition designs at reduced convergence ratios, which would be far more robust to instability growth, broadening the available implosion design-space. Mitigation of LPIs and particularly CBET is an active area of investigation related to direct-drive ICF and there are several distinct approaches which have been explored.

For LPIs to occur, there must be a spatial, temporal and spectral matching between waves. Spatially and temporally, the light waves must be in the same place and at the same time to create the ponderomotive beat pattern that excites plasma waves. The spatial coherence can be reduced by changing the widths of the beams involved in the interaction. For example, a narrower beam profile can be used for direct-drive implosions, which leads to less refracted light travelling away from the target to act as a seed for CBET when crossing other beams [142]. This is explored in more detail in Chap. 4. Although CBET scattering can be significantly reduced by taking this approach, which greatly increases energy coupling, the performance of the implosion becomes limited by the lack of symmetry due to large beam-

modes. An improvement to this approach is to dynamically alter the size of the beams that drive the implosion to match the shrinking critical surface, which is known as zooming [143]. If the beam width shrinks to match the critical radius of the target, then beam-mode growth is not large, which increases energy coupling without loss of stagnation state symmetry [144]. Zooming does require specialised laser optics which are not present on the NIF or the OMEGA laser system however [145], and therefore it is not a mitigation strategy that can be employed on current direct-drive implosion experiments. Temporal coherence can be broken by modulating the laser pulses in time with a high frequency. These pulse types are known as STUD pulses and act to disrupt the growth of LPIs [146].

The majority of current mitigation research focuses on disrupting the spectral coherence required for LPIs to grow. This approach suggests either using laser systems with multiple discrete wavelengths [147, 148, 131], or employing broadband lasers [149, 150], to break the spectral coherence required for CBET and other LPIs. Particularly, bandwidth is thought to be the most promising technology for implementation on a future laser system, as it is possible to almost entirely eliminate CBET and SBS with sufficient bandwidth. Specifically for CBET, with sufficient bandwidth, power transfer can occur ‘backwards’ in the backscatter mode, whereby energy is also transferred from the outgoing light to the incoming light, leading to almost no net energy exchange [151]. Bandwidth on the order of $\Delta\omega/\langle\omega\rangle \sim 1\%$ is found from ray-based modelling to be sufficient to effectively mitigate CBET [152, 153], which is in agreement with wave-based [150] and kinetic [151] modelling. Eliminating CBET and other LPIs would also open up the design space for new implosions, because peak intensities would no longer be limited to $\sim 10^{15} \text{ W/cm}^2$. This would increase the viability of, for example, shock ignition designs, with high intensity ignitor pulses late in the implosion [25, 26, 27]. A single broadband beam ($\Delta\omega/\langle\omega\rangle > 1\%$), known as the FLUX laser is currently under development on the OMEGA laser facility, to explore the effectiveness of bandwidth to mitigate CBET [154]. Bandwidth has the additional benefit that it would reduce laser imprint when passed through Smoothing by Spectral Dispersion (SSD) optics.

3 SOLAS, a 3-D Laser Ray-Trace and Cross-Beam Energy Transfer Model

This chapter describes the SOLAS code, a 3-D laser ray-tracing module implemented in the Radiative-Magnetohydrodynamics (Rad-MHD) code CHIMERA. The chapter begins with a discussion of why ray-tracing is frequently used to model ‘long-pulse’ lasers for High Energy Density Physics (HEDP) experiments and why the standard framework is inadequate to model Laser-Plasma Instabilities (LPIs). It then goes on to describe the ray-trajectory solver, electric-field reconstruction and Cross-Beam Energy Transfer (CBET) components of the model in detail, including discussion of the validity of the model components. The numerical methods are also presented alongside an extensive set of validation problems to verify the implementation.

3.1 Ray-Tracing for Hydrodynamic Simulations of Fusion Plasmas

Nanosecond length ('long-pulse') lasers are often used as an external energy source in the field of HEDP, for example in laboratory-astrophysics experiments [155, 156, 157], equation of state studies [158, 159] or Inertial Confinement Fusion (ICF) implosions [53, 160, 76]. Experimental design and analysis for these experiments must often be supported with fluid Rad-MHD simulations. The laser must therefore be described in these codes by a theoretical framework that is both valid and computationally tractable. The physical processes by which lasers interact with plasma are a result of microscopic couplings between the laser field and particles or plasma waves. The detailed microphysics of these interactions are often studied using Particle in Cell (PiC) codes [161] or wave-based solvers [162], typically for scales and durations of tens of micrometres and hundreds of picoseconds. Coupling these tools directly to multidimensional Radiative-Hydrodynamics (Rad-Hydro) simulations, which are often used for millimetre and nanosecond scales, is usually computationally intractable.

The Geometric Optics (GO) assumptions are often applicable for HEDP laser-plasma interactions and therefore ray-tracing offers a computationally tractable approach to modelling lasers in these experiments. By assuming static Rad-Hydro profiles (which is normally valid for the propagation time of light through the computational domain) a laser beam can be discretized into a bundle of rays. The ray equations can then be integrated along their path to solve for the trajectory of the light, assuming that refraction dominates over diffraction. A discrete amount of power can also be given to each ray. If there is a suitable model for the power-absorption rate, this can also be integrated along the ray to provide an energy source for the plasma. In many laser-driven HEDP experiments, frequency-doubled or -tripled lasers and long scale-length plasmas mean that Inverse-Bremsstrahlung (Inv-Brem) is the dominant deposition process [163, 164]. There are well established models for Inv-Brem that are suitable for implementation in ray-tracing codes, because they only require knowledge of the local plasma conditions, which are easily accessible via interpolation from the Rad-Hydro grid to ray locations [66, 165]. The combination of the ray equations and Inv-Brem deposition is therefore the basis for the vast majority of laser-modules coupled to Rad-Hydro codes.

In laser-driven ICF experiments however, another class of interaction, LPIs, are vitally important to the energetics of the implosion. For example, CBET leads to a zeroth-order correction to the energy deposited in direct-drive experiments at the OMEGA laser facility, reducing coupled power late in the implosion to $\sim 50\%$ [67]. LPIs cannot be included in the simple framework described above for two reasons. Firstly, they are non-linear¹ and secondly, reduced theoretical models of the interaction rely on knowledge of the electric field or intensity of the light [130]. Implementation in a ray-tracing code therefore necessitates a method by which the separate light waves can 'talk' to each other. For example, the CBET code described in this chapter stores information for separate components of laser beams on a common grid, which can then interact via the 'pump-depletion iterations', described in Sec. 3.5.5. Additionally, calculating the electric field or intensity is problematic because this

¹Non-linear in this context means that the interaction involves multiple light and plasma waves.

is not an attribute which can be obtained from a GO ray. Heuristically, rays have an associated power, so an area is required to obtain an intensity. The evolution of a portion of the beam front's area is governed by a first order expansion of the Helmholtz equation, rather than the zeroth order expansion, which is most commonly used in ray-tracing packages for Rad-Hydro codes [119]. The first order expansion introduces an equation for the ray amplitude, which can be solved in a variety of ways and used to obtain the electric field of the light. Sec. 3.4 describes the approach taken to solving for the amplitude of the rays in SOLAS, which is to track the area of a triangle² of rays around a standard GO ray.

For direct-drive ICF simulations, it is also desirable to have a 3-D ray-trace, where rays travel and refract in three dimensions. In some computational direct-drive studies, particularly in 1-D Rad-Hydro simulations, simplified laser models are employed in which rays travel radially toward the target [166]. This simplification can lead to significant deviations from reality, as it neglects any refractive losses, which become increasingly significant late in the implosion as the target converges. Assuming that rays travel radially inward, also leads to deposition occurring closer to the critical surface compared to a true 3-D ray-trace, resulting in an overestimation of the drive. The growth of LPs also depends on vector summations of light and plasma wavevectors, so a 3-D ray-trace is necessary when modelling these effects. Predictive direct-drive simulations therefore necessitate a fully 3-D ray-trace, even when coupled to just 1-D hydrodynamics.

3.2 Existing Cross-Beam Energy Transfer Models

There is a variety of existing computational tools used to model CBET for ICF conditions. These include codes that are used to study the interaction from first principles, reduced models to investigate the effect of CBET when coupled to hydrodynamics and validation tools to test the implementation of these reduced models. A non-exhaustive list of existing codes is listed below, provided as context for the work presented later in this chapter.

3.2.1 Ray-Based Models

The most common tool to study the coupling of CBET effects into hydrodynamics are reduced ray-based models, which use the linear gain theory of Stimulated Brillouin Scattering (SBS), described in more detail in Sec. 3.5. A variety of different codes used to study this. The main difference between the models is the way that the electric field or intensity of the laser light is obtained.

Inverse Ray Tracing (IRT) This is an approach implemented in the IFRIT code, that creates a mapping between points on an initial beam front and arbitrary locations within the plasma [123, 167]. This is in contrast to Forward Ray Tracing (FRT), used in the other ray-based approaches listed below, where discrete points on the beam front are integrated forward to discrete points within the plasma. IRT is an efficient approach for convex, approx-

²For a 2-D ray-trace, a co-planar, pair of rays is used rather than a triangle.




Figure 3.1: 3-D ray trajectories through a spherically symmetric, OMEGA direct-drive scale density profile. The incident rays are plotted in cyan and the reflected rays, which spread out over a large solid angle, are plotted in dark blue. The critical density is represented by the black mesh.

imately spherically symmetric plasma profiles, but cannot deal well with beams that have multiple caustics³, where the FRT approach is better suited. IFRIIT has been coupled to the 3-D Rad-Hydro code, ASTER from Laboratory for Laser Energetics (LLE), and prior to the development of CHIMERA-SOLAS, ASTER-IFRIIT was the only code combination capable of conducting 3-D direct-drive simulations with in-line CBET [67].

Ray Statistics Approach The LASNEX [168], TROLL [169] and DRACO codes [131], developed at Lawrence Livermore National Laboratory (LLNL), Commissariat à l’Énergie Atomique et aux Energies Alternatives (CEA) and LLE respectively, implement field-reconstruction methods, which depend heavily on having many rays per computational cell. LASNEX and TROLL, used mainly for indirect-drive hohlraum simulations, obtain the intensity of light by first propagating rays through the mesh and obtaining the deposited power in each grid cell. The intensity is then obtained from the power in each cell using electromagnetic energy conservation. Obtaining the intensity from deposition means that the field cannot be reconstructed in vacuum regions where no deposition occurs.

DRACO uses a similar approach, but the intensity is estimated by multiplying the ray power by the path length in a cell and dividing by the cell volume, which gives a dimensional estimate for the intensity. Both of these approaches require many rays-per-cell to accurately obtain the intensity, $\mathcal{O}(100)$ [170]. For direct-drive simulations this is extremely computationally expensive, because backscatter CBET dominates over sidescatter. Therefore, the reflected field of each beam must be resolved, and each beam spreads out over a large solid angle after reflecting off a convex density profile, as is demonstrated in Fig. 3.1. This means

³Caustics are defined and discussed in detail in Sec. 3.4.1.

that orders of magnitude more rays are required, compared to approaches which can accurately obtain the field from a single ray per cell.

Another drawback of both of these approaches is that caustics of beams (regions where the amplitude of rays diverge) are not identifiable, and therefore fields cannot be capped to physically accurate, diffraction-limited values. If a significant amount of CBET occurs at caustics, such as in direct-drive ICF, then erroneous global multipliers to CBET gains must be applied which are effectively free parameters that must be tuned to obtain a pre-known reduction in absorbed energy. It is therefore difficult to trust this approach for predictive, pre-shot direct-drive simulations.

Paraxial Complex Geometric Optics (PCGO) In this approach, each ray has an associated Gaussian intensity profile, the width of which is integrated along the ray trajectory [119, 171]. A single ray can be used per cell because each ray therefore has an intensity which is interpolated to the mesh. However, the reconstructed field near caustics is not accurate and the width evolution is only valid for a short propagation distance. This approach was coupled to the CHIC 2-D Rad-Hydro code, but the approach is difficult to extend to 3-D as the implementation relied on interacting only rays whose centroids crossed, which does not occur in 3-D for non-coplanar rays [172].

Neighbouring Rays The BEAMCROSSER code obtains an area for each ray by co-tracing a triangle of neighbouring rays around it that can be converted into a ray amplitude and therefore electric field from electromagnetic energy conservation [147]. Integrating the amplitude along the ray trajectory means that the caustics can be identified and therefore fields in those regions capped to diffraction limited values [173]. Each ray also has an individual field value, and it is therefore less dependent on ray-per-cell statistics than geometric models, such as that used in the DRACO code described above [2]. Sec. 3.4 describes the implementation of this approach into the CHIMERA 3-D Rad-MHD code. A similar approach, simplified by assuming uniform illumination of a direct-drive target, has been implemented into the 1-D LILAC Rad-Hydro code [174, 175]. It had previously however, not been coupled to a multidimensional Rad-Hydro code.

3.2.2 Wave Solvers

Solving Maxwell's equations in a plasma background is a useful tool for the study of LPIs. The main code used in the ICF community that uses this approach is LPSE, which propagates light waves through a prescribed, spatially varying density, temperature and velocity profile in 1-D→3-D [162, 176]. It then solves the nonlinear coupling of electromagnetic waves by allowing first order plasma perturbations (obtained from the ponderomotive beat pattern between light waves), which then feed back into the wave propagation. The perturbative approach limits LPSE to linear problems and the temporal and spatial resolution required to resolve the beat frequency mean that coupling to multidimensional Rad-Hydro simulations is not feasible. However, it is an extremely useful tool to validate implementation of

CBET models, especially for laser caustics where diffraction is important, such as the test case presented in Sec. 3.5.7.2. It can also be used for many other important studies, such as the mitigation of LPIs via laser bandwidth and the effect of beam smoothing techniques on the growth rate of LPIs [177]. The PF3D code is also frequently used to perform multi-dimensional simulations of ICF relevant LPIs [178]. PF3D models laser beams in the paraxial (weakly-refracting) limit and is therefore better suited to indirect drive problems where laser light typically undergoes weaker refraction than for direct-drive in an extended plasma corona.

3.2.3 Particle in Cell Codes

Both ray-based codes and LPSE are ill-suited to the study of LPIs in the non-linear regime, where the laser intensity becomes sufficiently large that the ponderomotive imprint on the plasma can no longer be treated perturbatively. Understanding the growth and saturation of LPIs in this regime is particularly important for ICF schemes with high peak intensities, such as during the ignitor spike in shock ignition pulses [26]. Often kinetic effects such as ion-trapping become important in non-linear saturation, where ions are trapped and then accelerated in the CBET induced Ion Acoustic Wave (IAW), leading to changes in the IAW phase velocity and therefore a loss of resonance [161]. PiC codes operate by using a finite number of macro-particles to represent the distribution function of the plasma, coupled to an electromagnetic solver, discretised on a mesh. They are therefore well suited to model this kinetic saturation, albeit over short timescales and in simulations that are not truly representative of direct-drive conditions, due to computational expense [179]. Kinetic modelling of CBET has demonstrated that the growth of LPIs can be a much more complex, time-dependent problem than is assumed by the linear-models used in ray-based codes, leading to larger net energy transfers [151].

3.3 SOLAS 3-D Laser Ray-Trace

This section describes how the ray equations, derived from the Helmholtz equation in Sec. 2.4.2 are solved in the SOLAS module for the CHIMERA Rad-MHD code. Details of the mesh used for spherical simulations, load balancing options as well as validation problems are provided.

A short summary of the computational derived-types is presented here, to aid understanding of the following sections. In SOLAS, a domain balanced mesh is created from the hydrodynamic grid in CHIMERA. The mesh is made up of *laser-cells*, which hold information about the geometry, hydrodynamic-quantities, laser-quantities⁴ and neighbouring cells on the SOLAS grid. *Beam* objects are created which hold user-defined geometric information about the beam-port and pointing location and also store a list of rays from the beam. Each ray is defined by a position ($\mathbf{x}$), wavevector ($\mathbf{k}$), phase (φ)⁵, angular frequency (ω) and power (P). If

⁴Laser quantities include, for example, Inv-Brem deposition and electric field strength.

⁵Note that the phase φ has units of space. It can be converted to the dimensionless phase, $\psi = k_0 \varphi$, where k_0 is the vacuum wavenumber.




Figure 3.2: A flowchart illustrating the basic interfacing between CHIMERA and SOLAS. The interpolation steps refer simultaneously to the re-domain balancing and cell combination procedures described in Sec. 3.3.2.

field reconstruction is performed, then each ray also stores up to three additional *neighbour-rays*, from which the electric field of the ray can be reconstructed, as described in Sec. 3.4.

3.3.1 Outline of SOLAS-CHIMERA Interfacing

A simple flowchart for the basic SOLAS operational loop and the interfacing with CHIMERA is shown in Fig. 3.2. At the start of a Rad-Hydro timestep, t_n , the hydrodynamic variables (stored on the Eulerian CHIMERA grid), which are necessary to compute ray trajectories and CBET gains, are interpolated to the SOLAS mesh. Rays are then initialised on the beam port locations and extrapolated to the computational domain, before being traced through the grid. SOLAS uses an Message Passing Interface (MPI) domain balanced approach to parallelisation, the reasons for which are discussed in more detail in Sec. 3.3.2, therefore rays are traced up to the internal borders of the MPI rank subdomain. Rays are then passed between ranks and moved through their new subdomain until they exit the entire simulation domain. As rays move through the domain, they lose energy via Inv-Brem, contributing to a running deposition total in each grid cell. When all rays have exited the domain, the SOLAS quantities, such as deposited energy which have been built up from the ray-trace, are then interpolated onto the CHIMERA grid, where the Inv-Brem deposition is used as a source term for the electron energy in the subsequent Rad-Hydro step, t_{n+1} . The interpolation, initialisation and propagation steps are outlined in sections 3.3.2, 3.3.3, 3.3.4 respectively.

3.3.2 SOLAS Mesh Structure

For direct-drive ICF simulations, which model the laser using ray-tracing (especially those which include CBET), the choice of an appropriate computational mesh on which to con-




Figure 3.3: An illustrative diagram demonstrating the MPI re-domain balanced grid employed by SOLAS compared to the CHIMERA domain balanced grid for a cylindrical geometry.

duct the laser ray-trace is crucial to obtaining an accurate and noise-free energy deposition. A well-chosen grid can also significantly reduce the run-time and memory requirements of the calculation. The mesh used for hydrodynamic simulations of these experiments is often not well suited for ray-tracing calculations. For example, hydrodynamic grids often have excessive resolution in the corona to resolve the gradients in the laser quantities, which greatly increases the expense of tracing rays through the mesh, without significantly enhancing the accuracy of the solution. This section describes the approach taken in the SOLAS code to create a grid structure that accurately resolves laser energy deposition, both with and without CBET, in a memory-efficient manner.

3.3.2.1 Re-Domain Balance for Radial Geometries

Rad-MHD codes including CHIMERA, often employ a domain balanced approach to parallelisation, where each computational rank solves a portion of the entire spatial domain, for each discrete timestep. Additional ‘ghost-cells’ are stored in the subdomains and used to calculate gradients on the boundaries between ranks, which are updated with inter-rank MPI communications each timestep. The hydrodynamic grid for CHIMERA is Eulerian, with options for Cartesian, cylindrical or spherical-polar grids.

The optimal domain balancing minimises the number of ghost cells at subdomain boundaries, which leads to cubic⁶ subdomains. If computing a laser ray-trace through this domain for a typical direct-drive calculation however, initially rays travel approximately radially and therefore regularly encounter processor boundaries where they must be transferred between ranks. A more optimal domain decomposition for the ray-trace minimises radial splitting to avoid excessive passing of rays. For spherical and cylindrical simulations, SOLAS therefore

⁶Cubic in cell dimension, not necessarily in physical space for non-Cartesian geometries.

takes the hydrodynamic variables on the CHIMERA grid and re-domain balances the grid for the ray-trace, such that the splitting does not occur in the radial coordinate. An example of this re-domain balancing is shown in Fig. 3.3 for an illustrative cylindrical mesh. The CHIMERA domain decomposition divides the radial and azimuthal extent into four and three ranks respectively, whereas SOLAS' division is purely in the azimuthal direction. For spherical and cylindrical direct-drive simulations where there is a defined, global minimum critical radius, SOLAS' mesh excludes the grid cells below this minimum radius beyond which the rays cannot reach in order to reduce the memory burden of the re-gridding.

An alternative approach to MPI-based domain-decomposition parallelisation of the module is to use the Open Multi-Processing (OpenMP) package, for which ranks share memory across a computational node. In this approach, the entire laser grid is stored once on the shared memory of the node and separate ranks trace rays through the entire domain without the need for transfers. While this is a preferable approach for a standard ray-trace to the MPI procedure described above, including a model for CBET is more challenging. This is because CBET requires communication between beams and therefore large amounts of information must be stored on the grid, which can lead to large memory overheads. Therefore, using multiple computing nodes is often a necessity for 3-D CBET calculations for which OpenMP-MPI hybrid approaches are required. This was deemed too significant an undertaking for the scope of the work presented in this thesis.

When a 3-D ray-trace is coupled to a lower dimensional hydrodynamics simulation (for example, a 1-D spherical direct-drive simulation) then a sparse 3-D mesh at a user-specified resolution is created for the ray-trace. This 3-D mesh allows the re-domain balancing routine to be employed to load balance the simulation. For CBET simulations, it also allows the field, which is required to compute the power change of rays due to CBET, to be stored as a function of angle when discretised on this sparse 3-D mesh. Storing the field discretised in angle enables the interaction between beams around the sphere to be correctly computed. After the ray-trace, if these extra dimensions have been added to the ray-trace grid, then the deposited power is integrated over the additional grid directions to obtain the source term on the hydrodynamic grid.

3.3.2.2 Semi-Structured Eulerian Grid with Combined Cells

Ideal ICF implosions are spherically symmetric and departures from this symmetry are usually higher-order corrections. Simulations of these experiments therefore typically employ a computational grid with spherical symmetry. CHIMERA typically simulates direct-drive implosions with a spherical-polar Eulerian mesh, shown in Fig. 3.4.a. This grid has the advantage of simplicity to take gradients across cells, however it has the disadvantage that cell edge lengths tend to zero as the radial coordinate, $r \rightarrow 0$ and the polar coordinate, $\theta \rightarrow 0, \pi$. This limits the hydrodynamic timestep, because stable explicit timesteps are inversely proportional to edge lengths, increasing the cost of 3-D spherical simulations. The issue is circumvented at late times in the implosion by remapping onto a Cartesian grid, which does not have vanishing edge lengths and face areas [91].




Figure 3.4: Illustrative diagrams demonstrating a) the spherical polar mesh used by CHIMERA for hydrodynamic calculations, b) the cell combination mechanism employed by SOLAS to obtain a roughly equal area grid for spherical simulations and c) the adaptive radial cell combination to reduce resolution in regions where CBET and refraction are unimportant.

Computing CBET requires at least a single ray from each interacting beam to pass through each computational grid cell where the interaction should be important. For spherical-polar meshes the vanishing cell volume therefore sets extreme minimum ray number limits for the calculation to fully resolve all the fields, particularly the reflected field in direct-drive simulations, which spreads out over 4π steradians, as is shown in Fig. 3.1. Hydrodynamic resolutions are often also excessive for ray-trace calculations to resolve the necessary refraction and energy exchange of the light. Rays must also stop at each cell boundary in order to deposit the correct amount of energy into each grid cell and therefore the expense of the ray-trace is directly proportional to the number of grid cells that the rays see.

To circumvent these issues, cells can be combined around the spherical grid angles in order to make a semi-structured Eulerian grid, as is shown in Fig. 3.4.b. Hydrodynamic grid cells are merged together on the SOLAS mesh in each grid direction until a pre-specified res-




Figure 3.5: Example initial locations for ~ 4000 rays for an OMEGA beam port with super-Gaussian intensity profile, with a shape defined by $n_s = 5.2$ and $\sigma = 352 \mu m$. The inverse-projection ray locations are shown in sub-figure a) which demonstrate a higher ray density for rays pointed to the pole at $x \sim -0.15 mm$. b) shows the randomly sampled rays where all rays are initialised with the same power, and c) shows uniformly sampled rays.

olution, set by the user, is reached. In Fig. 3.4.b, many cells are combined in each direction to clearly display the effect. Typically, for spherical simulations however, the cells are not combined in the polar direction, θ and n_φ cells are combined in the azimuthal direction until the azimuthal and polar resolutions match, explicitly $n_\varphi r \Delta\varphi \sin\theta \approx r \Delta\theta$ at a given radius, r .

Fig. 3.4.c demonstrates the capability of the meshing algorithm to adaptively combine cells in the radial direction based upon density gradients. This allows large cells to exist in the coronal plasma where the light refracts minimally and deposits little energy, while the sharp turning point regions close to the critical surface can be well resolved. To find the number of radial cells to combine at a given radius n_r , maximum gradient length scale from the combined hydrodynamic cells, $L_{n_e} = n_{crit}/|\nabla n_e|_{\max}$ is calculated. The algorithm then finds the optimal number of cells to merge together such that the new cell resolution $n_r \Delta r \approx C_L L_{n_e}$, where C_L is a user parameter⁷ which can be reduced to limit cell combining. A maximum cell size is also set to prevent excessively large cells in regions with small density gradients.

Note that many options for equal area computational grids exist which, unlike the semi-structured grid described here, are completely uncorrelated from the Rad-Hydro mesh [180, 181]. These have the advantage that the grid can be chosen purely based upon quantities important to the laser ray-trace. However, the mesh described in this section has the advantage that interpolation to the hydrodynamic mesh is extremely straightforward as the SOLAS grid cells completely overlap the CHIMERA cells. This also minimises artefacts from interpolation between the grids, which have the potential to introduce spurious high order modes to the deposition source term.

3.3.3 Ray Initialisation

To obtain a noise-free energy deposition source term from a ray-trace, both the grid and choice of initial ray number and location is crucial. The problems are closely related as a

⁷Default value set to $C_L = 0.05$.

higher resolution grid requires a larger density of rays to get equivalent ray-per-cell statistics. It is therefore wise to choose the number of rays used in a simulation to be some function of the grid used for the ray-trace. Several methods for ray initialisation which have been implemented are briefly outlined here, along with a summary of their strengths and weaknesses.

All beams profiles described in this thesis have circularly symmetric super-Gaussian intensity profiles described by the equation,

$$I(r) = I_0 \exp\left(-\left|\frac{r}{\sigma}\right|^{n_s}\right), \quad (3.1)$$

where r is the distance from the centre of the beam port, perpendicular to the beam normal, I_0 is the peak intensity, σ is the beam width⁸ and n_s is the super-Gaussian exponent.

Uniform Sampling In this method of ray initialisation, demonstrated in Fig. 3.5.c, each beam is assigned a number of rays and a maximum initialisation radius⁹, beyond which rays are not initialised. Rays are then placed on a uniform, square grid in this plane with a power proportional to the intensity value from by Eq. 3.1. For all three procedures described in this subsection, the total summed power of the rays is normalised to the incident beam power after the initialisation of all rays is complete. Typically, for direct-drive, the total number of rays used for spherical simulations is chosen to be,

$$N_{\text{ray}} = C_N \max(N_r N_\varphi, N_r N_\theta, N_\theta N_\varphi), \quad (3.2)$$

where $N_{i=r,\varphi,\theta}$ is the number of SOLAS grid cells in each grid direction and C_N is a user-parameter multiplier which to give better ray statistics for CBET simulations when required. $C_N \sim 2$ is found to give converged deposition when including CBET and so this value is used by default.

When there is sufficiently low ray-per-cell statistics, beat phenomena can occur between the ray spacing and the grid resolution. In this event, especially when using a nearest neighbour interpolation for ray power deposition onto the grid, significant spurious modes can be introduced to the power deposition. To resolve this, ray locations can be ‘dithered’ so that they take a random position within the polygon¹⁰ defined by neighbouring ray locations. This option is always used for simulations using uniform sampling in this thesis.

Random Sampling Rays can also be randomly sampled according to the intensity profile. Example ray locations from this method are shown in Fig. 3.5.b. Note that the intensity profile purely emerges from the ray locations in this method as all rays have equal power. For no-CBET simulations this is a useful method as it minimises coherent build-up of noise from ray-spacing, grid-resolution beating. However, the wings of the intensity profile have poor ray statistics so for direct-drive CBET calculations, resolving the reflected field and therefore the dominant backscatter CBET is excessively expensive.

⁸The radius at which $I = I_0 e^{-1}$.

⁹The maximum initialisation radius is usually set to be the radius at which $I = I_0 e^{-3}$.

¹⁰This polygon is a square for uniform sampling, but not for inverse-projection

Inverse-Projection Inverse-projection is the final method that has been implemented for ray initialisation in SOLAS. The algorithm is described in detail in Appendix A of Ref. [131]. The method works by finding several surfaces, defined by fractions of the critical density and then creating aim points in each cell on the surface. These aim points and their associated area on the surface are back-projected onto each beam port and rays are created if the aim point is not obscured by the surface on which it was created. The power of the ray is then the back-projected area multiplied by the intensity at the beam port location. For a spherical polar grid, this gives a ray distribution which varies in ray density according to which region on the beam port maps to regions on the grid with higher or lower cell-face areas. The higher ray density region on the beam port maps to the polar region of the grid, at $x \sim -0.16 \text{ mm}$ on Fig. 3.5.a have a correspondingly lower power compared to rays at equal radii to give the same intensity.

While inverse-projection does guarantee good ray statistics at given surfaces for the in-bound component of the field, the algorithm does not extend to beyond ray caustics, or the reflected field component. Therefore, ray statistics are not guaranteed to be good for the reflected field, so the advantages of this method are not evident for direct-drive CBET simulations when backscatter must be resolved. However, for no-CBET, direct-drive simulations, this is a useful option for ray initialisation.

For all simulations in this thesis, the uniform sampling method with ray dithering is used. This method is found to give the best ray-per cell statistics across the semi-structured Eulerian grid described in Sec. 3.3.2 while minimising the overall number of rays used required to resolve CBET.

After rays are initialised on each of the beam ports, they are extrapolated to the edge of the computational domain. Beam focussing is typically neglected, so all rays are assigned velocities parallel to the beam normal. This approximation is widely used for codes that simulate OMEGA-scale direct-drive implosions, where the lasers are approximately collimated on the implosion scale [67, 131]. It is assumed that outside of the simulation domain is vacuum and rays are therefore extrapolated to the edge of the computational domain in straight lines, using a simple root finding algorithm to obtain the intersection of the rays with either a spherical, cylindrical or rectangular domain depending on the simulation geometry.

3.3.4 Equations of Rays and Adaptive Integration

The partial differential equations that are integrated along the trajectory of each ray are,

$$\begin{aligned}
 \frac{d\mathbf{x}}{d\tau} &= \mathbf{k}, \\
 \frac{d\mathbf{k}}{d\tau} &= \frac{1}{2} \nabla \varepsilon(\mathbf{x}), \\
 \frac{d\varphi}{d\tau} &= \varepsilon(\mathbf{x}), \\
 \frac{d\omega}{d\tau} &= \frac{\omega}{2c} \frac{\partial (n_e/n_{cr})}{\partial t}, \\
 \frac{dP}{d\tau} &= -\kappa_{IB},
 \end{aligned} \tag{3.3}$$

where $\varepsilon = 1 - n_e/n_{cr}$ is the dielectric permittivity of the plasma, n_e and n_{cr} are the electron and critical densities, c is the speed of light, t is time, κ_{IB} is the Inv-Brem absorption kernel, given explicitly in Eq. 2.98 and τ is the ray path length. Note that the phase here is defined such that it has the same dimension as path length and therefore has spatial units. The wavevector is normalised such that $|\mathbf{k}| = \sqrt{\varepsilon}$. The optical path length is related to the arc length, ds by the variable change $ds = d\tau \sqrt{\varepsilon}$. In some formulations of the ray-tracing equations, the differential equations are written in terms of the arc length [131, 182], which is equivalent, but more complicated and without benefit. Note that the time, t is not related to the ray path length τ , because an operator split approach is taken to the ray-trace, such that the hydrodynamic variables are frozen and the ray-trace finds the time independent trajectory of the light through these profiles. The path length τ is therefore best seen as a parameterisation of the ray curve.

The frequency shift arises due to light propagating through a time varying refractive index. This causes successive wavefronts to bunch up (blue-shift) or rarefy (red-shift), which occurs for a temporally increasing or decreasing n_e respectively [120]. This effect is only significant for the work presented in this thesis when calculating CBET. For direct-drive calculations at the OMEGA laser facility scale, this results in wavelength shifts of $\Delta\lambda/\lambda_0 \sim \mathcal{O}(0.1\%)$, which is typically negligible for computing ray trajectories. However, the induced wavelength shift is relatively significant when calculating the CBET resonance and in direct-drive can alter the spatial location where scattering is significant [67]. Neglecting time-dependent ionisation effects, which is a good assumption during peak-power for the approximately steady-state direct-drive corona where CBET occurs, the time derivative can be obtained by assuming that n_e is advected with the bulk fluid,

$$\frac{\partial (n_e/n_{cr})}{\partial t} = - \frac{\iiint_V \nabla \cdot [(n_e/n_{cr}) \mathbf{u}] dV}{\iiint_V dV}, \tag{3.4}$$

where $\mathbf{u}$ is the fluid velocity and the integral is over a computational cell volume. Using the Divergence Theorem, this can be rewritten,

$$\frac{\partial (n_e/n_{cr})}{\partial t} = - \frac{\oint_S (n_e/n_{cr}) \hat{\mathbf{n}} \cdot \mathbf{u} dS}{\iiint_V dV}, \tag{3.5}$$

where the integral on the numerator is now over the cell bounding area, with normal $\hat{\mathbf{n}}$ [131].

Two algorithms have been implemented to integrate Eqs. 3.3, an adaptive RK45 algorithm and the Kaiser algorithm. These are described below, alongside the method for ray-deposition-to-cell interpolation and time-stepping considerations.

Adaptive RK45 Algorithm The default ray evolution algorithm in SOLAS is to solve Eq. 3.3 using an adaptive RK45 algorithm with stepsize control [183]. Ray steps are either limited by the error from this algorithm, or by distance to the next impact with a cell face. Evaluations of the right-hand side of these equations employ trilinear interpolation for n_e , ∇n_e and T_e to obtain varying values for these quantities at different locations throughout the cell. All other quantities are either defined at the ray location, or use nearest neighbour interpolation. The `bspline-fortran` library has also been implemented to allow tricubic interpolation of n_e , which yields a quadratically varying ∇n_e across the grid [184]. This is slow to evaluate many times per ray step as is required in an RK45 algorithm and therefore limits performance, but is used throughout this chapter as a higher order solution to validate accuracy of the default interpolation.

Linear interpolation of n_e and T_e is required for accurate computation of κ_{IB} , especially for ‘cold-start’ simulations where light impacts upon a solid target and the plasma profiles have steep gradients. Interpolation of ∇n_e was found to be necessary to obtain low-noise ray amplitudes from neighbouring rays for CBET evaluation, which is described in more detail in 3.4. Linear interpolation of both n_e and ∇n_e is technically inconsistent, because ∇n_e should be the gradient of n_e , however the test cases presented in Sec. 3.3.4.1 found that this did not reduce accuracy compared to the Kaiser algorithm, which employs a self-consistent linear interpolation of n_e and constant ∇n_e within a computational cell.

Kaiser Algorithm The Kaiser algorithm for ray integration has also been implemented [182]. In this algorithm, each cell has a single ∇n_e and an n_e value defined at the cell centre, allowing for linear interpolation of n_e throughout the cell. In a constant ∇n_e region, the trajectory of rays can be analytically shown to follow a parabola. The Kaiser algorithm thus evolves rays over parabolic segments between cells by finding the intersection of the parabola with cell faces, using a root finding algorithm. Assuming constant ∇n_e cells leads to discontinuities in n_e at cell interfaces when the true gradient is not linear and thus Snell’s Law is used to refract rays upon cell exit. While this method is often slightly more efficient compared to the adaptive RK45, it was found that when attempting to reconstruct the ray amplitude from the area of neighbouring rays, as outlined in Sec. 3.4, excessive noise was introduced in the amplitude from the discontinuous refraction. Therefore, the adaptive algorithm was used for all work in this thesis, other than where explicitly outlined.

Shape Functions for Ray Deposition Nearest neighbour interpolation of power deposition, which is defined along the ray trajectory, onto the computational grid can result in significant levels of noise in the power deposition profile when ray statistics are not sufficiently high. This is similar to problems experienced in PiC codes when interpolating macro




Figure 3.6: Shape function smearing for ray power interpolation to the SOLAS grid for a cylindrical mesh. The shape function has the same size and geometry as the SOLAS grid cell which the ray is located in. A top-hat shape function is used, so the fraction of the deposited power interpolated to a given cell is proportional to the volume overlap of the shape function with the cell.

particles to the grid without the use of shape functions [185, 87]. A PiC-inspired shape function approach was therefore taken, where the power deposition from rays is smeared across neighbouring cells. A top hat shape function is used, with a shape function that has the size and shape of the cell that the ray is located within, as is shown in Fig. 3.6. The power deposited into each cell that has an overlap with the shape function bounds is therefore proportional to the volume overlap of the shape function with the cell. This approach is more explicitly outlined, particularly for the case of the non-logically rectilinear grids which are present for SOLAS, in Ref. [186]. This smearing is somewhat numerically diffusive and therefore can be optionally disabled. For the work presented in this thesis, the shape function smearing is employed for multidimensional direct-drive simulations, but otherwise disabled.

Hydrodynamic Time Step Limiter Laser heating of a plasma results in a temperature increase of the material, which also reduces the Inv-Brem kernel. If the hydrodynamic timestep, δt is too large, excessive heating occurs in initially cold cells, which should have the magnitude of absorption gradually reduced as they warm up. This can lead to spurious oscillations in temperature at fine grid resolutions. A limiter on the hydrodynamic time-step is therefore introduced to prevent large heating in any cell within one timestep. The limit detailed by Haines in Ref. [187] is enforced, which estimates the increase in temperature of a cell from Inv-Brem heating by assuming an ideal gas equation of state to derive,

$$\delta t = \min_{n \in \text{cells}} \left(\tilde{\delta t}, C_{\delta t} \frac{3}{2} \frac{n_{e,n} e T_{e,n} V_n}{P_n} \right), \quad (3.6)$$




Figure 3.7: Results of the quadratic trough ray trajectory test problem for the adaptive solver with 3 different cases. These are default settings (linear interpolation of both n_e and ∇n_e), the Kaiser algorithm (linear interpolation of n_e and uniform ∇n_e in a cell) and the adaptive algorithm with tricubic interpolation of n_e). The top plot shows ray trajectories and analytic trajectory and the bottom plot shows absolute errors for the 3 cases.

where $\tilde{\delta t}$ is the stable hydrodynamic timestep from other physical processes, $n_{e,n}$, $T_{e,n}$, V_n and P_n are the electron density, electron temperature, volume and deposited power into cell n , and $C_{\delta t}$ is a user defined stability parameter. By default, the parameter $C_{\delta t}$ is set to 1, which is found to be adequate for most situations [187]. Note that Eq. 3.6 is only applied in cells where $n_{e,n} > 1 \times 10^{-4} n_{cr}$. Small amounts of deposition in these cells can lead to large temperature increases, which severely limits δt , despite the cells often being far from the region of interest, normally near critical, where the maximum deposition occurs. For simulations where a laser heats a solid, initially cold target, Eq. 3.6 mostly only affects the hydrodynamic timestep early in the simulation, after which the laser-heated plasma corona remains in a mostly steady-state unless there are steep changes in incident power.

3.3.4.1 Ray Solver Validation

Several validation problems have been conducted to verify that the ray solver has been implemented correctly. Here the quadratic trough and cylindrical helix test problems are presented, which compare the path of a single ray to an analytic solution in order to verify that the ray solvers correctly obtain the trajectory of light. The blast wave problem is also presented which is a test of Inv-Brem absorption to an analytic solution in the absence of thermal conduction and hydrodynamic motion.

Quadratic trough The quadratic trough is a test of ray trajectory in a quadratic density trough, which admits an analytic solution of a periodically oscillating ray [182, 187]. The

density profile used for the test is defined as,

$$n_e(y) = \frac{n_{\text{cr}}}{2} \left(1 + \frac{y^2}{y_c^2} \right), \quad (3.7)$$

which for light with wavelength $\lambda = 351$ nm, is initialised with $n_{\text{cr}} = 9.049 \times 10^{21} \text{ cm}^{-3}$ and $y_c = 5$ cm. The domain has bounds $x \in [0, 25]$ cm and $y \in [-5, 5]$ cm and a ray enters the domain at $[x_0, y_0] = [0, 3]$ cm. Analytic integration of the first 2 lines from Eq. 3.3, yields the analytic trajectory as a function of ray path length, τ ,

$$\begin{aligned} x(\tau) &= \tau \sqrt{1 - \frac{n_e(y_0)}{n_{\text{cr}}}}, \\ y(\tau) &= y_0 \cos \left(\frac{\tau}{\sqrt{2} y_c} \right). \end{aligned} \quad (3.8)$$

The analytic trajectory is compared to the solution from the adaptive solver (using both default interpolation and tricubic interpolation of n_e) and the Kaiser algorithm in Fig. 3.7. The associated error, defined as the difference in y from the analytic value at a given x is also plotted. Note that all results were obtained using a grid resolution of 100×100 cells. The top panel demonstrates that all trajectories are identical to the analytic solution by eye. The corresponding errors show that the tricubic interpolation effectively perfectly recreates the analytic solution and the Kaiser error is slightly more significant compared to the default adaptive error. The tricubic error is insignificant because the cubic interpolation perfectly recreates the true density profile, so the errors are purely numeric, not due to the resolution. Errors from the other algorithms would therefore decrease more quickly with increasing resolution as the density profile the ray sees becomes more similar to a quadratic trough. The error from the default adaptive interpolation (linear interpolation of both n_e and ∇n_e) is lower than that of Kaiser (linear n_e and uniform ∇n_e), indicating that the inconsistent interpolation is not a significant issue for resolving ray trajectories.

Cylindrical-Helix The quadratic trough test ensures that the ray evolution algorithm functions correctly in Cartesian geometry. An additional test was desired in non-Cartesian geometry, because CHIMERA also operates with cylindrical and spherical grids. The cylindrical-helix test was therefore devised, where a ray enters an axially symmetric density trough that keeps the ray at a constant, cylindrical radius, r , as is shown in Fig. 3.8.b. As is shown in Fig. 3.8.a, the z extent of the simulation is chosen such that the ray has performed one complete helical spiral when it exits the domain. By requiring that a ray undergo circular motion in the $x - y$ plane and specifying the initial ray direction such that $k_{z0} = \sqrt{k_{x0}^2 + k_{y0}^2}$, it can be derived that a ray entering the domain at r_0 will undergo one complete helical spiral over a z length $\Delta z = 2\pi r_0$, if the density profile has the shape,

$$n_e(r) = n_{\text{cr}} \left(1 - \exp \left(\frac{-r^2}{2r_0^2} \right) \right). \quad (3.9)$$




Figure 3.8: Results from the cylindrical-helix test problem, obtained using the default adaptive RK45 algorithm. a) shows the trajectory of a ray from a simulation with 100 radial cells in 3-D space. b) shows the trajectory of the ray projected on the $x - y$ plane, along with the n_e profile. c) shows the error of the test, defined as the fractional difference in radius of the ray as it exits the domain from the initial radius, as a function of number of radial cells.

For a ray entering the domain at $[x_0 = 0, y_0 = 1, z_0 = 2\pi]$ mm, with initial wavevector $[k_{x0} = \sqrt{\epsilon_0/2}, k_{y0} = 0, k_{z0} = -\sqrt{\epsilon_0/2}]$, the ray position as a function of path length τ is,

$$\begin{aligned} x(\tau) &= r_0 \sin\left(\sqrt{\frac{\epsilon_0}{2}} \tau\right), \\ y(\tau) &= r_0 \cos\left(\sqrt{\frac{\epsilon_0}{2}} \tau\right), \\ z(\tau) &= z_0 - \sqrt{\frac{\epsilon_0}{2}} \tau, \end{aligned} \quad (3.10)$$

where $\epsilon_0 \equiv \epsilon(r = r_0)$.

This problem was run in cylindrical geometry for a variety of radial resolutions, employing the default adaptive RK45 solver with linear interpolation of both n_e and ∇n_e . The domain bounds for the problem were $r \in [0, 1.5]$ mm, $\varphi \in [0, 2\pi]$ and $z \in [0, 2\pi]$ mm, with a single cell in both the φ and z directions. For all radial resolutions, the ray was observed by eye to exit the domain at the expected location, verifying that the ray-trace is functional in non-Cartesian geometries. Plotting the fractional error in ray exit location as a function of number of radial cells, N_r in Fig. 3.8.c, demonstrates that the error in final position of the ray scales as N_r^{-1} . This is because larger N_r more closely recreates the density profile in Eq. 3.9. The error scaling is dictated by the lowest order interpolation used in the algorithm. This is the linear interpolation of n_e for the default adaptive RK45 solver, which leads to the N_r^{-1} scaling.

Blast wave The blast wave problem is a validation test of the implementation of Inv-Brem absorption [188, 187]. In this problem, a 527 nm, 2 ns laser is incident on a uniform density ($\rho = 1 \text{ mg/cm}^{-3}$), cold gas ($T_{e0} = 1 \text{ eV}$) with fixed ionisation at $Z = 6$, that has an ideal gas equation of state. The $\ln \Lambda$ value used for the κ_{IB} coefficient is fixed at 7. The laser travels in




Figure 3.9: Results of the blast wave Inv-Brem absorption test problem. The colour-plot shows the temperature of the 1-D simulation as a function of time (x -axis) and space (y -axis). The heatfront obtained from CHIMERA, which is the maximum x location where $T_e > T_{e0}$, is in good agreement with the analytic solution.

the x direction and has an intensity of $6.4 \times 10^{13} \text{ Wcm}^{-2}$. Hydrodynamic motion and transport is disabled in the simulation. The laser is initially strongly absorbed in the gas, so energy is not transported beyond the initial layer, perpendicular to the laser propagation direction. As the gas heats up, κ_{IB} decreases and therefore more energy is transported further into the domain by the laser, moving the heatfront forward. The heatfront $x_{\text{analytic}}(t)$, defined as the largest x coordinate at a time, t where $T_e > T_{e0}$, has an analytic solution, derived by Denavit and Phillion in Ref. [188],

$$x_{\text{analytic}}(t) = \frac{2}{3\kappa_{\text{IB}}} \left(\frac{5}{3} \frac{\kappa_{\text{IB}} I}{n_e k_B} \right)^{3/5}, \quad (3.11)$$

which can be compared to the result from a ray-tracing simulation.

Fig. 3.9 shows the results for $T_e(x, t)$ and the heatfront location from a CHIMERA-SOLAS simulation of the blast wave problem. The heatfront from the simulation compares well to the analytic heatfront, validating the implementation of the Inv-Brem absorption kernel.

3.4 Ray-Based Field Reconstruction and Ray Sheets

In order to compute CBET, or other LPIs in a ray-tracing calculation, additional information about the electric field or intensity of the light is required than can be obtained from the ray directly. This section outlines the method used in SOLAS to obtain the electric field of the light along the path of each ray, which broadly follows the implementation from Ref. [2].




Figure 3.10: A 3-D illustration of the ray-amplitude construction for a single ray in a spherically-symmetric, direct-drive like plasma profile. The critical surface is shown as the grey wire mesh. The main-ray trajectory is plotted in red, and its three neighbour-ray trajectories are in blue. The area of the main-ray at 5 successive phase-fronts is plotted, *i.e.* when the main- and neighbour-rays all have the same phase, φ , is illustrated by the shaded triangles. This is inversely proportional to the ray-amplitude at this location.

3.4.1 Ray-Amplitude and Field Estimate from Neighbour-Rays

Recalling the formula for the ray-amplitude, A , Eq. 2.86 from Sec. 2.4.2,

$$A(\tau) = A(0) \left| \frac{D(0)}{D(\tau)} \right|^{1/2}, \quad D(\tau) = \begin{bmatrix} \frac{\partial x}{\partial \zeta_1} & \frac{\partial x}{\partial \zeta_2} & \frac{\partial x}{\partial \tau} \\ \frac{\partial y}{\partial \zeta_1} & \frac{\partial y}{\partial \zeta_2} & \frac{\partial y}{\partial \tau} \\ \frac{\partial z}{\partial \zeta_1} & \frac{\partial z}{\partial \zeta_2} & \frac{\partial z}{\partial \tau} \end{bmatrix}, \quad (3.12)$$

where $[x, y, z]$ and $[\zeta_1, \zeta_2, \tau]$ are the ray real-space and phase-space coordinates respectively and D is the Jacobian for the coordinate transform from phase-space to real-space. When refraction forces rays closer together, $|D(\tau)|$ will decrease and equally, when refraction forces rays to separate, $|D(\tau)|$ will increase. The amplitude can therefore be approximated by assuming proportionality between the determinant of the Jacobian and the area, S of an infinitesimally small bundle of rays surrounding the main-ray,

$$A(\tau) = \varepsilon^{-1/4} \sqrt{\frac{S(\tau=0)}{S(\tau)}}, \quad (3.13)$$




Figure 3.11: A demonstration of how the ray-amplitude is obtained from the area of neighbouring-rays for a beam propagating up a linear density gradient. Fig. a) explicitly shows the trajectory of a main-ray and its neighbour-rays in red and blue respectively, explicitly highlighting the beam's caustic region, where the amplitude of the light is capped. Fig. b) plots the area of the main-ray in blue (which goes through a minimum at the caustic location), as a function of the phase of the ray. Also shown are the uncapped and capped amplitudes of the ray.

where the $\varepsilon^{-1/4}$ term accounts for swelling of the field due to the increased optical path through plasma with finite density [2]. Note that the amplitude is a purely geometric quantity, which holds no information about power changes of the light due to absorption or LPIs. The electric field outside the caustic region, which is discussed in more detail in Sec. 3.4.2, can then be obtained with the formula,

$$\frac{|E(\tau)|}{|E(\tau=0)|} = \sqrt{\frac{P(\tau)}{P(\tau=0)}} A(\tau), \quad (3.14)$$

where P is the power of the ray and the initial (vacuum) electric field $|E(\tau=0)| = \sqrt{2I_0/c\varepsilon_0}$, where I_0 is the intensity of the beam at the ray initialisation point on the beam port.

This method, which has been used successfully to model direct-drive CBET in the BEAM-CROSSER post-process code, is the approach taken to estimate the field in SOLAS. Explicitly, for every ray that initialised for a SOLAS calculation, an additional 3 'neighbour-rays' are initialised in an equilateral triangle around it, perpendicular to the initial direction of propagation. Note that when a 2-D ray-trace is used, only two neighbour-rays are required as the rays cannot change separation in the out of plane direction. It is important to distinguish the dimension of the ray-trace, which is the number of dimensions in which rays can move in a simulation, to the dimension of the hydrodynamics. For example, a 1-D spherical direct-drive Rad-Hydro simulation still requires a 3-D ray-trace where rays can move in x , y and z for full accuracy.

These neighbour-rays are then co-traced up to the phase of the 'main-ray' to evaluate the area of the main-ray on its phase front. Fig. 3.10 plots an example of this co-tracing procedure for a single ray propagating in a 3-D direct-drive profile. It can be seen that as the re-

flects away from the critical surface, the area diverges. Note that artificially large neighbour-ray separation was used for this plot as it is purely illustrative. Practically, the separation distance of the bundle of neighbour-rays (i.e. the side-length of the triangle, Δl) does not need to be infinitesimally small, but small enough that it accurately captures the local change of $|D(\tau)|$. In terms of the hydrodynamic profiles, this means that $\Delta l \ll L_{n_e}$, where L_{n_e} is the density length scale of the plasma. For OMEGA scale direct-drive simulations, an initial separation of $\Delta l \lesssim 1 \mu\text{m}$ is found to give converged behaviour. If the separation is too small however, floating point arithmetic can lead to noise in the ray-amplitude estimation. For all simulations presented in this thesis, $\Delta l = \lambda_0/10 \ll 1 \mu\text{m}$ was used.

As mentioned in Sec. 3.3.4, it was found that the Kaiser algorithm for ray propagation, which uses a single ∇n_e in each grid cell and discontinuously refracts rays at cell boundaries, led to large levels of noise in the ray-amplitude profiles. In order to obtain smooth profiles for A , linear interpolation of ∇n_e was found to be necessary. The continuously varying ∇n_e , which does not have discontinuities across cell interfaces, allows a smoothly varying area for each ray to be obtained.

Fig. 3.11.a demonstrates this amplitude reconstruction process in a purely illustrative example of a beam travelling up a linear density gradient at an angle, leading to a caustic at its turning point. The trajectory of several main-rays are plotted in dark-dashed red, and a single main-ray is highlighted in a brighter shade of red. The two neighbour-rays of this highlighted main-ray are plotted in blue. Note that in this figure, the default initial main-neighbour-ray separation of $\Delta l = \lambda_0/10 = 0.035 \mu\text{m}$ was used, but the separation has been magnified on the plot to visualise the amplitude reconstruction from the area more clearly. The locations of the main- and neighbour-rays after a single step of each main-ray are plotted as black and dark-blue spots respectively. The ‘area’ of the neighbour-rays S , which in this 2-D ray-trace example here is simply a distance, is plotted on 3.11.a as a dashed blue line between the neighbour-rays. The neighbour-rays are pushed until they have the same phase as the main-ray, at which point the area is evaluated.

3.4.1.1 Caustics and Ray-Sheets

Caustics are an important concept in GO, and the reconstructed field in their vicinity must be handled correctly in order to accurately model CBET in ray-based codes. Fig. 3.11.b plots the area of the main-ray as a function of main-ray phase in blue. At $\varphi_{\text{ray}} \sim 12 \mu\text{m}$, $S \rightarrow 0$, i.e. neighbouring-rays cross over each other, which is known as the ray-caustic. In the absence of caustics, rays never cross over each other and therefore the mapping from ray phase-space (beam-port locations, where $\tau = 0$) to real-space is single-valued, or in other words each point in space will be reached by (at most) one ray from each beam. In the presence of caustics, this single-valued mapping breaks down and a single real-space location is accessible by multiple ray launch locations. A single-valued mapping is still possible if each beam is separated into distinct ‘sheets’ after caustics, where each sheet still has a single-valued projection from phase-space to real-space.

This is illustrated in Fig. 3.12.a, where rays from a single beam are plotted propagating up




Figure 3.12: An illustration of the concept of caustics and ray sheets. In a), the trajectory of rays traversing up a linear density gradient are plotted. Separate colours are used for the incident and reflected sheets, which are separated by a caustic. In b), rays propagate up a density profile, $n_e = n_{cr} \exp [-(r_{\mu\text{m}} - 20)/100]$. Unlike in a), the ray caustics do not occur at their turning point, *i.e.* the minimum radius.

a linear density gradient with $n_e = n_{cr} y_{\mu\text{m}}/15.5$. As the density gradient is purely in the y direction, each ray has translational symmetry in x and therefore each ray-caustic is located at the ray turning point, $y \sim 14 \mu\text{m}$. The portions of the ray trajectories that fall before and after the caustic are separated into the ‘incident’ and ‘reflected’ sheets respectively and plotted in separate shades of blue. Fig. 3.12.b also plots rays from a single beam, but now propagating through a cylindrically symmetric n_e profile. As can be seen from this figure, the caustic is not the same as the ray turning-point¹¹, but it is the location that the amplitude diverges.

Unless the amplitude is tracked along each ray trajectory, the location of the caustic cannot be identified and rays cannot be separated into distinct sheets, both of which have important ramifications for ray-based CBET models. Firstly, because the amplitude diverges, the field values can become large and therefore must be capped to diffraction limited values (the SOLAS methods for this are discussed in more detail in Sec. 3.4.2). The power change of rays due to CBET scales exponentially with field strength squared, so erroneously large field values are extremely problematic to accurately computing power changes. Additionally, if the beam cannot be separated into distinct sheets, then self-CBET, where the reflected component of a beam interacts with the incident component of the same beam, must be neglected. In SOLAS, the neighbour-ray method allows the amplitude to be tracked along each ray, enabling caustic location identification and sheet-separation. Currently, it is assumed that each sheet has two components, *i.e.* an incident and reflected field, but it should be possible to extend this to multiple caustics to accurately model complex CBET interactions such as target-stalk simulations [189, 190, 59]. This is an advantage of the forward ray-tracing approach to CBET modelling, compared to inverse-ray-tracing, where only plasma profiles that result in a single caustic can be modelled [67].

The phase of the light is also evolved along each ray trajectory in equations 3.3, and therefore the coherent sum of all electric fields can be reconstructed which involves interference

¹¹Defined here as the location of maximal n_e experienced along the ray trajectory.

between sheets. The total electric field is,

$$E = \sum_j^{\text{sheets}} |E_j| \exp(i(k_0 \varphi_j - \pi \alpha_j / 2)), \quad (3.15)$$

where k_0 is the vacuum wavevector of the light and α_j accounts for the $\pi/2$ phase shift that occurs when a ray changes sheet, *i.e.* $\alpha_j = 0$ for the incident sheet and $\alpha_j = 1$ for the reflected sheet [2].

The fields are discretised on the SOLAS mesh by nearest neighbour interpolation from the ray to the cell. When multiple rays from the same sheet pass through a cell, then the inverse-distance weighted average of the mid-point along the ray step to the cell centre is used to compute the field value for the cell. This leads to a field structure, $|E_j|(\mathbf{x}, \mathbf{k}, \omega, \phi)$, where j refers to each sheet, $\mathbf{x}$ is the spatial location of the field which refers to the cell the field is stored in, and $\mathbf{k}, \omega, \phi$ are the wavevector, frequency and phase of the ray. For many beam simulations, there is a large associated memory cost to store this information on the mesh, which is discussed in more detail in Sec. 3.5.4. When computing the effect of CBET, it is also necessary to interpolate the fields stored at the cell centre to the ray trajectories. Nearest neighbour interpolation is also used for this step, so rays experience a constant value for $|E_j|(\mathbf{x}, \mathbf{k}, \omega, \phi)$ throughout the cell. It is technically possible to use higher order interpolation for both stages, either from an unstructured mesh of the ray locations, or from the SOLAS mesh, which is close to logically rectilinear. The second of these options was tested but proved both to be the dominant computational cost for many beam CBET interactions and also challenging to robustly implement for non-Cartesian meshes, because the angular resolution for typical spherical direct-drive simulations was too coarse for the required accuracy.

3.4.2 Caustic Field Capping

In the vicinity of beam-caustics, the electric field magnitude of a beam is limited by diffraction which is neglected in GO. Therefore, ray-based models which do not cap fields near caustics can experience divergent electric fields. For direct-drive simulations, a significant amount of CBET occurs near caustics and therefore sensibly capping the field value is crucial for accurate simulations [167]. This cannot be achieved without knowledge of the caustic location, which is a severe limitation to the predictive capability of CBET codes that do not track the amplitude. This section describes two methods that have been implemented to cap the field values in the vicinity of laser caustics, the Field Limiter (FL) and Etalon Integral (EI) approaches.

Field Limiter This is a straightforward hard cap on the maximum value of the electric field reconstructed during the ray-trace. The electric field of light propagating up a linear density gradient has the analytic solution of an Airy function. The field limiter approach is to cap all

electric fields to this value,

$$\begin{aligned} \frac{|E|}{|E(\tau=0)|} &= \sqrt{\frac{P(\tau)}{P(\tau=0)}} \min \left[A, \sqrt{\zeta} \left(\frac{n_t}{n_{cr}} \right)^{1/4}, \sqrt{\zeta} \sqrt{\frac{S(\tau=0)}{S(\tau)}} \right], \\ \zeta &= 0.9 (\omega L / c)^{1/3}, \\ L &= n_{cr} / |\nabla n_e|_t, \end{aligned} \quad (3.16)$$

where n_t is the electron density at the caustic and $|\nabla n_e|_t$ is the magnitude of the electron density gradient at the caustic location [175, 173, 191]. The first term in the minimum is the standard field reconstruction from the amplitude away from caustics. The second term is the maximum of the Airy function and the third term is an improvement to the second term for near-normally incident rays [2]. Fig. 3.11.b plots the uncapped and capped amplitudes obtained from the field limiter approach for light propagating at an angle to a linear density gradient.

Etalon Integral The EI method is an improvement to the FL, which allows for deviations in the density profile away from linearity [123, 192, 193]. Several distinct types of caustic exist for different geometries of problem. For example in direct-drive, the light reflecting from the critical surface forms a ‘fold’-type caustic, whereas focussing light leads to a more complex ‘cusp’-type caustic [194]. If the form of the caustic is assumed, then an approximation to the total field in the vicinity of the caustic is formulated, which allows for deviations from the ideal case. In the example of the direct-drive fold caustic, the field is assumed to be the sum of an Airy function and its derivative, which accounts for the deviations from linearity. An expression for the total field can be derived in terms of the ray-amplitude and phase before (A_1, φ_1) and after (A_2, φ_2) the caustic,

$$\begin{aligned} E_T &= \sqrt{\pi} [(-\xi)^{1/4} (A_1 + A_2) \text{Ai}(\xi) - i(-\xi)^{1/4} (A_1 - A_2) \text{Ai}'(\xi)] e^{i(k_0 \chi - \pi/4)}, \\ \chi &= \frac{1}{2} (\varphi_1 + \varphi_2), \\ \xi &= - \left[k_0 \frac{3}{4} (\varphi_2 - \varphi_1) \right]^{2/3}, \end{aligned} \quad (3.17)$$

where Ai and Ai' are the Airy function and its derivative. In SOLAS, Ai and Ai' are computed numerically using the special-functions library [195]. This total field can then be divided between the incident and reflected sheet by inversion of Eq. 3.15 assuming that $|E_1| = |E_2|$,

$$\frac{|E_j|}{|E_{j,0}|} = \frac{\sqrt{W_j} |E_T|}{\sqrt{2} [1 + \sin(\varphi_2 - \varphi_1)]^{1/2}}. \quad (3.18)$$

The assumption of equal fields is approximately valid for OMEGA scale direct-drive conditions, where the caustic region is small compared to the plasma scales and therefore ray powers are assumed to not vary significantly over this distance. The Etalon integral is ap-

plied in the ‘caustic-region’, which in Refs. [196, 197, 123], is stated to be equivalent to,

$$|\varphi_1 - \varphi_2| \leq \lambda_0/2, \quad (3.19)$$

where λ_0 is the vacuum wavelength of the light.

Applying the EI method requires evaluation of Eq. 3.19 to determine if a ray is inside its caustic region, which relies on interpolation of the phase of the other sheet from the same beam onto the ray location. If using nearest neighbour interpolation, such as in SOLAS for interpolation of all field quantities onto rays, a grid resolution of $\Delta \ll \lambda_0$ must be used, which is a much higher resolution than is required for ray-tracing and CBET without the EI. Therefore, the EI method is only used in SOLAS simulations for high resolution grid test problems, such as those Sec. 3.4.3. The alternative approach is to implement linear interpolation from the field quantities of all beams onto ray locations, such as in Ref. [2]. This was not deemed a viable approach however, because the mesh resolution for spherical simulations was found to not be well-enough angularly resolved to interpolate the phase with sufficient accuracy to the rays. The standard procedure for CBET simulations in SOLAS is to use the FL approach, which was found to give satisfactory results compared to the EI method.

3.4.3 Field Reconstruction Validation

In this section, several test problems will be presented that compare electric fields obtained using the SOLAS field solver in the absence of CBET, to those from the wave based solver LPSE in direct-drive relevant plasma profiles. The LPSE results are presented in Ref. [2], and are available from the repository [198]. In all the following test problems, the laser has a vacuum wavelength, $\lambda_0 = 0.351 \mu\text{m}$ and a super-Gaussian intensity profile, defined in Eq. 3.1. All 1-D and 2-D validation problems assume lasers polarised out of the simulation plane along the z axis. The plasma density profiles are given by,

$$n_e(r) = n_0(d/r)^m, \quad (3.20)$$

where $n_0/n_{\text{cr}} = 1.165$, $d = 343 \times S \mu\text{m}$, $m = 3.78$ and S changes the scale of the density profile, where $S = 1$ corresponds to scale lengths from implosions using the full OMEGA laser-energy. These values were obtained from fitting the equation to 1-D Rad-Hydro calculations using the LILAC code.

3.4.3.1 1-D Reflected Beam

The first validation problem is for the case of light propagating normally up a density gradient. For this problem, 1-D Cartesian geometry was used, so the density profile was obtained by transforming the spatial coordinate from equation 3.20, $r \rightarrow x - 36.42 \mu\text{m}$. The beam is travelled in the $+\hat{x}$ direction, with $I_0 = 14 \text{ W/cm}^2$. Only a single ray was used because the ray-trace is completely 1-D. A scale factor of $S = 1/16$ was used, resulting in a critical surface at $x \sim 14.14 \mu\text{m}$. The simulation had bounds $x \in [-16, 20] \mu\text{m}$ with a resolution of $\Delta x \sim 1 \text{ nm}$. A high resolution was used here in order to resolve the highly oscillatory coher-




Figure 3.13: Results of the 1-D field reconstruction test, comparing SOLAS fields with different caustic field capping methods to the LPSE results. Panel a) shows the oscillatory behaviour of the field solution which arises due to the interference between the incident and reflected sheet. Panel b) plots the same results, but zoomed in on the caustic region to demonstrate the differences in fields obtained from the different methods.

ent field sum and also enabled use of the EI field capping method. Power changes of the rays due to Inv-Brem and CBET were also neglected.

The laser in this problem travelled up the density gradient, reflecting from the critical surface. For light travelling normally up a density gradient, the critical density is the location of the caustic, because the ray area, $S(\tau)$, is constant and therefore the amplitude diverges when $\varepsilon \rightarrow 0$. The total field is given by the coherent sum between the incident and reflected sheets. Fig. 3.13.a shows the LPSE result compared to the field obtained from SOLAS without a cap on the caustic field; using the FL approach; and using the EI approach. The only difference between the SOLAS methods occurs in the caustic region¹², shown more clearly in Fig. 3.13.b, which in this problem is for $x > 13.64 \mu\text{m}$. Failure to cap the caustic field leads to a divergent $|E_z|$ at the critical surface, but both the EI and the FL methods show good agreement with LPSE. This justifies the default use of the simpler FL method for SOLAS CBET calculations.

Standard GO cannot capture the field beyond the critical surface because rays cannot propagate beyond $\varepsilon = 0$. The evanescent field in this region can be reproduced in the complex-GO framework, which integrates rays which have complex-valued properties and can thus be evolved beyond the critical surface [123]. This, however, is not deemed to be necessary for accurate CBET modelling of OMEGA-scale direct-drive implosions.

3.4.3.2 2-D Reflected Beam

A second field reconstruction test was also implemented in 2-D geometry. Unlike the previous 1-D test, the reconstructed field for this problem also depends on the divergence of neighbouring rays. This is because apart from normally incident rays at the centre of the beam, $\mathbf{k} \nparallel \nabla n_e$, so the area, $S(\tau)$, of the rays vary. Simulations at 1/64 and 1/4 scales are

¹²*i.e.* where equation 3.19 is satisfied.




Figure 3.14: Results of the 2-D field reconstruction test. Panel a) plots the field from LPSE on the left side and the SOLAS field on the right for the $1/64$ OMEGA scale simulation. Panel b) shows the SOLAS field from the $1/4$ scale simulation which, when compared to a), demonstrates that as scale increases, the relative size of the caustic region decreases. Panels c), d) and e) are lineouts near the caustic region along y from the $1/64$ scale simulations at $x = 0, 3$ and $6 \mu\text{m}$ respectively, the positions of which are indicated by the white lines on a). Note that the circular features in the field in panels a) and b) arise due to aliasing artefacts between the true solution scale, with the resolution of the image and/ or simulation grid.

presented here. The $S = 1/64$ simulation bounds were $x, y \in [-20, 20] \mu\text{m}$ and resolution $\Delta \sim 20 \text{ nm}$. The $S = 1/4$ simulation bounds were $x, y \in [-200, 200] \mu\text{m}$ and resolution $\Delta \sim 100 \text{ nm}$. For both scales, a single beam propagated parallel to $+\hat{y}$ and was centred on $x = 0 \mu\text{m}$. The beam widths were $\sigma = 8$ and $115 \mu\text{m}$ for the $1/64$ and $1/4$ scales respectively. The standard uniform sampling procedure described in Sec. 3.3.3 was used with $C_N = 2$. CBET and Inv-Brem were again both neglected.

Fig. 3.14.a shows the LPSE and SOLAS, FL fields for the $S = 1/64$ setup. Qualitative agreement is good between LPSE and the ray-based field in the sub-critical plasma. The field from the SOLAS FL $S = 1/4$ simulation is plotted in Fig. 3.14.b, from which it can be seen that the caustic region, which has a characteristic width $\sim \lambda_0$, shrinks in relative size as the scale increases. This suggests that at increasing scale, the importance of accurate caustic modelling may decrease in importance. Plotted in Figs. 3.14.c, 3.14.d and 3.14.e are lineouts from the $1/64$ simulations taken at $x = 0, 3$ and $6 \mu\text{m}$ respectively. They all include LPSE and SOLAS results with FL, EI caustic field capping along with no limiting. These plots demonstrate that caustic field limiting is necessary to compare favourably to the LPSE solution and that the FL approach slightly overestimates the caustic field compared to LPSE and the EI method. The discrepancy between the FL and EI results is again deemed sufficiently small to justify the default use of the simpler FL approach. To remind the reader, this choice is because the




Figure 3.15: The SOLAS CBET model operation loop. F_i is shorthand for the total field from all sheets at pump depletion iteration i and f_{CP} is the modifier to the caustic pump multiplier, described in Sec. 3.5.6 of the j^{th} energy conservation iteration. ΔE denotes the energy conservation error, which is the total incident power minus deposited power and power exiting the grid.

implementation of the EI method in SOLAS requires grid resolution $\Delta \ll \lambda_0$, the reason for which is described at the end of Sec. 3.4.2.

3.5 Ray-Based CBET Model

This section describes the implementation of the CBET model in SOLAS. Fig. 3.15 shows a flowchart which illustrates a broad overview of the operational loop of the CBET model when coupled to CHIMERA. The initial ‘field reconstruction’ ray-trace in the absence of CBET, follows the procedure described in Sec. 3.4 to obtain the field for every sheet on the SOLAS mesh. Ray-based CBET models, including the model in SOLAS, typically must perform repeated ray-traces through the same hydrodynamic profiles to account for pump depletion and optionally energy conservation.

The pump depletion loop, described in more detail in Sec. 3.5.5, is a fixed point iteration method to obtain the new CBET laser fields. This loop is necessary as each iteration produces a new background field which leads to a different CBET interaction, which converges when CBET has been fully accounted for. The resultant solution will not necessarily conserve energy if CBET occurs in the presence of laser caustics, due to errors in field reconstruction

and CBET scattering in these locations. To remedy this, additional energy conservation iterations can be performed, where a multiplier on the CBET gain in caustics can be varied to conserve laser energy. Although this is an ad-hoc correction, it is found to lead to better agreement with higher fidelity solver, as is explicitly demonstrated in Sec. 3.5.7.3.

CBET does not affect the trajectory of light, only its power and therefore repeating the calculations of ray trajectories is both computationally costly and unnecessary. Ray locations are therefore stored in a memory-efficient, linked-list to be re-used in both loops. It was found in SOLAS that caching the ray locations in this manner typically leads to an order of magnitude speed up compared to performing a new ray-trace for every iteration.

3.5.1 Power Change of Rays due to CBET

Ray-tracing calculations assume steady-state hydrodynamic profiles for rays propagating through the plasma and that the field of a ray propagates as a plane wave along the ray path [199]. The linear gain theory of SBS, which describes the steady-state energy change of locally plane wave light which passes through uniform plasma conditions, can therefore be used [130, 162]. ‘Linear’ in this context means that the plasma response to the driving fields is sufficiently small to be treated by a linear expansion. The uniform assumption here means that the plasma conditions must not vary significantly over the interaction time or length scales. For direct-drive conditions, hydrodynamic time scales are typically ~ 100 ps and CBET saturates on the order of ~ 10 ps. Ray steps are also limited to grid cell boundaries, which are smaller than hydrodynamic length-scales, so it is safe to assume a locally uniform plasma background when this model is applied only over a single ray step at a time.

The steady-state interaction between two parallel-polarised sheets (i, j) over a path length $d\tau$ can be written in terms of the field magnitude of the sheet ($|E_{i,j}|$) or equivalently ray power ($P_{i,j} = A_{i,j}|E_{i,j}|^2$),

$$\begin{aligned}\frac{d|E_i|^2}{d\tau} &= (-\kappa_{IB} + \gamma_{ij}|E_j|^2)|E_i|^2, \\ \frac{dP_i}{d\tau} &= (-\kappa_{IB} + \gamma_{ij}A_jP_j)P_i,\end{aligned}\tag{3.21}$$

where κ_{IB} is the Inv-Brem absorption kernel, A_j is the amplitude of sheet j and γ_{ij} is the CBET gain of sheet i from the interaction with sheet j . The sheet/ ray which changes energy in a given interaction (here i), is often called the *probe* while the other beam is termed the *pump*. The gain is a function of both the plasma conditions and the local field profiles for both sheets in the interaction, so a gain must be calculated between every pair of sheets at every location where both are present. Going from the first to the second equation assumes that $dA_i/d\tau = 0$, which is consistent with the assumption from the linear gain theory that the fields are approximately uniform over a small step. Different models exist for the gain, which either use a fluid or kinetic treatment for the plasma IAW response.

When multiple different probe beams are present, the total CBET interaction is treated simply as the sum of all interactions from different sheets,

$$\frac{dP_i}{d\tau} = (-\kappa_{IB} + \kappa_{CBET})P_i,\tag{3.22}$$

where the CBET kernel has been defined $\kappa_{\text{CBET}} = \sum_{j \neq i} \gamma_{ij} |E_j|^2$. This summation assumes that the fields from different sheets are uncorrelated. For direct-drive conditions, this assumption is broadly correct, apart from when a probe sheet undergoes a CBET interaction in the caustic region of the probe. In this case, the *Coherent Caustic* correction is applied, which is described in Sec. 3.5.3.

Note that when computing both Inv-Brem and CBET, the partition of power lost to each interaction can be computed as,

$$\begin{aligned} \Delta P_{i,\text{IB}} &= -\langle P_i \rangle \frac{\kappa_{\text{IB}}}{\kappa_{\text{CBET}} - \kappa_{\text{IB}}}, \\ \Delta P_{i,\text{CBET}} &= \langle P_i \rangle \frac{\kappa_{\text{CBET}}}{\kappa_{\text{CBET}} - \kappa_{\text{IB}}}, \\ \langle P_i \rangle &= P_i + P_i (\kappa_{\text{CBET}} - \kappa_{\text{IB}}) d\tau / 2, \end{aligned} \quad (3.23)$$

where $\Delta P_{i,\text{IB}}$ and $\Delta P_{i,\text{CBET}}$ are the powers lost to Inv-Brem and CBET respectively and $\langle P_i \rangle$ is the average power of the ray over the step, $d\tau$. This can be derived by solving Eq. 3.21 for P_i and then taking the limit $d\tau \rightarrow 0$ [131]. The deposition, $\Delta P_{i,\text{IB}}$, can then be interpolated to the grid as an electron energy source as is described in Sec. 3.3.4.

The two formulations of the gain are presented below. Broadly, the gain describes the plasma response to the 2 driving fields and therefore dictates the energy transfer between them. They differ in that the fluid formulation arises by deriving a gain from linearised fluid equations for the plasma response, whereas the kinetic formulation is obtained from a linearisation of the Vlasov equation. The kinetic gain is more complete and should be used for comparison to experimental observables, whereas the fluid gain is useful for comparison with LPSE, which directly solves the linearised fluid equations.

Fluid CBET Gain The linear fluid gain, derived by Randall *et al.* in Ref. [130] is,

$$\begin{aligned} \gamma_{ij} &= \frac{n_e e}{4m_e c \omega_i} \frac{1}{T_e (1 + 3T_i/ZT_e)} \frac{R(\eta_{ij})}{v_{\text{ia}}}, \\ R(\eta_{ij}) &= \frac{(v_{\text{ia}}/\omega_s)^2 \eta_{ij}}{(\eta^2 - 1)^2 + (v_{\text{ia}}/\omega_s)^2 \eta_{ij}^2}, \\ \eta_{ij} &= \frac{(\omega_j - \omega_i) - \mathbf{k}_s \cdot \mathbf{u}}{\omega_s}, \end{aligned} \quad (3.24)$$

where T_e and T_i are the electron and ion temperatures in eV respectively, ω_i and $\mathbf{k}_i$ are the frequency and wavevector of sheet i respectively, Z is the ionisation state, $\mathbf{u}$ is the fluid velocity, $\omega_s = |\mathbf{k}_s|c_s$ is the IAW frequency, $\mathbf{k}_s = \mathbf{k}_j - \mathbf{k}_i$ is the IAW wavevector, $c_s = \sqrt{e(ZT_e + 3T_i)/m_i}$ is the sound speed and v_{ia} is the IAW damping rate. This formulation is useful for comparison with LPSE, which directly solves the linearised fluid equations and therefore in situations where the GO assumptions are valid, the two methods should give similar results. However, it is generally not considered to be a good choice of model to compare to experiment. The damping rate, v_{ia} does not have an analytic formula for arbitrary plasma conditions and therefore must be prescribed. Additionally, Eq. 3.24 assumes an average ion treatment of the

species in the plasma if multiple ion species are present. In reality, both fast and slow modes of IAW can be driven in a 2-species plasma [128], which can independently lead to CBET scattering. This effect cannot be captured using the fluid gain, where the plasma is assumed to be made of a single ion species with a number density averaged ion mass. Codes that use this formulation typically use a damping value $\nu_{ia}/\omega_s = 0.2$ for CH plasmas, which is the value used for all fluid gain simulations in this chapter.

Kinetic CBET Gain The linear kinetic CBET gain, described by Michel in Ref. [129] is,

$$\begin{aligned}\gamma_{ij} &= \frac{e^2 |\mathbf{k}_s|^2}{4m_e^2 c \omega_i^3} \text{Im}(K_{ij}), \\ K_{ij} &= \frac{\chi_e(1 + \chi_i)}{(1 + \chi_e + \chi_i)}, \\ \chi_e &= \frac{-1}{2|\mathbf{k}_s|^2 \lambda_{De}^2} Z' \left(\frac{(\omega_j - \omega_i) - \mathbf{k}_s \cdot \mathbf{u}}{\sqrt{2} |\mathbf{k}_s| v_{Te}} \right), \\ \chi_i &= \frac{-1}{2|\mathbf{k}_s|^2 \lambda_{De}^2} \frac{T_e}{\langle Z \rangle T_i} \sum_{\alpha}^{\text{species}} f_{\alpha} Z_{\alpha}^2 Z' \left(\frac{(\omega_j - \omega_i) - \mathbf{k}_s \cdot \mathbf{u}}{\sqrt{2} |\mathbf{k}_s| v_{T\alpha}} \right),\end{aligned}\tag{3.25}$$

where χ_e and χ_i are the electron and ion susceptibilities respectively, f_{α} and Z_{α} are the number fraction and ionisation of ion species α respectively, Z' is the derivative of the plasma dispersion function, $v_{Tn} = \sqrt{eT_n/m_n}$ is the thermal velocity of species n with temperature T_n in eV and mass m_n , $\lambda_{De} = \sqrt{\epsilon_0 e T_e / n_e e^2}$ is the electron Debye length and $\langle Z \rangle = \sum_{\alpha} f_{\alpha} Z_{\alpha}$ is the average ionisation. Note that the susceptibilities $\chi_{i,e}$ are complex. In order to compute Z' in SOLAS, the CALGO library is used to numerically compute the value of the *Faddeeva function*, $w(x)$ [200]. This is then related to the plasma dispersion function and its derivative by,

$$\begin{aligned}Z(x) &= i\sqrt{\pi} w(x), \\ Z'(x) &= -2[1 + xZ(x)],\end{aligned}\tag{3.26}$$

where $i = \sqrt{-1}$ and $w(x)$, $Z(x)$ and $Z'(x)$ are all complex numbers [201]. Unlike the fluid gain, this theory correctly captures the resonance and has no effective free parameters, making it the judicious choice for predictive simulations. The plasma dispersion function is however relatively slow to numerically evaluate and dominates computational runtime if evaluated every ray step. SOLAS therefore calculates γ_{ij} between each sheet in every grid cell at the end of the field reconstruction ray-trace (from Fig. 3.15) and then interpolates this pre-calculated value onto the ray locations for the pump-depletion iterations. This reduces computational runtimes by orders of magnitude for many beam direct-drive simulations, but has an associated memory overhead, which is discussed in more detail in Sec. 3.5.4. All simulations in this thesis use the kinetic CBET gain unless explicitly stated otherwise.

Random Polarisation Correction The gains calculated above all assume that the interacting fields have parallel polarisations. For many laser systems, beams may have arbitrary or random polarisations, such as at the OMEGA laser facility where Distributed Polarization Ro-




Figure 3.16: Illustration of the Caustic Gain Truncation (CGT) and caustic region identification algorithms. Panel a) shows a laser, separated into incident and reflected sheets by a caustic, propagating up and reflecting in a linear density ramp. CGT is illustrated in b) where a cell is separated into a *lit* and *unlit* region by the reconstructed *caustic plane* (dashed magenta line). Panels c) and d) illustrate the method used to identify the *caustic region* (shaded blue), where the coherent caustic modifier is applied. c) shows that cells are geometrically separated into *capped* (blue squares) and *uncapped* (red squares) regions and the coherent caustic modifier to the gain is applied in this capped, caustic-region.

tators (DPRs) separate each beam into 2 orthogonally polarised sub-beams which overlap to effectively have one spot with random polarisation [65]. If a simulation is conducted with multiple beams of random polarisation, a *Polarisation Smoothing* multiplier, M_{PS} , must be applied to the value of γ_{12} which accounts for random polarisation angles between the field from different sheets,

$$M_{PS} = \frac{1}{4}(1 + \cos^2 \theta), \quad (3.27)$$

where θ is the angle between the wavevectors of the interacting sheets.

3.5.2 Caustic Gain Truncation

For rays propagating through SOLAS cells, nearest neighbour interpolation is used to interpolate field quantities to the rays. This is generally an adequate approximation for field magnitudes, frequencies and wavevectors, which vary over greater length scales than typical cell resolutions $\mathcal{O}(1\mu\text{m})$, apart from in the vicinity of a caustic [2]. Near caustics, rays turn sharply and fields have a sharp cut-off, dropping to zero sharply, as be seen in the lineouts from Fig. 3.14. The sharp turning of rays mean that $\mathbf{k}$ can vary rapidly within a cell. Although SOLAS does have variable radial resolution, as described in Sec. 3.4, this resolution is limited to the minimum radial cell size from the hydrodynamic grid, which is not dictated from laser parameters and cells are typically larger than the caustic region width $\mathcal{O}(\lambda_0)$. The issue of $\mathbf{k}$ variation is not addressed in SOLAS, but linear interpolation or a truly adaptive mesh, separate from the CHIMERA grid could be implemented to address this.

The field cut-off is an issue however, as it can significantly affect energy conservation in ray-based CBET models. CGT is an algorithm that effectively improves the resolution of the grid near laser caustics [173]. This method separates cells into *lit* and *unlit* regions. If rays travel in the unlit region, then the CBET gain is not applied. For example, Fig. 3.16.a shows a beam propagating up a linear density gradient, separated by the caustic into incident and reflected sheets. Fig. 3.16.b shows a zoomed in view of in the vicinity of the caustic of the incident sheet ray locations, plotted as grey dots. If a ray from another beam were present with $14.3 \lesssim y < 15 \mu\text{m}$, nearest neighbour interpolation would lead to the ray undergoing a CBET interaction in this *unlit* region, where no rays from the pump sheet are present. Field magnitudes in this caustic region are typically large, so even though this is a small region of space, it can lead to significant errors, especially at lower resolutions.

The CGT algorithm described by Follett in Ref. [173] defines the *lit* region for each sheet by creating a mesh of the caustic locations from every ray on the sheet and only applying CBET gains in on the side of this mesh that rays traversed. A different algorithm has been developed for SOLAS, which geometrically splits each cell into *lit* and *unlit* regions by a *caustic plane* if at least one ray caustic occurs in the cell. The *caustic plane*¹³ is defined by a point $\mathbf{x}_{\text{caus}}$ and normal $\hat{\mathbf{n}}_{\text{caus}}$. The location, $\mathbf{x}_{\text{caus}}$, is simply the average of all ray caustic locations from the sheet present in the cell, for example shown in Fig. 3.16.b by the small pink circles for each ray caustic and the larger pink square for the cell averaged location. The normal is estimated by the average over all ray caustics in the cell of the ray tangent vector just after the caustic $\hat{\mathbf{t}}_{\text{After}}$, minus the velocity just before $\hat{\mathbf{t}}_{\text{Before}}$,

$$\hat{\mathbf{n}}_{\text{caus}} = \left\langle \frac{\hat{\mathbf{t}}_{\text{After}} - \hat{\mathbf{t}}_{\text{Before}}}{|\hat{\mathbf{t}}_{\text{After}} - \hat{\mathbf{t}}_{\text{Before}}|} \right\rangle, \quad (3.28)$$

such that the normal points to the lit side. This is plotted in Fig. 3.16.b as small pink arrows for each ray caustic normal and the large arrows in each cell for the cell averaged value. The planes defined by these points and normals for each cell that contains a sheet caustic is

¹³In cylindrical and spherical simulation geometries, the ‘planes’ are defined in the native coordinate system, which creates manifolds that better reflect the assumed geometry of the problem.

shown as a pink dashed line. If a probe ray at location $\mathbf{x}_{\text{probe}}$ is propagating in a cell where there is a pump sheet caustic, it will only experience a CBET interaction if $\hat{\mathbf{n}}_{\text{caus}} \cdot (\mathbf{x}_{\text{probe}} - \mathbf{x}_{\text{caus}}) > 0$. Rays from other sheets also have their steps limited by the caustic plane.

This is a slightly lower fidelity implementation of the more complete CGT model from Ref. [173]. The model implemented by Follett can represent caustics which, within a grid cell, are not locally well approximated by planes. The SOLAS model is however simpler to implement and also acts to effectively increase the resolution of the cells near caustics. The caustic test problem presented in Sec. 3.5.7.2 illustrates the effectiveness of the SOLAS-CGT model at reducing energy conservation errors for caustic CBET interactions where the grid cells are significantly larger than the caustic region.

3.5.3 Coherent Caustic Correction and Caustic Region Identification

In Sec. 3.5.1, in order to obtain Eq. 3.22, it was assumed that when a ray undergoes CBET interactions with multiple pump sheets, the field from the pump sheets can be treated independently. This is equivalent to assuming that over a ray propagation path, $d\tau$, the fields from the sheets do not add coherently,

$$\int d\tau \sum_j^{\text{sheets}} |E_j|^2 \approx \int d\tau \left| \sum_j^{\text{sheets}} E_j \right|^2, \quad (3.29)$$

where E_j is the field from the j^{th} sheet. Eq. 3.29 is valid for all regions of space where the phase from different sheets are uncorrelated. For direct-drive configurations, this assumption holds everywhere apart from the caustic region of the pump sheet, where fields add coherently. Thus, the coherent sum of fields in the caustic region is larger than the incoherent sum, so without correcting for this, CBET scattering through a pump beam in its caustic region will be underestimated.

The ratio of the coherent and incoherent sums of incident (E_{inc}) and reflected (E_{refl}) fields of the pump sheet in its caustic region is called the *coherent caustic multiplier*,

$$M_{\text{CP}} = \frac{|E_{\text{inc}}|e^{i\phi_{\text{inc}}} + |E_{\text{refl}}|e^{i\phi_{\text{refl}}}}{|E_{\text{inc}}|^2 + |E_{\text{refl}}|^2}. \quad (3.30)$$

By making the same assumption as used to obtain Eq. 3.18, that $|E_{\text{inc}}| \sim |E_{\text{refl}}|$, this simplifies to,

$$M_{\text{CP}} = 1 + \sin(\phi_{\text{refl}} - \phi_{\text{inc}}), \quad (3.31)$$

which should be applied in the caustic region, (defined in Eq. 3.19), as the region where $|\varphi_1 - \varphi_2| \leq \lambda_0/2$. As stated in the discussion in Sec. 3.4.2, grid cells are typically larger than the size of the caustic region, $\mathcal{O}(\lambda_0)$, and SOLAS does not include linear interpolation of field quantities. Therefore, Eq. 3.31 cannot robustly be evaluated at the relevant scales. The caustic region is identified by a geometric algorithm, similar to the *caustic plane* method discussed in Sec. 3.5.2 and the ‘average’ value of the coherent caustic multiplier is used,

$$M_{\text{CP}} = 1 + 2/\pi, \quad (3.32)$$

when a ray undergoes CBET with a pump sheet in its caustic region. This coherent caustic correction increases the CBET scattering and is an important addition to make sure that ray-based CBET models are energy conserving.

The geometric approach that was developed to estimate when a ray is in the caustic region of another beam, ($|\varphi_1 - \varphi_2| \leq \lambda_0/2$), is described here. Similar to the *caustic plane* method from Sec. 3.5.2, this algorithm works by finding planes that divide cells and thus effectively increase their resolution, improving the accuracy of the nearest neighbour interpolation. The key assumptions of the method are that the boundary of the caustic region can be treated locally in grid cells as a plane and that the caustic region is approximately where ray amplitudes are greater than the FL capped amplitude, $A_{\text{ray}} > A_{\text{FL}}$. Fig. 3.16.c plots the trajectory of incident sheet rays from Fig. 3.16.a, with the locations of capped and uncapped rays plotted as small transparent red and blue dots respectively. The average value of each of these locations is represented by the larger blue ($\langle \mathbf{x}_{\text{cap}} \rangle$) and red squares ($\langle \mathbf{x}_{\text{uncap}} \rangle$) respectively, which are stored (only where present), in each SOLAS grid cell. The region of the cell occupied by these capped (blue) ray locations is assumed to be equivalent to the caustic region. The validity of the assumption was tested by conducting high resolution simulations of the 2-D reflected beam field reconstruction test in Sec. 3.4.3 to accurately reconstruct the phase and amplitudes of incident and reflected sheets on a high resolution mesh. It was found that the assumption is adequate, although it typically slightly overestimates the width of the caustic region by $\mathcal{O}(10\%)$. This is not considered to be a significant issue, because energy conservation iterations are employed which minimise errors by varying M_{CP} over multiple iterations. Therefore, errors in the initial iteration are compensated by this minimisation.

To geometrically estimate the caustic region in cells which contain both capped (blue) and uncapped (red) ray locations, a *caustic region plane* is defined to separate these regions, defined by a point, $\mathbf{x}_{\text{CR}}$, and normal, $\hat{\mathbf{n}}_{\text{CR}}$,

$$\begin{aligned} \mathbf{x}_{\text{CR}} &= \frac{\langle \mathbf{x}_{\text{cap}} \rangle + \langle \mathbf{x}_{\text{uncap}} \rangle}{2}, \\ \hat{\mathbf{n}}_{\text{CR}} &= \frac{\langle \mathbf{x}_{\text{cap}} \rangle - \langle \mathbf{x}_{\text{uncap}} \rangle}{|\langle \mathbf{x}_{\text{cap}} \rangle - \langle \mathbf{x}_{\text{uncap}} \rangle|}, \end{aligned} \quad (3.33)$$

where $\hat{\mathbf{n}}_{\text{CR}}$ is defined such that it points toward the caustic region. For a probe ray propagating through a mesh which contains a pump field, the ray is inside the caustic region if $\hat{\mathbf{n}}_{\text{CR}} \cdot (\mathbf{x}_{\text{probe}} - \mathbf{x}_{\text{CR}}) > 0$, under which circumstances, the pump-probe CBET gain is multiplied by the average value of M_{CP} defined in 3.32. This is illustrated in Fig. 3.16.c by the purple dashed lines. If a cell contains both a *caustic region plane* and a *caustic plane* (from Sec. 3.5.2), then the ray must be inside the caustic and on the lit side of the caustic to undergo CBET with M_{CP} applied. Fig. 3.16.d plots the caustic region, as obtained by the algorithm, as a transparent blue shaded area.

3.5.4 Dynamic Memory for Storing Fields and CBET Gains

CBET gains must be calculated between every pair of sheets that cross at a given location. For N_b beams, each with 2 possible sheets, discretised on a grid with N_x cells in each dimension, a total of $N_f \sim 2N_x^3 N_b$ *field data points* must be stored. Each of these data points must store the field magnitude, phase, wavevector and information used to evaluate CGT and the coherent caustic correction where necessary, as described in Sec. 3.5.2 and 3.5.3 respectively. This amounts to large computational memory requirements to store the fields. Additionally, SOLAS stores the CBET gains between each laser sheet to improve the code execution time. A separate gain exists between each sheet and therefore the memory cost to store this information scales as $N_g \sim 2N_f^3 N_b^2$.

N_g is a single number and N_f represents all field information, so for simulations with small numbers of beams, N_f is the dominant memory cost, whereas for large number of overlapping beams such as direct-drive simulations, N_g dominates. It is therefore important to use a memory efficient method to store this data so that large scale, 3-D simulations can be completed. SOLAS utilises pointers in each grid cell to dynamically allocate the field data, such that only where a sheet is present is its corresponding field magnitude, phase *etc.* stored. Additionally, gains between sheets are an attribute of the field data and are stored as single-precision floats, rather than the native SOLAS double precision, which also reduces memory costs for large simulations. These measures significantly reduce memory overhead for multi-dimensional CBET simulations, however high resolution 3-D simulations are still extremely memory intensive, which is an issue also experienced by IFRIIT, the other 3-D CBET code coupled to hydrodynamics.

3.5.5 Pump Depletion Iterations

The field reconstruction ray-trace (the initial ray-trace performed in the absence of CBET) returns the field from each sheet $|E_j^{i=0}|(\mathbf{x}, \mathbf{k}, \omega, \phi)$, where j is the sheet index and $i = 0$ indicates the field reconstruction ray-trace. This field is discretised on the SOLAS mesh using nearest neighbour interpolation, so that a ray passing through a cell updates the field in that cell. Rays can then be re-traced through this mesh and undergo changes in power due to CBET interactions with the fields from other sheets. The solution obtained from a single ray-trace with CBET is not the true solution however, because the field that the rays see throughout the CBET ray-trace is not dynamically updated by the change in ray power due to CBET as they propagate. This necessitates iteratively solving the ray-trace to account for CBET in *pump depletion iterations*, whereby for iteration i , the rays from sheet j propagate through the mesh and undergo CBET through $|E_{sheet \neq j}^{i-1}|$, while building up a new field, $|\tilde{E}_j^i|$. The new *build-up* fields, $|\tilde{E}_j^i|$, are introduced for the next iteration. This loop is repeated until convergence of the field has been achieved. Convergence is tested each iteration ($i > 0$) against the total power of rays $\tilde{P}_{j,n}^i$, from each sheet j travelling through each cell, explicitly defined by




Figure 3.17: Results of successive pump depletion iterations for an illustrative simulation of two beams crossing and undergoing a CBET resonance in a uniform plasma background. Each beams' rays from the field reconstruction ray-trace are plotted in a), where more opaque rays indicate a higher power. Panel b) plots the reconstructed intensity from this ray-trace, where $i = 0$ indicates the field reconstruction ray-trace and therefore no CBET occurs. Panels c) and d) plot the intensity reconstructed from the 1st and 7th pump depletion iterations respectively, where the effect of transfer of powers between beams due to CBET is evident. Panel e) plots the convergence parameter from the simulation over successive pump depletion iterations, with convergence defined here as $K_{\text{Conv},0} = 10^{-6}$.

the convergence parameter,

$$K_{\text{Conv}}^i = \max_{\substack{j \in \text{sheets} \\ n \in \text{cells}}} \left| \left(\tilde{P}_{j,n}^i / \tilde{P}_{j,n}^{i-1} \right) - 1 \right| < K_{\text{Conv},0}, \quad (3.34)$$

where $K_{\text{Conv},0}$ is a user defined parameter, typically set to 10^{-5} unless stated otherwise.

This procedure is illustrated in Fig. 3.17 for a simple, illustrative problem of two beams crossing in a uniform, stationary plasma background with a frequency difference between the beams tuned to give a CBET resonance. The effect of Inv-Brem deposition has been neglected in this simulation. Fig. 3.17.a shows the ray trajectories for the two crossing beams and Fig. 3.17.b plots the intensity from the field reconstruction iteration, without CBET. The intensities after one and seven pump depletion iterations are plotted in Fig. 3.17.c and Fig. 3.17.d respectively, which show that a single re-trace through the no-CBET fields does not return the converged solution. Fig. 3.17.e plots the convergence parameter after each iteration, illustrating that seven iterations were required to obtain the converged solution, plotted in Fig. 3.17.d.

For many beam simulations with complicated beam trajectories and many overlapping crossings such as direct-drive calculations, the problem is asymptotically unstable. This




Figure 3.18: Energy conservation iterations for a two-beam caustic test problem, which is presented in Sec. 3.5.7.2. 4 iterations were required to achieve energy conservation, which was defined to be $\delta E < 1 \times 10^{-4}$. The order of the iterations is indicated by the text next to each point.

means that for iteration i , allowing rays to undergo CBET in the full $|\tilde{E}_{\text{sheet} \neq j}^{i-1}|$ causes unbounded gains on ray powers leading to divergent results for $|\tilde{E}_j^i|$. Numerical damping is therefore applied for these calculations to stabilise the solution. The actual fields that a ray will see for the subsequent iteration, $i + 1$, are composed of the weighted sum of the two *build-up* fields,

$$|E_j^i| = f_{\text{damp}} |\tilde{E}_j^i| + (1 - f_{\text{damp}}) |\tilde{E}_j^{i-1}|, \quad (3.35)$$

where f_{damp} is a user-defined damping value, which is by default set to a (relatively conservative) value of 0.75, to ensure that the solver does not diverge for complex problems.

3.5.6 Energy Conservation Iterations

The converged field solution from the pump depletion iterations returns the steady-state field and deposition profiles, accounting for CBET, given a field reconstruction algorithm, interpolation scheme *etc.* If there are errors in, for example, the field reconstruction algorithm, then the pump depletion iterations will return a steady-state solution, however this is not guaranteed to be accurate, *i.e.* the same as the result from a higher fidelity solver. The main cause of errors in field reconstruction is the treatment of laser caustics, as demonstrated for example in Fig. 3.14. This figure demonstrates that the SOLAS field reconstruction algorithm obtains essentially perfect agreement with LPSE everywhere apart from in the vicinity of the caustic. It can therefore be assumed that any errors in the CBET profiles, returned from the pump depletion iterations, occur due to errors in the field reconstruction and CBET scattering in the caustic region. The energy conservation error, ΔE , is a good proxy for the accuracy of the fields obtained from CBET, as in the absence of laser caustics, energy is conserved in ray-based algorithms [2]. This is defined as,

$$\Delta E = \frac{P_{\text{dep}} + P_{\text{out}} - P_{\text{inc}}}{P_{\text{inc}}}, \quad (3.36)$$

where P_{dep} , P_{out} and P_{inc} are the power deposited by Inv-Brem, the ray-power leaving the domain and the incident power respectively. If this value is not small, it suggests that caustics are not accurately treated in the CBET scattering.

As in the BEAMCROSSER code, the assumption in SOLAS is made that minimising ΔE by altering the CBET gains in the caustic region will bring the result into better agreement with the true solution [2]. M_{CP} is modified to include a new *caustic pump multiplier*, f_{CP} , which is applied to probe ray CBET gains, when they are in the caustic region of the pump sheet,

$$\tilde{M}_{\text{CP}} = \begin{cases} 1 + f_{\text{CP}} \left(\frac{2}{\pi} \right), & \text{if in caustic region} \\ 1, & \text{otherwise} \end{cases} \quad (3.37)$$

where the same, geometric identification for the pump sheet caustic region is used as in Sec. 3.5.3. $f_{\text{CP}} = 1$ initially and is varied between $f_{\text{CP}} \in [-\pi/2, 2]$ over several *energy conservation* iterations¹⁴ to minimise ΔE . The lower bound on f_{CP} is to prevent ‘negative scattering’ in the caustic region, where the sign of the gain flips. The upper bound was selected by experimentation to prevent excessively large gains which result in the exponential power gain of rays blowing up to infinity.

An implementation of the Secant method is used to estimate the next value of f_{CP} . The first and second energy conservation iterations in a simulation use $f_{\text{CP}} = 1$ and 0, to obtain a gradient with which to estimate values of the multiplier. The Secant method was chosen for its simplicity and the fact that bracketing values of the root are not required. An example of the minimisation results are shown in Fig. 3.18. These results are from a two-beam caustic test problem, presented in Sec. 3.5.7.2, but the specifics of the problem are not important for the discussion here, simply that CBET in the presence of a caustic occurs. Therefore, a caustic region exists inside which f_{CP} will lead to modified gains which reduce energy conservation errors. The curve of $\Delta E(f_{\text{CP}})$ is simple and monotonic, which is observed for all simulations conducted with SOLAS, regardless of the complexity of the laser and target configuration. This means a more complex algorithm than the Secant method is not required for minimisation. When the CBET model runs in-line with the hydrodynamics, the value of f_{CP} from the last timestep that CBET was calculated is used, which means that energy conservation iterations are often not necessary. This is shown explicitly in Fig. 3.24.c, which shows the results of a 1-D CHIMERA-SOLAS simulation of OMEGA shot 89224. The bottom panel in the figure shows that energy conservation is frequently achieved by simply re-using the value of f_{CP} from the previous timestep.

To summarise, the total CBET gain that a probe ray (i) experiences in a pump sheet (j) is,

$$\tilde{\gamma}_{ij} = \gamma_{ij} M_{\text{PS}} \tilde{M}_{\text{CP}}, \quad (3.38)$$

where γ_{ij} is the value defined in Sec. 3.5.1 (depending on if using the kinetic or fluid gain model), the modified coherent caustic multiplier $\tilde{M}_{\text{CP}}$ is defined in Eq. 3.37 and the polarisation smoothing multiplier, M_{PS} , defined in Eq. 3.27, is applied if the beams are assumed to

¹⁴Each energy conservation iteration contains a new set of pump depletion iterations, as is shown in Fig. 3.15.




Figure 3.19: Results of the 2-D CBET test in the absence of laser caustics. Panels a) and b) show the electric field from the SOLAS and LPSE simulations respectively, after the pump depletion iterations have occurred. Panels c) and d) plot lineouts from the inbound fields of the pump and seed beams respectively for both SOLAS and LPSE. e) and f) plot the same for the outbound fields.

be randomly rather than parallel-polarised.

3.5.7 CBET Validation

This section presents results from several CBET test problems where the results from SOLAS are compared to the wave based solver LPSE. Because LPSE is a wave-solver coupled to a linearised fluid response, the fluid CBET gain was used for all simulations in this section. Beams in the simulations all had a vacuum wavelength, $\lambda_0 = 0.351$ nm and a super-Gaussian spot profile, which was defined in Eq. 3.1. For consistency with LPSE, the Doppler shift of laser frequency was not calculated in these simulations so its effect on CBET is neglected, *i.e.* the right-hand side of Eq. 3.5 was forced to be zero. The plasma backgrounds for all initialisations were fully-ionised CH with $Z = 3.1$ and $T_i = T_e/2$. All tests are presented in Ref. [2] and the LPSE data is available from the repository in Ref. [198].

3.5.7.1 2-D CBET Without Caustics

The first test problem is for the case of two beams crossing in a non-uniform, non-stationary plasma in the absence of laser caustics. The grid used for the problem was $x, y \in [-21, 22]$ μm and 2160 cells were used in each direction. This led to a resolution of $\Delta x \sim \lambda_0/20$ such that

Table 3.1: The ratio of the power of the seed beam as it exits the domain to its initial power.

	$P_{\text{seed,out}}^{\text{LPSE}} / P_{\text{seed,in}}^{\text{LPSE}}$	$P_{\text{seed,out}}^{\text{BEAMCROSSER}} / P_{\text{seed,in}}^{\text{BEAMCROSSER}}$	$P_{\text{seed,out}}^{\text{SOLAS}} / P_{\text{seed,in}}^{\text{SOLAS}}$
CBET + Absorption	1.315	1.325	1.320
CBET only	1.551	1.561	1.558
Absorption only	0.856	0.856	0.852

the beat field between the two beams could be resolved. The electron temperature was uniform, $T_e = 1$ keV. The density and fluid velocity profiles were,

$$n_e(y) = n_{\text{cr}}(0.006y_{\mu\text{m}} + 0.25), \quad (3.39)$$

$$\mathbf{u}(y) = -\hat{\mathbf{y}}(1.4 - 0.008y_{\mu\text{m}})c_s, \quad (3.40)$$

where $y_{\mu\text{m}}$ is the coordinate in microns. The beams both had widths $\sigma = 8.41 \mu\text{m}$ and peak intensities $I_0 = 2 \times 10^{18} \text{ W/m}^2$. The beams were defined as *pump* and *seed* and launched from $\mathbf{x} = [-21, 2] \mu\text{m}$ and $\mathbf{x} = [0, -21] \mu\text{m}$ along the $+\hat{\mathbf{x}}$ and $+\hat{\mathbf{y}}$ directions respectively. The beam and plasma configuration were chosen to create a strong CBET resonance where the beams cross one another.

Fig. 3.19.a and Fig. 3.19.b show the electric field obtained from the SOLAS and LPSE simulations respectively. It can be seen that the density gradient in $\hat{\mathbf{y}}$ causes the seed beam to refract somewhat and that CBET results in a significant energy transfer from the pump to the seed beam. The lineouts in Fig. 3.19.c, Fig. 3.19.d, Fig. 3.19.e and Fig. 3.19.f all show excellent agreement between SOLAS and LPSE. This is further demonstrated in Tab. 3.1, which lists the ratio of the outgoing seed beam power to the incident power for LPSE, SOLAS and the BEAM-CROSSER post-processor, the results from which are also presented in Ref. [2]. The energy conservation error after the pump depletion iterations was small, $\Delta E \sim 10^{-6}$, as expected because the CBET interaction does not involve any caustics.

3.5.7.2 2-D CBET With Caustics

Another two-beam problem was implemented to test CBET in the presence of a caustic. For this simulation, the grid was $x, y \in [-21, 21] \mu\text{m}$ and $x, y \in [-60, 60] \mu\text{m}$ with a minimum resolution of $\Delta x = \Delta y \sim \lambda_0/20$. A uniform electron temperature was used, $T_e = 2$ keV. The pump beam travelled straight up a linear density gradient and therefore did not refract, whereas the seed beam reflected in the gradient and travels out. The beam geometry was chosen so that the pump beam crossed through the middle of the seed beam caustic and the density and flow profiles were tuned to create a resonance at this crossing location,

$$n_e(y) = \begin{cases} 0.02y_{\mu\text{m}} + 0.3, & x \geq -15 \mu\text{m}, \\ 0, & x < -15 \mu\text{m}, \end{cases} \quad (3.41)$$

$$\mathbf{u} = -1.14c_s\hat{\mathbf{y}}. \quad (3.42)$$




Figure 3.20: Results of the 2-D CBET test in the presence of a laser caustic. Panels a) and b) show the electric field from the LPSE and SOLAS simulations respectively. Also plotted is an inset zoom on the caustic region of the seed beam, which demonstrates that SOLAS does not capture the evanescent field. Panels c) and e) plot lineouts from the inbound fields of the pump and seed beams respectively for both SOLAS and LPSE. d) and f) plot the same for the outbound fields.

Both beams again had widths $\sigma = 8.41 \mu\text{m}$. The pump and seed beams were launched from $\mathbf{x} = [18, -21] \mu\text{m}$ and $\mathbf{x} = [-35, -21] \mu\text{m}$ and pointed along $\hat{\mathbf{k}} = [0, 1]$ and $\hat{\mathbf{k}} = [0.707, 0.707]$ respectively. The peak intensities were $I_0 = 2 \times 10^{18} \text{ W/m}^2$ and $I_0 = 2 \times 10^{17} \text{ W/m}^2$ for the pump and seed respectively.

Fig. 3.20.a and Fig. 3.20.b plot the LPSE and SOLAS¹⁵ total fields respectively. The inset zooms demonstrate that the evanescent fields are not captured in the GO framework. Fig. 3.20.c, Fig. 3.20.d, Fig. 3.20.e and Fig. 3.20.f plot the lineouts of the LPSE and SOLAS field solutions at their exit from the grid, which all demonstrate excellent agreement. The LPSE seed beam amplification was $P_{\text{seed,out}}/P_{\text{seed,in}} = 1.94$ [2]. This test was run both with and without the CGT model described in Sec. 3.5.2 and at a variety of resolutions, to demonstrate how energy conservation at caustics are partially a resolution problem and how CGT mitigates for this somewhat.

Fig. 3.21 plots the results of the resolution scans both with and without CGT. The seed beam amplification after energy conservation iterations is plotted in Fig. 3.21.a, which demonstrates that even at lower resolutions, SOLAS is manages to reproduce the LPSE result to an accuracy of $\mathcal{O}(1\%)$. The energy conservation error from the first energy conservation iteration (*i.e.* after the first set of pump depletion iterations, but before energy conservation was achieved), is plotted in Fig. 3.21.b. This plot shows convergence in the physics at higher resolutions, although the converged solution (without enforced energy conservation) indicates small errors in the handling of the caustics, which energy conservation is able to correct in an *ad-hoc* manner. Fig. 3.21.c plots the caustic pump multiplier that was required to achieve energy conservation. Both Fig. 3.21.b and Fig. 3.21.c show that CGT improves the conver-

¹⁵The field from SOLAS simulation with the minimum resolution, $\Delta x \sim \lambda_0/20$ is plotted in Fig. 3.20.b.




Figure 3.21: Results from a resolution scan of the two-beam caustic CBET test both with CGT and without. Plotted in a), b) and c) are the power amplification of the seed beam, initial energy conservation error before energy conservation iterations and caustic pump multiplier used to obtain energy conservation respectively.




Figure 3.22: Results from the 16 beam CBET test problem. Panel a) plots the deposited power both without (left) and with (right) CBET from the simulation ran in cylindrical coordinates at $1 \mu\text{m}$ radial resolution. Panel b) plots the results of the resolution scan for both Cartesian and cylindrical coordinates. The circles and squares are the values of power deposition without and with enforced energy conservation respectively. The shaded regions illustrate the magnitude of energy conservation error from the initial energy conservation iteration, which can be thought of as the (single-sided) error on the circular points.

gence of the SOLAS solution when the caustic region width $\sim \lambda_0$ is not properly resolved by the mesh.

3.5.7.3 16 Beam CBET in 2-D

The final CBET test problem presented as validation is a planar, azimuthally symmetric CBET test with 16 beams in a direct-drive like plasma at the OMEGA scale [2]. This problem involves many beams, all of which undergo laser caustics and CBET scattering with multiple other laser sheets. It is also well suited to test the implementation of the SOLAS CBET model in non-Cartesian geometries, unlike the previous problems which did not have an angle of

symmetry. The same setup can be easily initialised in both Cartesian and cylindrical geometry to cross-compare the results.

All beams were pointed toward the origin with an initial starting radius of 1 mm, equally spaced around the circle with a width $\sigma = 386.8 \mu\text{m}$. The peak intensity of each beam was $I_0 = 0.68 \times 10^{18} \text{ Wm}^{-2}$ and each had an incident power $P_{\text{inc}} = 76.3 \text{ TW}$. The electron temperature was set to $T_e = 2 \text{ keV}$ and the ion temperature $T_i = T_e/2$. The electron density used the profile from Eq. 3.20, with the scale factor, $S = 1$. A radial velocity profile was used, which was obtained from a fit to a 1-D LILAC simulation,

$$\mathbf{u}(r) = c_s \left[M_0 + M_1 \log \left(\frac{r - x_0}{x_1} \right) \right], \quad (3.43)$$

where $M_0 = 1.41$, $M_1 = 1.37$, $x_0 = 204 \mu\text{m}$ and $x_1 = 343 \mu\text{m}$. The Cartesian simulations had a bounding domain $x, y \in [-800, 800] \mu\text{m}$, with a resolution ranging from $\Delta_{x,y} = 1 \rightarrow 12.5 \mu\text{m}$. The cylindrical simulations had domain $r \in [0, 800] \mu\text{m}$ and $\phi \in [-\pi, \pi]$, with radial resolutions $\Delta_r = 1 \rightarrow 12.5 \mu\text{m}$ and a fixed ϕ resolution of $\Delta_\phi = 0.025$. These simulations are compared to an LPSE simulation of the same problem below.

In the absence of CBET, LPSE calculated 90.5% power deposition. All SOLAS simulations agree with this result to within $\mathcal{O}(0.1\%)$. Including CBET, LPSE obtained an absorption of 52.7%. Deposited power from the $\Delta_r = 1 \mu\text{m}$ simulation is plotted in Fig. 3.22.a, both without CBET, and with CBET after energy conservation was achieved. Power deposition from the Cartesian simulation at $\Delta_x = 1 \mu\text{m}$ is not plotted, but looks almost identical, apart from the small grid artefacts. The reduction in magnitude of deposition is clearly visible, along with a slight radial broadening of the deposition region. The angled-line features in both plots are from the caustic region of the beam, where a 16-fold symmetry can be seen in both of the results, due to the beam geometry of the problem. The finer scale thin structures in both of these plots is due to enhanced deposition occurring at the beam caustic, where the ray amplitude swells to a maximum, enhancing the field strength.

Plotted in Fig. 3.22.b are the values of deposited power in the presence of CBET for both the cylindrical and Cartesian simulations, compared to the LPSE result. Circular markers are the values obtained for CBET deposition after the first set of pump-depletion iterations, *i.e.* before energy conservation was achieved. The larger square markers are the results after energy conservation was achieved. Shaded regions correspond to the magnitude and sign of the energy conservation error, $\Delta E_{j=0}$, before energy conservation was achieved, *i.e.*, for the circular points. This region corresponds to the deposition that could be achieved if any, arbitrary approach was taken to enforce conservation, for example if the value $\Delta E_{j=0} = \pm 10\%$ for one simulation, then up to 10% could be added to (subtracted from) P_{dep} . The trend, especially at lower resolution demonstrates that enforcing energy conservation brings the SOLAS result closer to LPSE, motivating the use of the approach in later chapters of this thesis. Although the cylindrical results typically appear to have a larger value for $\Delta E_{j=0}$ compared to the Cartesian results, the post conservation results are all deemed sufficiently close to both the Cartesian result and LPSE. The saturation of error for cylindrical geometry at higher radial resolutions seem to indicate that $\Delta E_{j=0}$ is dominated by the ϕ resolution in this region,




Figure 3.23: The initial conditions for OMEGA shot 89224. The shot was a cryogenic, DT implosion with materials and radii shown in panel a). Panel b) plots the incident laser pulse shape.

which was kept fixed across all simulations.

3.6 Simulations of CBET for OMEGA Direct-Drive Implosions

This Section presents CHIMERA-SOLAS simulations of OMEGA direct-drive implosions. In Sec. 3.6.1, 1-D Rad-Hydro simulations of OMEGA shot 89224 coupled to a 3-D ray-trace from SOLAS will be presented, both with and without the effect of CBET. This illustrates that the implementation of the CBET model in SOLAS correctly captures the loss of drive efficiency. In Sec. 3.6.1, 3-D SOLAS post-processes of spherically symmetric hydrodynamic data from the LILAC code are presented both with and without CBET. It is compared to results from the 3-D CBET codes IFRIIT and BEAMCROSSER. These simulations demonstrate that the SOLAS CBET model also correctly captures the enhancement of beam-mode drive asymmetry introduced by CBET in spherical illumination geometries.

3.6.1 1-D hydrodynamics, 3-D Ray-Trace Simulation of OMEGA Shot 89224

Table 3.2: Integrated metrics from simulations of OMEGA 89224. LILAC metrics are from Refs. [202, 203, 204].

	Yield ($\times 10^{14}$)	$\langle T_i \rangle$ (keV)	t_{bang} (ns)	V_{imp} (km/s)
LILAC, CBET	4.00	4.8	2.17	476
CHIMERA, CBET	1.78	4.2	2.16	490
CHIMERA, no CBET	6.77	6.2	1.94	661

A 1-D CHIMERA-SOLAS Rad-Hydro simulation of OMEGA 89224 was performed to illustrate the effect that CBET has on drive efficiency of direct-drive implosions. Drive efficiency is lost




Figure 3.24: CHIMERA-SOLAS simulation of OMEGA shot 89224. Panel a) shows the deposited power both from 2 separate simulations of 89224, one with and one without CBET as red and blue solid curves respectively. The shell radius, defined at the radius of maximum density, is also plotted on the right axis as dashed curves. Panel b) plots the time resolved partition of non-absorbed light for these simulations. Plotted in c) is the energy conservation error after energy conservation was achieved, alongside the value of f_{CP} that was used to achieve conservation.

primarily due a reduction in the magnitude of deposited energy. A secondary effect is that the radii where deposition occurs is pushed slightly further away from the ablation surface. This happens because beams only lose energy inside the CBET resonance volume, which for backscatter CBET occurs at the Mach-1 surface, as explained in Sec. 2.5.3. Therefore, the deposited power vs. radius curve is shifted to slightly higher radii. Both of these effects are captured by a 3-D CBET model coupled to 1-D Rad-Hydro.

The initial conditions for shot 89224 are plotted in Fig. 3.23.a, with vacuum outside the outer radius. The incident laser pulse is also plotted in Fig. 3.23.b, which has a total incident energy of 28.7 kJ. An experimental yield of 1.17×10^{14} neutrons was observed for this shot. A tabulated Sesame equation of state was used for each material [205] and a $P_{1/3}$ radiation transport algorithm was used with tabulated opacities and emissivities from the SPK

code [206]. Thermal conduction was solved with using a fully implicit algorithm for flux-limited Spitzer conductivities [207]. The electron flux limiter was set by using the default CHIMERA direct-drive method,

$$f_{\text{lim},e} = \begin{cases} 0.06, & \text{if in picket pulse} \\ 0.15, & \text{otherwise.} \end{cases} \quad (3.44)$$

The simulation domain was $r \in [0, 1680] \mu\text{m}$, with a fixed radial resolution $\Delta_r = 1 \mu\text{m}$. Two simulations were conducted, one with the effect of CBET on power deposition and one without. In order to load balance the ray-trace and discretise CBET interactions between different beams around the spherical 4π sr, SOLAS used a sparse angular grid with $N_\phi = 51$ and $N_\theta = 30$ azimuthal and polar cells respectively. At times earlier than 10 ps, a *cold start* routine was employed, where 50% of the incident power was deposited at the critical radius to create a finite scale length of underdense plasma, which rays could deposit and refract in. Ray-traces were conducted every hydrodynamic timestep, without computing CBET up to 150 ps¹⁶ after which a full ray-trace routine, including the effect of CBET was conducted every 10 ps. For hydrodynamic steps in between these CBET ray-traces, the most recently computed P_{dep} profile (normalised to 1 and then scaled to the incident power), was used as the laser source term for the timestep. Parameter scans of the ray-trace frequency confirmed that the behaviour of OMEGA simulations is well converged with this frequency.

The spatially integrated absorbed power is plotted in Fig. 3.24.a as solid curves. 62.8% and 83.9% of the incident energy was absorbed for the CBET and no-CBET simulations respectively. Also plotted as dashed curves are the shell trajectories from the two simulations. It can be seen that the reduction in drive efficiency leads to a significantly slower implosion with a delayed bangtime. Tab. 3.2 gives temporally and spatially integrated implosion metrics from the CHIMERA-SOLAS simulations, along with a comparison to a post-shot LILAC simulation which includes a 1-D model for CBET. Including CBET in the CHIMERA-SOLAS simulation gives significantly better agreement across all metrics, which is primarily due to correctly capturing the reduction in absorption. The most interesting metrics to compare across the codes when looking at the effect of CBET in 1-D are the implosion velocity and bangtime, which are in excellent agreement with LILAC. This is because these metrics are primarily functions of the absorbed laser energy, whereas the yield and burn averaged temperature additionally depend strongly on shock timing and compression [70, 208]. Shock timing between codes is particularly sensitive to differences in the thermal conduction modelling, equation of states, grid resolution and preheat modelling, all of which are likely to have been different between the LILAC and CHIMERA simulations. The strong agreement between the CHIMERA CBET simulation and the LILAC results for the bangtime and implosion velocity metrics provide evidence that direct-drive implosions can be accurately modelled without using artificial multipliers to the laser power.

Fig. 3.24.b plots the partition of the non-absorbed power for both simulations. This is divided between *blowby* light (B), which is light that left the computational domain travel-

¹⁶Note that the low intensity from this section of the pulse means that the CBET gains are negligible.




Figure 3.25: CBET deposition diagnostics from $t = 1.6$ ns into the CHIMERA-SOLAS simulation of OMEGA 89224, with CBET. Both sub-plots plot the power deposition with and without CBET, along with the g_{CBET} , which is defined as the sum of the absolute value of all power lost or gained by rays via CBET in the cell. The electron density normalised to critical density and radial velocity normalised to the sound speed are also plotted in panel a) and b) respectively. The blue shaded region in b) represents the full-width half-maximum of the g_{CBET} curve.

ling quasi parallel to the beam normal and *reflected* light, which leaves the domain travelling quasi antiparallel. Note that the plot is on a log axis. Examining first the blue curves which are from the no-CBET simulation, it can be seen that initially during the main pulse, from $t \sim 1.0 \rightarrow 1.6$ ns, the reflected light fraction decreases while the blowby goes up. This is because the critical radius converges, so the wings of the laser beam are less efficiently absorbed. The reflected light fraction increases sharply at $t \sim 1.6$ ns and $t \sim 1.8$ ns for the no-CBET and CBET simulations respectively. This corresponds to the time when the critical density crosses from the CH ablator material to DT fuel. DT has a significantly lower ionisation state compared to the ablator and therefore the electron-ion collision frequency decreases, corresponding to a drop in the Inv-Brem absorption and therefore an increase in the reflected light fraction. Both simulations show that blowby light is the most significant loss of coupling both with and without CBET, especially late in the implosion as the capsule converges.

Plotted in Fig. 3.24.c are the laser energy conservation, ΔE , and the modification to the caustic pump multiplier f_{CP} that was found to achieve energy conservation. The grey bounds indicate the accepted criteria to satisfy energy conservation, which was 0.5% of the incident power. The plot of f_{CP} illustrates that the value remained constant over many timesteps and therefore only a single set of pump depletion iterations was required for the majority of timesteps. This is not true late in the implosion, when the critical radius converges and CBET is responsible for $\sim 40\%$ reductions in deposited power. At these times, caustic CBET interactions are important and therefore ΔE is sensitive to the value of f_{CP} . As shown in Sec. 3.5.7.3 however, the iterations are believed to improve the fidelity of the modelling which is validated here by strong agreement with tuned, post-shot LILAC simulations.

Plotted in Fig. 3.25.a are the instantaneous power deposition profiles from the CHIMERA-

SOLAS simulation with CBET at $t = 1.6$ ns, alongside the n_e profile in the under-dense region. Note that the no-CBET profile plotted here is not from the no-CBET simulation of OMEGA 89224, but the power deposition from the field reconstruction ray-trace of the CBET simulation, *i.e.*, the deposition through the same hydrodynamic profiles, but without the effect of CBET on the power deposition. At this time, SOLAS estimates that CBET reduces the deposition versus by about 40% compared to the no-CBET deposition. No rays go beyond n_{cr} and therefore all deposition occurs outside of the critical radius. Also plotted in dark blue is the CBET scattering, g_{CBET} , defined as the sum of the absolute value of all power lost or gained by rays due to CBET in a SOLAS cell,

$$g_{CBET}(\mathbf{x}) = \sum_i^{\text{rays}} |\Delta P_{i,CBET}(\mathbf{x})|, \quad (3.45)$$

where $\mathbf{x}$ represents the position of the power gain/ loss, discretised within cells on the SOLAS grid. The deposition is broadly reduced only inside the peak of the scattering curve, because backscatter CBET transfers energy from inbound light to outbound light in direct-drive. Deposition is marginally increased outside the peak of the scattering curve, because the reflected and blowby light has been amplified due to CBET.

These power profiles are also plotted alongside the coronal radial velocity profile, normalised to the sound speed. The full-width half-maximum of the scattering curve is shown as the shaded blue region. Resonance occurs close to the Mach-1 surface as expected, which is where the plasma flow shifts the frequencies of radially inbound and outbound light such that they can excite an IAW. The peak scattering also is shifted slightly further out, because the outbound light mostly does not travel radially outward. Assuming the fluid CBET gain formulation and neglecting the $(\omega_j - \omega_i)$ in Eq. 3.24, it can be shown that non-radial light shifts the resonance to larger u_r and therefore greater radii.

3.6.2 3-D Post Process of Absorption Non-Uniformity

A final problem used to validate the CBET model from SOLAS was to compute the instantaneous 3-D effect of CBET on power deposition through spherically symmetric hydrodynamic data from a LILAC simulation of an OMEGA implosion. The same problem has been computed with both the IFRIT and BEAMCROSSER codes allowing cross-code validation of the 3-D deposition asymmetry. The LILAC profiles were taken from Ref. [67], and the IFRIT and BEAMCROSSER results are presented in Ref. [134]. CBET is known to enhance the beam-mode asymmetry which arises from the number and geometry of laser beams used to illuminate the target [147]. Correctly capturing this enhancement of asymmetry relies on a field reconstruction algorithm with good treatment of laser caustics and sufficient ray statistics in every cell across the computational grid to reconstruct the field everywhere that CBET is important.

This problem is a particularly stringent test of the grid discretisation for field reconstruction. The use of an underlying spherical polar grid for the SOLAS semi-structured mesh leads to cells that vary in volume from pole to pole. Some differences in the reconstructed field are




Figure 3.26: 3-D SOLAS post-process of spherically symmetric hydrodynamic data obtained from a LILAC simulation of an OMEGA direct-drive implosion. Figures a) and b) show the deviation in radially integrated deposited power without and with CBET respectively. Figures c) and d) plot the spherical harmonic modal decomposition for these maps alongside results from the IFRUIT and BEAMCROSSER codes without and with CBET respectively.

therefore inevitable between the capsule waist and pole purely from this grid discretisation. If this difference was insufficiently small, the power deposition with CBET would contain a conspicuous mode-2 which, for simulations integrated with the hydrodynamics, would lead to undesirable stagnation asymmetries, purely seeded from the computational mesh.

The simulation had a radial domain $r \in [350, 1362] \mu\text{m}$ with underlying grid resolution $\Delta_r = 1 \mu\text{m}$ ¹⁷. The full 4π sr was simulated, where an underlying polar mesh was used for the SOLAS grid with $N_\phi = 160$ and $N_\theta = 82$, which was then combined to give approximately equal area cells as described in Sec. 3.3.2. Fig. 3.26.a and Fig. 3.26.b show the deviation in radially integrated power deposition from the mean value for the no-CBET and CBET simulations respectively. The mode-10 pattern is clearly visible in both plots, although the magnitude is greatly increased by CBET, as expected. No clear mode-2 from the capsule pole to waist is seen which indicates that the SOLAS mesh is sufficiently close to equal-area to not result in significant grid artefacts.

A function, f discretised over azimuthal (ϕ) and polar (θ) angles can be decomposed into spherical harmonics,

$$f(\phi, \theta) = \sum_{\ell=0}^{\ell_{\max}} \sum_{m=-\ell}^{m=\ell} a_{\ell m} Y_{\ell}^m(\phi, \theta), \quad (3.46)$$

where $Y_{\ell m}$ are the spherical harmonic basis vectors and $a_{\ell m}$ are the associated coefficients

¹⁷Note that the adaptive radial resolution of the SOLAS mesh had a maximum radial resolution $\Delta_{r,\max} = 10 \mu\text{m}$.

[134]. The modal asymmetry of a given mode number, ℓ , is then defined as,

$$\sigma_\ell = \sqrt{\frac{1}{4\pi} \sum_{m=-\ell}^{m=\ell} a_{\ell m} a_{\ell m}^*}, \quad (3.47)$$

where the $a_{\ell m}^*$ is the conjugate of $a_{\ell m}$. Fig. 3.26.c plots the modal decomposition of the no-CBET radially integrated power deposition deviation from the mean, *i.e.* σ_ℓ up to mode $\ell = 20$. Excellent agreement is observed between SOLAS and the other codes. The dominant mode is $\ell = 10$, as is expected due to the beam port geometry.

Fig. 3.26.d plots the same, but now for the power deposition including CBET. As is also seen in the skymaps, the asymmetry is significantly amplified by CBET. Good agreement is again observed between SOLAS, IFRIIT and BEAMCROSSER across the spectrum. CBET appears to induce small amplitude modes other than just the $\ell = 10$ such as $\ell = 6$, $\ell = 12$ and $\ell = 18$. These are all multiples of 6, and therefore are likely induced by some OMEGA beam ports having 6 nearest neighbours [58, 57]. Small differences are observed between all codes, which is expected due to the different solvers and grids employed by each. SOLAS appears to slightly overestimate the modes 9 and 11, which have negligible values for both IFRIIT and BEAMCROSSER. This is potentially due to the ray shape function smearing routine used, that was described in Sec. 3.3.4, which smear power from a ray in one cell into neighbours. This could bleed some $\ell = 10$ into neighbouring modes. In future, a more robust spherical harmonic smoothing could be introduced to address this issue, although the effect is not likely to be significant for integrated simulations due to the small amplitude.

The incident and reflected sheet fields¹⁸ for a single beam both with and without CBET are plotted in Fig. 3.27. The incident fields without and with CBET are plotted in Fig. 3.27.a and Fig. 3.27.c respectively, and share the left colourbar. The inbound CBET sheet field shows a clear depletion of the field strength near the critical surface compared to the no-CBET field, which is the result of energy transfer from radially inward propagating light to outward travelling light via backscatter. This leads to the reduction in absorbed energy. Fig. 3.27.b and Fig. 3.27.d plot the reflected fields without and with CBET respectively, and share the right colourbar. The reflected sheet travels mostly radially outward and therefore has been significantly amplified by the CBET backscatter mechanism compared to the no-CBET case. In the absence of CBET, the field strength of the reflected sheet is significantly lower than the incident sheet field strengths for two reasons. Firstly, peak absorption occurs at the maximum radial penetration depth of the light into plasma, which is typically close to where rays experience their caustic and change sheet. Additionally, the rays diverge after reflecting from the critical surface leading to a reduction in ray amplitude and therefore field strength.

¹⁸Note that incident and reflected here refer to which sheet the ray belongs to, *i.e.* if it is before or after the caustic location.




Figure 3.27: Electric fields from the 3-D post-process. Panels a), b), c) and d) plot the incident sheet without CBET, the reflected sheet without CBET, the incident sheet with CBET and the reflected sheet with CBET respectively. The critical surface is plotted with the inbound sheets as a transparent spherical wireframe. Note that the outbound sheets are plotted on a log scale.

3.7 Conclusions

This chapter presented the implementation and validation of a 3-D ray-trace and CBET model, SOLAS, for the 3-D Rad-Hydro code CHIMERA. Prior to this work, a simple 1-D ray-trace was used to model the laser for direct-drive implosions which was strongly detrimental to the ability of the code to accurately model these experiments. Radially tracing rays inward, even for 1-D Rad-Hydro simulations, misses the important interplay between the convergence of the target and loss of deposition. The lack of a CBET model, which is responsible for significant energy coupling loss at OMEGA laser facility energy scale implosions, also meant that deposited power had to be reduced in an *ad-hoc* manner to obtain agreement with experiments. These factors severely limited the predictive capability of the code to model these experiments.

Sec. 3.3 described the implementation of the 3-D ray-trace in the SOLAS laser package. An adaptive solver is used to trace rays through a semi-structured Eulerian mesh in Cartesian, cylindrical or spherical geometries. Validation of this implementation was also presented. The 3-D laser ray-trace enables CHIMERA to model direct-drive implosions and other laser-driven HEDP experiments. It is also a crucial prerequisite to a ray-based CBET model.

Issues with ray-based CBET models often stem from a poorly reconstructed laser electric field profile. The approach taken in SOLAS to obtain the electric field was described in

Sec. 3.4, which is to trace a set of neighbouring rays around each ray and then relate the area of the ray to the field strength. Validation problems were conducted for this section which illustrated that this approach yields excellent agreement with the wave-based solver LPSE. Caustics are also identified and the field near them is capped to a diffraction limited value. Validation problems showed that the field solver was robust and suitable for use within the CBET model.

The implementation of the SOLAS CBET model was described in Sec. 3.5. This model employs the linear kinetic or fluid gain formulations and the reconstructed field to estimate power changes of rays due to CBET. Fixed point iteration is then used to find the equilibrium solution and additional energy conservation iterations can be used to account for small errors in the caustic field reconstruction. Validation problems against LPSE and the ray-based BEAMCROSSER post-processor demonstrated that SOLAS correctly computes CBET for direct-drive conditions in Cartesian and cylindrical geometries.

A 1-D, spherical Rad-Hydro, CHIMERA-SOLAS simulation of the OMEGA shot 89224 was presented in Sec. 3.6.1, alongside a comparison to a LILAC simulation. The CBET model behaved as expected, significantly reducing the power deposition and slowing the implosion. Excellent agreement was obtained between CHIMERA-SOLAS and LILAC for integrated parameters which primarily depend on the amount of energy coupled to the target. Sec. 3.6.2 presented a 3-D post-process of spherically symmetric hydrodynamic data from a LILAC simulation. The power deposition asymmetry was compared to the IFRIT and BEAMCROSSER codes and good agreement was found with those models. CBET was shown to significantly enhance the deposition asymmetries as expected. The results from this section showed that the SOLAS CBET model functions as expected in spherical geometry and therefore CHIMERA-SOLAS is capable of performing 3-D spherical direct-drive implosions with CBET which could previously only be conducted with the ASTER-IFRIT code combination.

4 Cylindrical Simulations to Study the Effect of Beam Radius in Direct-Drive

This chapter describes a cylindrical, direct-drive implosion simulation platform and its use to study the effect of the beam radius initial condition on OMEGA laser facility experiments. Although results from the cylindrical simulations do not have the same convergence properties of spherical implosions, the essential physics for studying the effect of beam radius is preserved. The main benefit of the geometry is that a 2-D ray-trace can be used to model the lasers, which yields several orders of magnitude speed-up, compared to spherical 3-D implosions. The reduction in computational expense enables an ensemble of Cross-Beam Energy Transfer (CBET) simulations to be performed, which would be exceedingly expensive for 3-D spherical calculations. Beam radius strongly affects CBET and therefore including a model for the interaction in computational studies is crucial [142].

The chapter begins with a review of existing literature concerning the beam radius initial condition for direct-drive implosions, with an emphasis on the use of this parameter in statistical modelling of OMEGA campaigns. A description of the cylindrical simulations is then provided, which includes a discussion of its advantages, weaknesses and applicability to current OMEGA experiments. A tuning procedure which was followed to obtain hydrodynamically similar implosions at different beam radii is then described. The main results of the chapter are then presented, which include calculations of the power deposition asymmetry both with and without CBET and an explanation of why CBET typically amplifies the asymmetry. CBET is also shown to introduce *modal-flips* of the deposition, where the power deposition asymmetry inverts. Stagnation state asymmetries of the hydrodynamic profiles are then studied for all implosions and these demonstrate that while increasing beam radius in the absence of CBET reduces beam-mode asymmetries, the opposite behaviour is observed in the presence of CBET, although the exact relationship is complex. The chapter concludes with a summary of the simulations and suggestions for further work that could be undertaken using the same cylindrical platform are provided.




Figure 4.1: The trajectory of rays from two beams through a direct-drive like coronal plasma. The density profile for both simulations is $n_e = n_{cr} \exp[-(r_{\mu m} - 20)/100]$. Panels a) and b) plot rays from beams with widths $\sigma = 10$ and $18 \mu m$, respectively. Ray trajectories are separated for each beam by colour depending on their sheet. Rays coloured in red and dark blue are from the incident sheet (before the ray caustic) and rays coloured in magenta and light blue are from the reflected sheet (after the ray caustic).

4.1 Introduction to Beam Radius in Direct-Drive Inertial Confinement Fusion

Neglecting the effect of random shot-to-shot variations, the dynamics and performance of an idealised direct-drive implosion is defined by a limited number of initial conditions. The target can be described by a set of materials and their thicknesses. Initial target parameters are intimately coupled to the physics of the implosion and, in part, dictate the propagation time of shocks through the target, hydrodynamic stability and absorption of the laser energy. The pulse shape defines the laser power which is incident on the target as a function of time. It can be designed to, for example, drive shocks by introducing sharp rises in the incident power with time, which leads to large changes of ablation pressure [27]. A given facility also has a number of beam ports, each of which has a specific origin and pointing location, that influence the magnitude of the *beam-mode* asymmetry, arising from the uniformity of laser absorption. The focal-spot intensity profile of each laser, specifically the beam radius, is an additional parameter which can be varied. It plays an important role in defining both the power which can be coupled to the target and the magnitude of beam-mode asymmetry.

As shall be explored in this chapter, increasing the beam radius alters the magnitude of energy lost via CBET, leading to a reduction in the maximum target mass that can be imploded at a given speed [142]. The width of the beam relative to the target is therefore often effectively varied from shot to shot, by changing the outer radius (and therefore mass) of the target, while fixing the width of the beam. This defines a dimensionless variable, which is the radius of the beam divided by the target radius, R_b/R_t . Typically, at the OMEGA laser facility,

this is explicitly defined as the radius of the beam which contains 95% of the incident power divided by the initial outer target radius [142, 152, 209],

$$R_b/R_t = \frac{r_{95}}{R_t}, \quad (4.1)$$

where r_{95} is defined by the integral,

$$\int_0^{r_{95}} e^{-|\frac{r}{\sigma}|^{n_s}} r dr = 0.95, \quad (4.2)$$

and the definition of a circular, super-Gaussian beam profile from Eq. 3.1 has been used. In the absence of CBET, it can intuitively be understood that increasing this parameter should improve the uniformity of the laser illumination, because beam spots overlap each other more on the target, and therefore reduce the beam-mode [73]. A larger R_b/R_t also leads to slightly less absorption in the absence of CBET, because a larger fraction of the incident light (especially at late time as the target converges) would reach lower density plasma and therefore not be absorbed. CBET significantly complicates this interpretation.

Fig. 4.1.a and Fig. 4.1.b plot results of ray-tracing calculations, in a direct-drive relevant, exponentially decaying density profile, with a smaller and larger beam, respectively. In direct-drive, backscatter CBET is the dominant mechanism that depletes absorption, which occurs when outbound light gains energy from inbound light¹. The outward rays from the small beam radius simulation in Fig. 4.1.a do not overlap the incident light from the other beam and therefore limited CBET between these beams will occur. The trajectories from the larger radius simulation in Fig. 4.1.b, however, do cross the inbound rays from the other beam, which could lead to a resonant CBET interaction, and significant reduction of the absorbed power. As was shown in Fig. 3.26, CBET substantially increases absorption asymmetry on the OMEGA laser facility. This means that in the presence of CBET, the effect of increasing R_b/R_t on illumination asymmetry is not clear. While in the absence of CBET, increasing R_b/R_t should lead to greater beam overlap, and therefore absorption symmetry. However, when accounting for CBET, increased overlap will result in more CBET, which degrades absorption uniformity.

Isolating the contribution of CBET in particular is important to enable extrapolation of experimental results to future facilities. It is hoped that next generation direct-drive facilities will have broad bandwidth lasers² which are expected to almost entirely eliminate CBET [149, 150, 153]. Studying how CBET specifically affects absorption and stagnation symmetry can not currently be done experimentally, as significant backscatter CBET occurs at all laser facilities which are capable of conducting compression experiments. Therefore, computational studies are well suited to investigate how R_b/R_t influences the role of CBET.




Figure 4.2: Soft x-rays emitted from the ablation surface of direct-drive implosions with various R_b/R_t values, as measured by an x-ray framing camera. All images were taken at a constant capsule radius of $R = 175 \mu\text{m}$. The figure has been reproduced with permission from Ref. [142].

4.1.1 Previous Work Studying the Effect of Beam Radius on OMEGA

Experimental and computational work has been conducted to explore the effect that the beam radius has on direct-drive implosions. Froula *et al.* conducted a series of implosions which systematically varied R_b/R_t to explore the balance between loss of coupling due to increasing CBET at large radii, and increased illumination non-uniformity at lower radii [142]. Soft x-ray emission data from a selection of implosions with different radii are plotted in Fig. 4.2, all of which are taken at the same target shell radius. These images show that at lower values of R_b/R_t , mid-mode perturbations³ become increasingly significant. The results of these experiments found that neutron yield was maximised at $R_b/R_t \sim 0.8$. 1-D modelling using the CBET model in LILAC was in good agreement with the experimental results, verifying that CBET was responsible for the decrease in coupled energy to the target [175]. The DT-liner campaign on OMEGA has also significantly enhanced the absorption of laser energy, by using larger outer radius capsules to reduce R_b/R_t and thus mitigate CBET [75]. This led to the recent demonstration of hotspot fuel gain exceeding unity for the first time in a direct-drive implosion [76].

During direct-drive implosions, the target implodes radially inward, therefore the critical radius decreases with time. At the OMEGA laser facility, beam radii are fixed for a single shot, and therefore it makes sense to parameterise the initial condition by the ratio of the beam radius to the initial target radius. A promising laser optics technique known as *zooming*, could significantly enhance performance by reducing the focal spot of the laser to track the critical radius as the target implodes [145]. The reduced beam radius late in the implosion reduces blowby light and leads to more deposition closer to the ablation surface, which enhances the ablation pressure and overall performance. Simulations by Trickey *et al.* in Ref. [208] have demonstrated that, assuming full CBET mitigation, zooming can enhance ablation pressures

¹‘Outbound’ here, means light travelling quasi-parallel to the approximately radially outward fluid velocity, and ‘inbound’ means quasi-anti-parallel.

²It is expected that $\Delta\omega/\omega_0 \sim 1\%$ will be sufficient to eliminate CBET in direct-drive for frequency-tripled lasers.

³In direct-drive, *mid-modes* are loosely defined as modes similar to $\ell = 10$.

by $\sim 50\%$ for ignition scale direct-drive experiments. If zooming is employed without full CBET mitigation, the fractional increases would be substantially larger, because CBET losses scale very strongly with R_b/r_{cr} , which is the ratio of beam radius to critical radius [152]. 2-D modelling has also demonstrated that zooming the beam to $\sim 60\%$ of its initial radius effectively mitigates CBET while maintaining low-mode uniformity [144].

4.1.2 Statistical Modelling of OMEGA Direct-Drive Implosions

In recent years, a lot of work has been carried out to develop a statistical modelling capability for direct-drive implosions on the OMEGA laser facility. This modelling serves several critical purposes, including enhancing the predictive capability of simulations [72], guiding experimental design to achieve higher performance implosions [74], identifying important sources of degradation on current facilities [73, 210] and validating simulation codes to help ensure they produce physically relevant results [211]. The first generation statistical model, described by Lees *et al.* in Ref. [72], uses a high dimensional fit, to create a mapping between experimental and 1-D simulation results in order to explain significant differences in their results. A next-generation statistical model has recently been developed by Ejaz *et al.*, which uses auto-encoders to reduce the design-space dimensionality and subsequently transfer learning to calibrate simulation outputs to experimental data. The subsequent review shall however, focus on the more widely studied results of the first model.

In the model described by Lees *et al.* [72], 1-D simulation results are fed into the model and degraded by a series of power law multiplications, which returns a more physically accurate experimental yield. Each power law multiplication represents a physical process for yield degradation with respect to 1-D physics, not included in the simulation. Each of these is termed a *yield-over-clean* (YOC_i), where the i refers to each source of degradation. The neutron yield from the 1-D simulation (Y_{1-D}^{sim}) can thus be converted to a prediction of the experimental yield [72],

$$Y^{exp} = (YOC_h YOC_f YOC_{\ell=1} YOC_b YOC_{res}) Y_{1-D}^{sim}, \quad (4.3)$$

where YOC_h is a degradation term from hydrodynamics and instability growth, YOC_f is degradation due to radioactive decay of the tritium fill, $YOC_{\ell=1}$ is degradation from $\ell = 1$ modes, YOC_b is degradation from a finite number and radius of beams and YOC_{res} is a residual size scaling, which is required to reduce performance of hydrodynamically downscaled implosions [212]. Each of these terms and their functional forms shall be discussed briefly, in order to provide context for the utility of the model and highlight that understanding the relevant physical processes which lead to degradation can improve its performance.

Hydrodynamic Degradation 1-D simulations do not capture short wavelength perturbations which grow via the Rayleigh–Taylor Instability (RTI) and reduce the yield of experiments by puncturing and breaking up the shell as the capsule implodes inwards. Instabilities may be seeded by laser imprint [213, 214, 215], or small scale defects in the target materials [216, 217]. Degradation can be reduced either by altering implosion design to in-

crease the shell adiabat which increases the ablative stabilisation of the RTI [218, 219], or by lowering the In-Flight Aspect Ratio (IFAR), which increases the distance that the instability must grow through to puncture the shell [220]. The degradation term includes scaling with the target convergence ratio, $C_R \equiv R_0/R_{\text{stag}}$, and with the ratio of outer to inner shell radius, $\hat{D} \equiv R_{\text{out}}/R_{\text{in}}$, which is believed to compensate for inaccuracies in modelling the shock propagation speed through the target. The hydrodynamic degradation term thus has the functional form,

$$\text{YOC}_h = \left[\frac{(\alpha/3)^{1.1}}{\text{IFAR}/20} \right]^{\mu_1} C_R^{\mu_2} \hat{D}^{\mu_3}, \quad (4.4)$$

where μ_i are the fitting parameters, which are obtained from nonlinear regression across many OMEGA shots. The fitting procedure demonstrated that experimental yields are very significantly reduced by these hydrodynamic degradations, with the most unstable shots yielding values of $\text{YOC}_h \sim 0.1$ [73].

Fill Age Degradation OMEGA cryogenic implosions contain a DT fuel gas fill with a surrounding ice layer. The tritium in this fuel is unstable and undergoes radioactive decay to ^3He . This occurs over a period of days to weeks, which typically elapse between the initial gas filling procedure to the shot day [221]. ^3He has a lower freezing temperature than DT and thus sublimates, accumulating in the fill region. The accumulation of helium in the gas reduces the final yield, both by increasing radiative losses due to its higher ionisation, and by increasing density of the vapour, which reduces compressibility and thus stagnation pressure in the hot-spot. Both of these effects can be captured by conducting 1-D simulations with a ^3He concentration (and corresponding reduction of tritium density) which is a function of fill age. The yield over clean due to the fill age and radioactive decay can then be taken as the ratio of these 1-D simulation yields,

$$\text{YOC}_h = \left(\frac{Y_{1\text{-D,He}}^{\text{sim}}}{Y_{1\text{-D}}^{\text{sim}}} \right)^{\mu_4}, \quad (4.5)$$

where $Y_{1\text{-D,He}}^{\text{sim}}$ is the yield from the 1-D simulation with accumulated ^3He and μ_4 is a fitting parameter. Good agreement is observed with a fitted parameter value of $\mu_4 = 1.3$. The value is larger than 1 (and the 95% confidence interval does not include 1), which suggests stronger degradation than observed in 1-D calculations. This could be due to radioactive decay damaging the shell, leading to increased hydrodynamic instability growth [73].

Mode 1 Degradation In direct-drive implosions on the OMEGA facility, $\ell = 1$ modes can be introduced to an implosion by a global offset of the capsule from the target chamber centre, mispointing of the laser beams or a power imbalance. These are random and uncontrollable and therefore the statistical models can only account for their effect after the shot has occurred. By conducting post-shot analysis of an implosion, the model returns an estimated yield which could have been achieved, if no $\ell = 1$ were present, $Y^{\text{exp}}/\text{YOC}_{\ell=1}$. Mode 1 asymmetries have a clear signature in the broadening of the neutron time-of-flight detector peaks, when observed from orthogonal lines of sight [136]. The width of the peaks from multiple




Figure 4.3: Experimentally inferred fusion yield degradation due to the finite beams on the OMEGA laser facility. The dotted orange curve is the fit obtained from just using the $\bar{R}_{b/t}$ relation, while the black dashed curve uses the full relation in Eq. 4.7. The vertical dotted line indicates the critical threshold, $\hat{R}_{b/t}^{\text{crit}} = 0.86$, after which the $\hat{R}_{b/t}$ also has an effect. The cross validation error from the $\bar{R}_{b/t}$ and full fit is -1.0% and -0.5% respectively. The figure has been reproduced with permission from Ref. [73].

lines of sight can be analysed to return an angularly resolved apparent ion temperature map, the asymmetry of which is dominated by the lowest mode of the hot-spot [222]. Thus, it is deduced that the ratio of the maximum to minimum apparent ion temperatures from the experimental neutron time-of-flight signal can be used as a proxy for the amplitude of the mode 1, $R_T = T_{\text{max}}/T_{\text{min}}$. This leads to a yield over clean expression for the $\ell = 1$ degradation source,

$$\text{YOC}_{\ell=1} = \hat{R}_T^{\mu_5}, \quad (4.6)$$

$$\hat{R}_T \equiv \max \left(\frac{R_T}{R_T^{\text{min}}} \right),$$

where μ_5 is the fitting parameter and the minimum threshold value, R_T^{min} is introduced due to imperfect reconstruction of the apparent ion temperature map and is fitted separately. Work has been conducted to minimise the effect of the $\ell = 1$ on OMEGA by repositioning the target after several initial shots to minimise the asymmetry in the apparent ion temperature measurement and thus increase performance [136].

Finite Beam Degradation The OMEGA laser facility has 60 beams arranged around a sphere, which generally gives good illumination uniformity on a hard sphere surface, less than the 1% deviation which is believed to be necessary to achieve ignition [127, 140]. An $\ell = 10$ remains in the deposition however, as is demonstrated in Fig. 3.26, which is often referred to

as the beam-mode. In the absence of CBET, increasing R_b/R_t increases the hard-sphere illumination uniformity [210]. As already described however, increasing beam radius leads to more blowby light and therefore more CBET. This reduces the coupled energy and potentially introduces additional asymmetry to the implosion. Additionally, increasing the overlap of beams on the target could reduce the amplitude of the imprint seed and therefore increase performance. The uncertainty as to which physical mechanisms are important, is highlighted by the complexity of the degradation parameter compared to those already described,

$$\begin{aligned} \text{YOC}_b &= \left(\bar{R}_{b/t} \right)^{\mu_6} \left(\hat{R}_{b/t} \right)^{\mu_7}, \\ \bar{R}_{b/t} &= \begin{cases} R_b/R_t & \text{if } R_b < R_t, \\ 1 & \text{if } R_b \geq R_t, \end{cases} \\ \hat{R}_{b/t} &= \begin{cases} \frac{R_b}{R_t R_{b/t}^{\text{crit}}} & \text{if } R_b/R_t < R_{b/t}^{\text{crit}}, \\ 1 & \text{if } R_b/R_t \geq R_{b/t}^{\text{crit}}, \end{cases} \end{aligned} \quad (4.7)$$

where μ_6 & μ_7 are fitting parameters, the threshold behaviour in $\bar{R}_{b/t}$ was chosen to fit a small number of shots (< 10) at $R_b/R_t > 1$ and the threshold behaviour at $R_{b/t}^{\text{crit}}$ was introduced to fit a physically unexplained transition between two regimes in the data.

The fitted curve from the model is shown in Fig. 4.3, as the black dashed curve alongside the inferred values from experimental data points. Also plotted in orange is a fitted curve obtained from just using the simple $\bar{R}_{b/t}$ degradation. Introducing the $R_{b/t}^{\text{crit}}$ threshold significantly reduces the cross validation error of the fit. The switch between the two regimes is found from the fitting procedure to occur at $R_{b/t}^{\text{crit}} = 0.86$. This is close to value of minimum illumination asymmetry for beams incident on a hard hard-sphere ($R_b/R_t = 0.82$), which suggests that the degradation at the lowest values of R_b/R_t is dominated by beam-mode, however this has not been experimentally or computationally verified. Experiments between $R_{b/t}^{\text{crit}} < R_b/R_t < 1$ could be influenced by changing behaviour due to CBET or imprint, which is not properly captured by the 1-D LILAC simulations included in the model.

The change in behaviour observed at $R_{b/t}^{\text{crit}} = 0.86$ is more clearly visible in recent extensions to the statistical model [223]. This extension of the model enables predictions of the absolute yield from experiments, rather than the fraction of expected yield with respect to 1-D simulations [223]. The prediction yield with respect to beam radius over target radius for the extended model are plotted in Fig. 4.4, which clearly demonstrates a transition in behaviour at $R_b/R_t = 0.86$.

The hypothesis tested in this chapter is that the change in yield over clean at $R_b/R_t = R_{b/t}^{\text{crit}}$ is due to increasing CBET as beam radius increases, which increases beam-mode asymmetry and therefore suppresses the YOC_b term. Although LILAC does include a model for CBET, it is a 1-D code and therefore the 3-D beam-mode perturbations cannot be inferred from its results. Qualitatively, this hypothesis can explain the observed behaviour in Fig. 4.3, which demonstrates that at $R_{b/t}^{\text{crit}}$, the gradient of YOC_b decreases. This behaviour should occur if CBET acted to amplify the asymmetry of the stagnation state, causing the simulation to be less similar to the 1-D LILAC results, than if CBET did not significantly amplify asymmetries.




Figure 4.4: Predictions for how the absolute yield varies across OMEGA direct-drive experiments with R_b/R_t . A change in behaviour is clearly seen at $R_{b/t}^{\text{crit}} = 0.86$, at which point the yield trend with R_b/R_t flips from monotonically increasing to decreasing. The data in this figure is obtained from Ref. [223].

4.2 Cylindrical Simulation Platform for Beam Radius Parameter Scan

The method employed to study whether CBET induced beam-mode asymmetry at $R_b/R_t \gtrsim R_{b/t}^{\text{crit}}$ contributes to the second distinct regime from Eq. 4.7, was to conduct a series of cylindrical direct-drive simulations. These simulations are in a different convergence regime to spherical implosions, but retain much of the key physics relevant to the study, such as CBET, target convergence and beam-mode perturbations of the target. Crucially, it circumvents the large computational run-times of 3-D spherical simulations which include a CBET model. Inspiration for the cylindrical implosion platform was taken from the work of Follett *et al.* in Ref. [2], where a similar, 16-beam, planar configuration was used for a 2-D CBET test problem. The SOLAS result for this validation problem is presented in Sec. 3.5.7.3.

4.2.1 Advantages and Validity Considerations of the Cylindrical Simulation Platform

Ideally, fully 3-D Radiative-Hydrodynamics (Rad-Hydro) simulations, coupled to a 3-D CBET model would be used for this work. This would retain the spherical convergence of the implosion and the 3-D nature of the beam-mode perturbation growth through stagnation, while also capturing how CBET influences these effects. These simulations are extremely expensive however and can take months to complete. A 2-D ray-trace (where rays can only move in two cardinal directions rather than 3) leads to large reductions in cost.

Ray-based CBET models require a ray from every laser sheet to pass through every computational grid cell where CBET scattering is significant. Therefore, by reducing the dimen-

sionality of the problem, savings are made proportional to the reduction in number of cells, which is $\mathcal{O}(100)$ from 3-D spherical direct-drive calculations to the 2-D simulations presented in this chapter. Additionally, 10 beams were used to produce a mode-10, rather than the 60 on OMEGA, yielding another factor of 6 fewer rays. Each of the eight 2-D CBET simulation presented here took $\mathcal{O}(10^3)$ CPU hours, whereas (assuming the above logic), $\mathcal{O}(10^6)$ CPU hours would be required for each equivalent 3-D spherical simulation. On 1000 processors, each 3-D simulation would therefore take over a month to finish, as opposed to these simulations which all ran to completion within a day on 128 cores. Note that the above scaling assumes that the laser ray-trace and CBET computation costs dominate the simulation run-time, which is found to be the case for all CHIMERA-SOLAS simulations.

The physics of the implosion is however different in cylindrical as opposed to spherical geometry. Firstly, the mass converges in only two directions rather than three as the target implodes, which results in increased convergence at stagnation, potentially altering the beam-mode asymmetry growth. The 2-D cylindrical perturbations will also only evolve in the simulation plane, unlike the true 3-D case where they interact with material ‘above’ and ‘below’ them as well. This could lead to the cylindrical simulations overpredicting the beam-mode degradation compared to 3-D spherical simulations. In the corona, the expanding plasma also diverges in only two directions rather than three as it rockets away from the capsule, which leads to reduced density gradients in the corona, where the laser propagates and deposits energy. This could have the effect of shifting the deposition to greater radii above the critical surface, reducing the drive efficiency.

Despite these differences, the simulations were, qualitatively, sufficiently hydrodynamically similar to spherical implosions, to suggest that parallels can be drawn from the work done here to spherical implosion data. Cylindrical implosion experiments are also conducted on laser facilities, the results of which are often explicitly extrapolated to spherical implosions, suggesting that the approach is valid and accepted within the community [224, 225]. Although the differences between the two regimes are not quantified in the work conducted in this chapter, future simulations could be performed, which extend the platform to a 2-D ‘plane’ in spherical rather than cylindrical geometry, which would capture the spherical convergence of the target and divergence of the coronal flow. This is discussed in more detail at the end of the chapter, in Sec. 4.5.

4.2.2 Pulse Shape and Target Initial Conditions

The base simulation initial conditions are plotted in Fig. 4.5. The initial conditions are symmetric about the azimuth and in the plane normal to the $\hat{z}$ direction. A target with the same initial layer radii as OMEGA shot 89224 was constructed with a DT gas fill, a layer of DT ice and a CH plastic ablator with vacuum outside, shown in Fig. 4.5.a. Ten beams were placed around the target, equally spaced in azimuthal angle and all were polarised in the out-of-plane, $\hat{z}$ direction. All beams had super-Gaussian spot profiles with $n_s = 5.2$ and σ set by Eq. 4.2. A simple 2.5 ns square pulse (including a 0.1 ns ramp to full power) was used for all simulations, plotted in Fig. 4.5.b. The maximum power of the pulse for each 2-D simulation




Figure 4.5: The a) target initial conditions with beam geometry, and b) pulse shape used for the 2-D cylindrical simulations. All beams were polarised out of the simulation plane, in the $+\hat{z}$ direction. Initial layer radii were taken from the initial conditions for OMEGA shot 89224, presented in Fig. 3.23.a.

was tuned from a separate set of 1-D simulations, such that the bangtime occurred at 2.5 ns.

By tuning the simulations such that the bangtime was consistent across all simulations, the coupled energy and implosion velocity was kept the same across all implosions. The difference between implosions was therefore primarily due to differences in the spatial location of the deposited power. If the incident energy were fixed, increasing R_b/R_t would lead to more CBET, which would result in less energy coupled to the target. Therefore, to compare simulations, the target would also have to be altered to reduce the imploding mass with less coupled energy. This was deemed beyond the scope of the work presented in this chapter and therefore only the incident energy was altered to maintain the 1-D implosion hydrodynamics.

Every simulation in this chapter used a grid with radial extent $r \in [0, 1600]$ with resolution $\Delta_r = 1 \mu\text{m}$ and 256 cells in the azimuthal direction. A tabulated Sesame table of state was used for each material [205] and thermal conduction routine was solved using an Alternating-Direction Implicit (ADI) method with flux-limited Spitzer conductivities [226]. The electron flux limiter was set using the default CHIMERA direct-drive setting, outlined in Eq. 3.44. Radiation transport was not included in simulations, because the small cells on the $r = 0$ boundary led to significant computational expense. 1-D calculations showed that the effect of including radiation transport was relatively small. The dominant effect was that radiation transport reduced the bangtime by ~ 0.1 ns, primarily due to temperature losses in the corona. Future work could therefore include the dominant radiation effect by using a radiative loss model rather than full transport.

4.2.3 1-D Implosion Tuning

As already mentioned, the energy of the laser was varied to maintain a consistent bangtime across all R_b/R_t values, so that the target parameters did not have to be separately

Table 4.1: Results of the 1-D Tuning Simulations.

R_b/R_t		$P_{\max}$ (TW/cm)	I_0 (10^{14} W/cm 2)	t_{bang} (ns)	Y_{DT} (10^{13} /cm)
0.75	No CBET	54.44	0.85	2.49	1.53
	CBET			2.51	1.44
0.80	No CBET	58.25	0.83	2.49	1.56
	CBET			2.51	1.45
0.85	No CBET	63.44	0.83	2.48	1.67
	CBET			2.51	1.43
0.90	No CBET	70.00	0.85	2.47	1.82
	CBET			2.50	1.41
0.95	No CBET	77.94	0.89	2.46	1.99
	CBET			2.49	1.49
1.00	No CBET	87.25	0.93	2.45	2.15
	CBET			2.49	1.60
1.05	No CBET	97.94	0.99	2.46	2.27
	CBET			2.50	1.61
1.10	No CBET	110.00	1.06	2.47	2.31
	CBET			2.51	1.52

optimised for each simulation. This was done via a series of 1-D, with-CBET simulations which varied the maximum power of each beam, $P_{\max}$, at each R_b/R_t to obtain an implosion with $t_{\text{bang}} = 2.50 \pm 0.01$ ns. For no-CBET simulations, the absorbed power vs time from the CBET simulation with the same R_b/R_t was enforced. Thus, when comparing any two simulations, the absorbed energy is equivalent, but the spatial location of the deposition is different. This manifests both as different azimuthal asymmetries in the deposition, which alters the stagnation state asymmetry, and different radial location of absorption. For example, the with-CBET power deposition occurs at slightly larger radii compared to no-CBET profiles, as is shown in Fig. 3.25. This means that the no-CBET implosions have a slightly increased drive efficiency compared to CBET implosions, because thermal conduction does not have to transport energy as far from the absorption region to the ablation surface.

Tab. 4.1 shows implosion metrics from all the tuned 1-D implosions. As shown, the incident maximum power of each beam for the CBET simulations, $P_{\max}$, increases with increasing R_b/R_t because larger R_b/R_t leads to more CBET and therefore less absorption, so more incident power is required to maintain the same absorbed energy. The maximum intensity of each beam at peak power, I_0 is non-monotonic, because although the maximum power increases, the beam radius also increases, which limits the increase in maximum intensity. Bangtimes and yields are similar across all simulations. Note that with increasing R_b/R_t , bangtime and yield difference between the CBET and no-CBET results at the same R_b/R_t increase. This is because more CBET occurs for the larger R_b/R_t simulations and there-




Figure 4.6: Streak plots from two of the 1-D tuning simulations including CBET. Panels a) and b) plot the electron density of the $R_b/R_t = 0.8$ and $R_b/R_t = 1.0$ simulations, respectively, as a function of time (x -axis) and radius (y -axis). Panels c) and d) plot the electron temperature for the $R_b/R_t = 0.8$ and $R_b/R_t = 1.0$ simulations, respectively.

fore the difference in deposition radius increases between CBET and no-CBET simulations, marginally improving the effective drive efficiency of the no-CBET results.

Streaks, which plot hydrodynamic quantities as a function of radius and time, of n_e and T_e are shown in Fig. 4.6 for 1-D, with-CBET simulations at two separate R_b/R_t values. Qualitatively, the 1-D implosion trajectories from these plots are similar. Small differences in shock timing exist between simulations, as is evidenced by the initial shock for the $R_b/R_t = 0.8$ hitting the $r = 0$ axis at $t \sim 1.9$ ns, which is about 0.1 ns earlier than the $R_b/R_t = 1.0$ simulation ($t \gtrsim 2.0$ ns). Despite these small differences in both the metrics from Tab. 4.1 and the streaks from Fig. 4.6, all simulations were deemed sufficiently similar that implosions could be cross-compared.

4.3 Asymmetry of Deposited Power

This section describes the asymmetry of the deposited power profile for simulations with and without CBET. The effect of this asymmetry on the in-flight and stagnation state hot-spot profiles are discussed in Sec. 4.4. Analysis of the deposited power profile shows that the growth of asymmetries in the target is the result of a complex, space- and time-dependent evolution of the deposition. In the absence of CBET, *modal-flips* of the deposition occur, where the phase of the driving asymmetry flips in time. This is due to the overlapping beam intensity changing in the region where Inverse-Bremsstrahlung (Inv-Brem) deposition is important. The pattern of these modal flips depends on the width of the beams, the time-dependent convergence of the target and the time-dependent coronal plasma altering the




Figure 4.7: Demonstration of the analysis workflow to obtain the key results for this chapter. The power deposition at $t = 1.12$ ns from the CBET (left) and no-CBET (right) simulations are plotted in panel a) for the $R_b/R_t = 0.85$ case. Panel b) plots the radially integrated deposition from the profiles in a) as a function of azimuthal angle. As can be seen from this plot, the CBET asymmetry (light-blue) is greater than the no-CBET asymmetry (red). The power spectrum of these profiles is then plotted in panel c). This demonstrates that the dominant modes in the spectrum are multiples of the number of beams.

radius above the target where deposition is important. It is observed that in the presence of CBET, due to the non-uniform resonance of CBET gains across inbound laser sheets, additional asymmetries in the deposition are seeded and lead to more modal-flips than are observed without CBET.

4.3.1 Analysis and Quantity Definitions

Initially, definitions of key variables used in the analysis of the results of the chapter shall be provided. These are introduced for the example of the with-CBET and no-CBET $R_b/R_t = 0.85$ simulations, plotted instantaneously at $t = 1.12$ ns in Fig. 4.7.

Fig. 4.7.a shows the volumetric deposited power for the CBET (left) and no-CBET (right) simulations. Note that, as described in Sec. 4.2.3, the no-CBET simulation is forced to absorb the same magnitude of power as a function of time as the CBET simulation. Therefore, the total absorbed power is identical for both simulations, even though the no-CBET plot appears more saturated on the colour scale. This difference in saturation is partially due to the non-linear colour scale used for the plot, and also because the CBET result has more power deposited at larger radii, which widens the profile and reduces saturation on the colour scale. The mode-10 in the deposition due to the number of beams is clearly visible on both plots. Significant deposition in the caustic region, especially for the CBET result, is visible as the cross-structure in the deposited power. This suggests (and shall subsequently be shown explicitly in Sec. 4.3.3), that the caustic fields are strongly amplified by CBET, leading to more Inv-Brem in this region.

In order to quantify azimuthal asymmetry, radial integrals of the deposited power and fuel

density are taken,

$$\begin{aligned} P_{\text{dep}} R(\phi) &= \int_{r=0}^{\infty} P_{\text{dep}}(r, \phi) \, dr, \\ \rho_{\text{DT}} R(\phi) &= \int_{r=0}^{\infty} \rho_{\text{DT}}(r, \phi) \, dr, \end{aligned} \quad (4.8)$$

where P_{dep} is power, in units $[\text{W}/\text{m}^3]$, unlike the volumetric power plotted in Fig. 4.7.a. The deviation from the mean of these profiles can then be taken,

$$\text{Deviation}(f[\phi]) = \frac{f[\phi] - \int_{-\pi}^{\pi} f[\phi] \, d\phi}{\int_{-\pi}^{\pi} f[\phi] \, d\phi}. \quad (4.9)$$

The deviation of the CBET and no-CBET deposited power profiles shown in Fig. 4.7.a is plotted in Fig. 4.7.b. It can be seen that at this time, CBET considerably amplifies the instantaneous deposition asymmetry. It also distorts the sinusoidal profile of the no-CBET simulation, marginally widening and narrowing the curve peaks and troughs, respectively. Interestingly, CBET has also resulted in a phase-inversion of the deposition profile, where the peaks of the CBET deviation occur at the angles of the troughs of the no-CBET curve. This behaviour shall be called a *modal-flip* throughout this chapter. Note that modal-flips are also observed in the same simulation through time, *i.e.* as the target converges, the beam overlap pattern changes which results in phase inversions of the deposition, relative to earlier deposition profiles.

A Discrete Fourier Transform (DFT) is then used to analyse the modes which contribute to the signal. A signal $f(\phi)$, which is sampled N times in the interval $\phi \in [\phi_{\min}, \phi_{\max}]$ where $n = 0 \rightarrow N - 1$, has a DFT defined by,

$$F_{\ell} = \sum_{n=0}^{N-1} f_n \exp\left(-i2\pi \frac{\ell}{N} n\right), \quad (4.10)$$

where f_n is the sample at $\phi = (\phi_{\max} - \phi_{\min})n/N$ and ℓ is the frequency mode number. The power spectrum, which gives the power of each mode, is then given by,

$$P_{\ell} = \frac{1}{N^2} |F_{\ell}|^2. \quad (4.11)$$

The power spectra of the deposited power deviations from Fig. 4.7.b are plotted in Fig. 4.7.c on a log scale. The no-CBET profile is dominated by the $\ell = 10$ mode, with only a small $\ell = 20$ present. This yields the sinusoidal curve in Fig. 4.7.b. Many more modes are present for the CBET power spectrum, and (compared to the no-CBET results) a clear amplification of multiples of the $\ell = 10$ is visible. The significant $\ell = 20$ distorts the curve in Fig. 4.7.b, slightly widening the peaks and narrowing the troughs. Modes with $\ell < 10$ are presumed to be mostly spurious and introduced by relatively small, instantaneous errors in the field reconstruction algorithm. Unlike the $\ell = 10, 20, 30, \dots$, the $\ell < 10$ exhibit oscillatory, random growth from timestep to timestep. Therefore, they do not significantly imprint on the hydrodynamic profiles over the timescale of the implosion.




Figure 4.8: Radially integrated deposited power from no-CBET simulations as a function of time (x -axis) and angle (y -axis), alongside amplitudes of the dominant modes from a Fourier power spectrum. Panels a) and b) plot the radially integrated deposited power and Fourier modes, respectively, for the $R_b/R_t = 0.8$ simulation. The same is plotted for the $R_b/R_t = 0.9$ simulation in c) and d), and for the $R_b/R_t = 1.0$ simulation in e) and f). The mode 10 from the number of beams is clearly visible in the radially integrated power plots as 10 peaks to troughs in angle at a given time, *i.e.* 10 cyclical perturbations along a vertical lineout.

4.3.2 Deposition Asymmetries in the Absence of CBET

This section presents results of the deposited power in the absence of CBET for several implosions. Plotted in Figs. 4.8.a, 4.8.c and 4.8.e are $P_{\text{dep}}R(t, \phi)$ for $R_b/R_t = 0.8, 0.9$ and 1.0 , respectively. Explicitly, this is the radially integrated power from Eq. 4.8, plotted as a function of time (x -axis) and azimuthal angle (y -axis). The $\ell = 10$ deposition asymmetry at a single time (for example the red curve in Fig. 4.7.b) is visible as ten peaks to troughs along a vertical lineout. As expected, comparison of the colour scale saturation between these three plots demonstrates that at smaller beam radii, no-CBET asymmetries are much more significant. Figs. 4.8.b, 4.8.d and 4.8.f plot the power spectrum amplitude of the dominant $\ell = 10$ and 20 modes, as a function of time, on a log scale.

At times $t \sim 0.7, 0.9$ and 1.1 ns for the $R_b/R_t = 0.8, 0.9$ and 1.0 simulations respectively, a modal-flip of the deposited power is observed. This occurs because the plasma scale length




Figure 4.9: Demonstration of a mode-flip in P_{dep} for the no-CBET $R_b/R_t = 0.8$ simulation. Panel a), b) and c) plot the power deposition just before, during and just after the mode-flip. Panel d) plots the deviation from the mean of the radially integrated profiles around the azimuthal angle for all 3 times. It can be seen from this plot that the deposition is very symmetric at $t = 0.7$ ns. In all four panels, the angles of the beams are shown by the dashed magenta lines. The overlap region where deposition rises as the scale length increases, is highlighted with a dashed circle in panels a), b) and c).

increases in time, widening the plasma region above the critical surface where Inv-Brem is significant. Thus, the wings of the beams, which do not penetrate as far radially in, contribute more to deposition after this longer scale-length coronal plasma region has evolved. This eventually leads to a modal flip, when more deposition occurs between beam angles than at the angle of the beam itself. The flips occur later in time for wider beams, because the wings of the beams penetrate less far into the plasma, so a longer plasma scale length must develop before deposition from these edge rays becomes significant.

This is shown explicitly in Fig. 4.9, which plots the P_{dep} profiles for the $R_b/R_t = 0.8$, no-CBET simulation at $t = 0.6$ ns, $t = 0.7$ ns and $t = 0.8$ ns, i.e. just before, during and after the modal-flip, respectively. Particularly, Fig. 4.9.d plots the radially integrated powers plotted in Figs. 4.9.a, 4.9.b and 4.9.c. Before the modal-flip, at $t = 0.6$ ns, maximum deposition occurs at the angles of the beams, shown by dashed magenta lines. During the flip, at $t = 0.7$ ns, very symmetric deposition is observed and just after, at $t = 0.8$ ns, maximum deposition occurs between beam angles. Examining Figs. 4.9.a, 4.9.b and 4.9.c, the highlighted ‘cross’ feature (white dashed circle) between beam angles becomes increasingly saturated as more Inv-Brem occurs here. This occurs due to the plasma scale length increasing, raising the density further away from the critical surface and thus increasing deposition where the wings of neighbouring beams overlap.

4.3.3 CBET Imprint on Incident Field

When CBET is included in these simulations, it acts to significantly alter the field structure of the inbound laser sheets. This leads to additional asymmetry in the deposition, which predominantly occurs closer to the critical surface than the CBET scattering. Extra asymmetry is introduced, because the resonant CBET interaction is spatially localised, as opposed to




Figure 4.10: Field structure which leads to CBET induced asymmetry on power deposition and its dependence on R_b/R_t and target convergence. Each panel plots the incident sheet field, which includes the effect of CBET, along with contours of the critical electron density (green) and the incident field, $|E_z^{\text{in}}| = 1 \times 10^{10} \text{ Vm}^{-1}$ contour of another beam (magenta). Panel a) and b) plot this for the $R_b/R_t = 0.9$ simulation at $t = 1.0 \text{ ns}$ and $t = 1.5 \text{ ns}$ respectively. The convergence of the target in this time interval, leads to beams wrapping around the target more, and therefore a change in the spatial location of the resonant CBET interaction. Panel c) plots the same for the $R_b/R_t = 0.80$ simulation at the same time as panel a). This demonstrates that the $R_b/R_t = 0.80$ beam is not wide enough at this time to lead to a resonant CBET interaction, unlike the wider beam in panel a).

azimuthally symmetric around the target. Fig. 4.10.a illustrates this by plotting the electric field magnitude from the incident sheet of ‘beam-1’ from the $R_b/R_t = 0.9$ CBET simulation at $t = 1.0 \text{ ns}$, along with the $|E_z^{\text{in}}| = 1 \times 10^{10} \text{ Vm}^{-1}$ contour of ‘beam-4’ (magenta line). These two beams are separated from each other by 108° . It can be seen that there are two ‘holes’ in the beam-1 incident field profile, on either side of the beam centre. These are due to CBET interactions with other beams, for example, the resonance with the caustic field region of beam-4 is responsible for the hole at $[x, y] \sim [-500, -400] \mu\text{m}$. Typically, a large fraction of CBET scattering occurs in the caustic field region of beams for direct-drive implosions, due to the refractive swelling of the ray amplitude and therefore large electric fields present [167]. Caustic regions are narrow structures, typically with width $\sim \mathcal{O}(\lambda_0)$, therefore the region of maximal CBET scattering is strongly localised, leading to the non-uniform imprint on the

fields of inbound beams.

The trajectory that the light follows, and therefore the location of the caustic, depends on the electron density profile. For a direct-drive implosion, the critical density falls inward throughout the implosion, leading to a caustic field structure which ‘wraps around’ the target more, *i.e.* rays at an equivalent impact parameter are deflected less, as the critical radius converges. This can be seen by comparing Fig. 4.10.a ($t = 1.0$ ns) and Fig. 4.10.b ($t = 1.5$ ns), with critical radii $r_{\text{cr}} \sim 430$ and $360 \mu\text{m}$, respectively. The shrinking of r_{cr} has allowed beam-4 to wrap around the target more, such that the caustic CBET interaction now occurs in the middle of the incident field of beam-1. Comparing the field structures from Fig. 4.10.a and Fig. 4.10.b, the deposition of beam-1 at $t = 1.0$ ns will peak at the incident beam angle due to the depleted field on either side, whereas at $t = 1.5$ ns, the depletion of the field in the beam centre will shift the $P_{\text{dep}}R$ maxima in azimuthal angle. In other words a modal flip of the power deposition occurs between these times, which is due to the localised caustic CBET interaction changing location as the target converges.

Fig. 4.10.c plots the same as Fig. 4.10.a, but for the $R_b/R_t = 0.80$ case, *i.e.* a much narrower beam. From this plot, it can be seen that unlike for the $R_b/R_t = 0.9$ case, beam-4 does not wrap around the target enough for its caustic region to deplete beam-1, leading to a relatively unperturbed incident field profile. This demonstrates that the CBET induced modal flips are different for implosions with different R_b/R_t , because narrower (wider) beam will wrap around the target less (more) at an equivalent critical radius. Therefore, larger R_b/R_t values allow more modal flips to occur, because at a given critical radius, wider beams wrap around more, and therefore can undergo CBET with beams which have a larger angular separation. It is important to note that although this analysis only considers a single beam, the coronal hydrodynamic profiles are approximately azimuthally symmetric and therefore the beams and CBET interactions have rotational symmetry.

4.3.4 Deposition Asymmetries in the Presence of CBET

Results shall be presented in this section of the P_{dep} asymmetries for simulations including CBET, contrasted to Sec. 4.3.2, which presented results for the no-CBET simulations. Similarly to Fig. 4.8, plotted in Figs. 4.11.a, 4.11.c and 4.11.e are $P_{\text{dep}}R(t, \phi)$ for $R_b/R_t = 0.8, 0.9$ and 1.0 , respectively. Note that in these plots, the colour scale saturates at 20% rather than 5%, which demonstrates that instantaneous deposition asymmetries are typically much larger when the effects of CBET are included. A small discontinuity is visible in these plots at $t \sim 2.15$ ns, which is the time that the critical density passes from the CH ablator material into the DT fuel. Mix between these materials is not modelled in CHIMERA, so this critical density traversal leads to short-lived, potentially spurious, radially outward ∇n_e near the critical radius, which are visible in the field reconstruction and deposition.

The $P_{\text{dep}}R$ plots in Fig. 4.11 have significantly more structure than the equivalent no-CBET plots in Fig. 4.8. There is an increase in the number of mode-flips, increasing R_b/R_t appears to increase deposition asymmetry (rather than decrease for no-CBET) and more significant higher modes also appear to be present in $P_{\text{dep}}R$. Figs. 4.11.b, 4.11.d and 4.11.f plot the $P_{\text{dep}}R$




Figure 4.11: This figure plots the same as Fig. 4.8, but now for the equivalent simulations which include the effect of CBET. Comparing these results and those in Fig. 4.8 demonstrates that CBET introduces additional modal-flips of the deposition and amplifies the magnitude of asymmetries.

Fourier power spectrum amplitudes of the modes $\ell = 10$ and 20 throughout the implosion. The mode-flips can be identified in these plots as the times at which the $\ell = 10$ sharply falls and then quickly rises again.

Comparing Fig. 4.8 and Fig. 4.11, CBET leads to a mode-flip in the deposition at $t \sim 0.4$ ns, which does not occur for the no-CBET case. The resulting $\ell = 10$ in the CBET deposition profiles between $t \sim 0.4$ ns and $t \sim 1.2$ ns strongly depends on the R_b/R_t value. For the $R_b/R_t = 0.8$, CBET simulation, the $\ell = 10$ amplitude is low compared to the larger R_b/R_t , CBET simulations. This is because the relatively narrow beams at this time do not sufficiently wrap around the target to give rise to strong caustic fields and subsequent CBET imprints on incident beams. This is seen explicitly, by comparing the incident field profiles at $t = 1.0$ ns for the $R_b/R_t = 0.9$ and 0.8 simulations plotted in Figs. 4.10.a and 4.10.c respectively.

Interesting behaviour is also observed later in the implosion. Comparing the $R_b/R_t = 0.9$ (Fig. 4.11.c) and 1.0 (Fig. 4.11.e) plots at $t \sim 1.8$ ns, it can be seen that an additional mode-flip is beginning to develop for the $R_b/R_t = 1.0$ case. This is also seen in Figs. 4.11.d and 4.11.f, by the reduction in amplitude of the $\ell = 10$ and rise of the $\ell = 20$. However, the $R_b/R_t = 1.0$ beams are not quite wide enough to fully translate to an inversion of the deposition asym-




Figure 4.12: Demonstration of a mode-flip in P_{dep} for the CBET, $R_b/R_t = 0.9$ simulation. Panels a), b) and c) plot the power deposition before, during and after the mode-flip. Panel d) plots the deviation from the mean of the radially integrated profiles around the azimuthal angle for all 3 simulations. Fig. 4.10.a and Fig. 4.10.b plot a single, incident sheet field from panels a) and c) respectively. It can be seen that the centre of the beam depletion at $t = 1.5 \text{ ns}$ leads to less deposition at the beam angles (magenta lines) on panel d). Panel d) shows that during the CBET induced mode-flip, a significant $\ell = 20$ occurs in the deposition. The reducing absorption at the centre of one beam is highlighted in by the dashed circles.

metry, so this results in a net decrease in $P_{\text{dep}}R$ asymmetry, compared to the $R_b/R_t = 0.9$ simulation at the same time.

Fig. 4.12 plots the development of the mode-flip which occurs at $t \sim 1.25 \text{ ns}$. Figs. 4.12.a, 4.12.b and 4.12.c plot the deposited power from all beams before, during and after the flip, respectively. As can be seen in the area highlighted by the dashed circles in each plot, the deposition at the angle of the beams (magenta dashed lines) becomes less significant, compared to the angles between beams. Figs. 4.12.a and 4.12.c are from the same time as the single field profiles plotted in Figs. 4.10.a and 4.10.b, respectively. These field profiles demonstrate that CBET scattering in the caustic region of the beams separated by 108° , is responsible for the depletion of the inbound beams on either side of the beam centre at $t = 1.0 \text{ ns}$ and in the beam centre at $t = 1.5 \text{ ns}$. The radially integrated deposition is plotted in Fig. 4.12.d, which shows the $\ell = 20$ of the deposited power during the mode-flip.

4.4 Stagnation State Asymmetry

This section describes how the asymmetry of the deposition profiles, described in Sec. 4.3, imprint upon the hydrodynamics. The stagnation profiles for various R_b/R_t are presented for CBET and no-CBET simulations. These demonstrate that CBET tends to increase the asymmetry of the stagnated state, due to the decreased absorption uniformity. It is observed that CBET changes the trend of the stagnation symmetry with R_b/R_t . While, for the no-CBET simulations, symmetry increases with increasing R_b/R_t , CBET generally reverses this behaviour, with the least symmetric stagnation states observed for the largest R_b/R_t value. The CBET trend is, however, non-monotonic which is due to development of addi-

tional mode-flips at specific values of R_b/R_t , slightly increasing the deposition uniformity late in the implosion.

4.4.1 Stagnation State Asymmetry Trend with Beam Radius

Fig. 4.13 plots the stagnation fuel density and ion temperature for both CBET and no-CBET simulations at four different R_b/R_t . The beam-mode $\ell = 10$ is clearly identifiable in all plots. Recall that all simulations are tuned to have the same amount of absorbed laser energy, via the tuning process described in Sec. 4.2.3. Therefore, all profiles have approximately the same convergence properties. If 1-D effects of CBET were included, which would involve a much more complex, simultaneous optimisation of pulse and target, the stagnation state of all simulations would likely look far less similar. No-CBET and smaller R_b/R_t results would have significantly more coupled energy, and would therefore have more mass in the final fuel configuration. Instead, these results effectively show, for hydrodynamically similar implosions, how CBET and increasing R_b/R_t affects deposition and stagnation *symmetry* only. This simpler optimisation also has the advantage that the final fuel configurations do appear more similar and are therefore directly comparable.

The no-CBET results clearly show a trend of increasing stagnation symmetry with increasing R_b/R_t , as expected from the plots in Fig. 4.8, which showed increasing absorption uniformity with increasing beam radius. The pressure of the stagnation state is approximately isobaric, and therefore the temperature profiles are inversely proportional to the density. Peak densities increase and the radii of stagnation decrease marginally as R_b/R_t increases. This improved compression is partially due to increased symmetry providing better compression, but more optimal shock timing for the larger R_b/R_t implosions also contributes. As is seen in the streak plots in Fig. 4.6, the tuning process led to the first shock hitting the axis slightly earlier at narrower R_b/R_t , than for the larger R_b/R_t implosions. Although the streak plots are for 1-D with-CBET simulations, the same trend is observed for the 2-D with- and without-CBET simulations. The more optimal shock timing for larger R_b/R_t also has the signature of higher on-axis densities. The contribution of small differences in shock timing is not quantified in the following analysis, but it is assumed that it will have a second-order effect on the stagnation uniformity, compared to asymmetry in deposition.

Increasing R_b/R_t when including CBET broadly shows the opposite behaviour to the no-CBET simulations, leading to highly non-uniform density profiles at large R_b/R_t . This trend was also observed in the deposition plots in Fig. 4.11, with the wider beam simulations leading to a more saturated plot on the colour scale. Higher order modes than $\ell = 10$ also become increasingly evident in the wider beam CBET simulations. Note that the $R_b/R_t = 0.8$ simulation with CBET is more symmetric than the no-CBET simulation at the same beam radius. This is believed to be due to the developing mode-flip, visible in Fig. 4.11.a from $t \sim 0.5 \rightarrow 1.3$ ns, which strongly reduces the absorption asymmetry and subsequent imprinting onto the density.

The standard deviations in azimuthal angle of the radially integrated stagnation state fuel




Figure 4.13: Densities of the DT fuel and ion temperatures for various R_b/R_t simulations both with and without CBET. Each row corresponds to a different R_b/R_t value; the left column contains simulations without CBET; and the right column contains simulations with CBET. It is visible from the density plots that increasing R_b/R_t improves stagnation symmetry for the no-CBET simulations, but degrades it for the CBET simulations.

density,

$$\sigma[\rho_{DT}R] = \sqrt{\langle \rho_{DT}R^2 \rangle - \langle \rho_{DT}R \rangle^2}, \quad (4.12)$$

for the CBET and no-CBET simulations are plotted as a function of R_b/R_t in Fig. 4.14.a. The




Figure 4.14: Trends of a) stagnation asymmetry and b) maximum (azimuthally averaged) fuel density for CBET and no-CBET simulations. The no-CBET improvement in symmetry with R_b/R_t is observed which also corresponds to improved compression. The symmetry trend including CBET is more complex, but broadly the stagnation state symmetry is worse with increasing R_b/R_t . Note that simulations had a similar stagnation areal density, $\rho_{DT} R \sim 2.5 \text{ g cm}^{-2}$. This nominal value is provided to help the reader to appreciate the magnitude of asymmetry plotted in panel a).

no-CBET trend from this metric clearly follow the behaviour seen in Fig. 4.13, that increasing beam width results in better symmetry of the hydrodynamic profiles at stagnation. Note that the $R_b/R_t = 0.75$ simulation value is more symmetric than the 0.8 result, suggesting that more complex behaviour may occur at smaller beam radii, $R_b/R_t < 0.8$. However, this is outside the range of typical OMEGA implosions and therefore not studied in detail. The with-CBET curve shows a broad trend of increasing asymmetry with wider beams, which is visible in the stagnation profiles in Fig. 4.13, however the behaviour is clearly more complex than the no-CBET simulations. A local maximum of asymmetry is observed at $R_b/R_t = 0.9$ and minima occur at $R_b/R_t = 0.8$ and $R_b/R_t = 1.0$. As shall be described in the subsequent section, Sec. 4.4.2, these features are explained by the complex, CBET-induced modal-flips of the deposition which were described in Sec. 4.3.4.

The azimuthally averaged maximum density,

$$\rho_{DT,max} = \max \left(\int_{-\pi}^{\pi} \rho_{DT}(r, \phi) d\phi \right), \quad (4.13)$$

is plotted in Fig. 4.14.b for both sets of simulations. The behaviour in this plot is inversely proportional to the standard deviation, which suggests that the asymmetry of the density has a detrimental effect on the compression of the target. Lower modes can prove detrimental to compression by introducing significant non-radial motion, resulting in less efficient compression [227, 228]. The (unintended) improved shock timing at larger beam radii does appear to somewhat compensate for the reduction in uniformity, *i.e.* the compression metric at $R_b/R_t = 1.05$ is about the same as $R_b/R_t = 0.8$, despite far worse symmetry. It is therefore difficult to quantify the contribution to the degradation of compression due to each of these




Figure 4.15: Time resolved $\ell = 10$ Fourier power spectrum amplitude for a) $\rho_{DT} R$ and b) $P_{dep} R$ for CBET simulations with 3 R_b/R_t values. The developing, but unrealised, modal-flip for $R_b/R_t = 0.8$ from $t \sim 0.5 \rightarrow 1.2$ ns reduces the $P_{dep} R_{\ell=10}$, leading to slow $\rho_{DT} R_{\ell=10}$ growth and ultimately a relatively symmetric stagnation state. Despite large values of $\rho_{DT} R_{\ell=10}$ initially, the developing modal flip of the $R_b/R_t = 1.0$ simulation from $t \sim 1.8 \rightarrow 2.1$ ns slows the density asymmetry growth.

effects, from this set of simulations.

4.4.2 Time Resolved Asymmetry Growth

The time-resolved, $\ell = 10$ mode growth of the fuel density from the $R_b/R_t = 0.8, 0.9$ and 1.0 , with-CBET simulations are examined more closely in this section in order to understand the non-monotonic behaviour of the CBET, fuel-symmetry curve from Fig. 4.14.a. These simulations were chosen as they are the local minima and maxima of the curve and therefore should most clearly display the important behaviour. Fig. 4.15.a plots the $\ell = 10$ growth of $\rho_{DT} R$ for each of these simulations. Note that all bangtimes were approximately $t = 2.5$ ns, so the value at this time is the $\ell = 10$ at stagnation. The $\ell = 10$ growth of $P_{dep} R$ for these simulations is plotted in Fig. 4.15.b, which are the same as the curves from the right panels of Fig. 4.11, but are plotted again here for direct comparison with the $\rho_{DT} R$ values. The mode-flips are clearly identifiable in this plot as sharp falls and rises in the $\ell = 10$, which corresponds to an inversion of deposition asymmetry.

Firstly, the growth of the fuel density asymmetry from the $R_b/R_t = 0.8$ simulation shall be studied. It can be seen from the green curve in Fig. 4.15.a, that the growth of this mode

is slow from $t \sim 0.5 \rightarrow 1.4$ ns, compared to the $R_b/R_t = 0.9$ and $R_b/R_t = 1.0$ simulations. This is explained due to the low deposition asymmetry in Fig. 4.15.b which is ultimately responsible for imprinting the asymmetry on the hydrodynamics. As was seen from the time-resolved, radially integrated deposition in Fig. 4.11.a, this low asymmetry of the deposition is due to the developing mode-flip. In this simulation, the developing mode-flip at this time ultimately reduces deposition asymmetry, because the beam is not wide enough for the target convergence to truly flip the deposition. Although the $\ell = 10$ of the $\rho_{DT}R$ starts to grow again after the second mode-flip at $t \sim 1.2$ ns, the long period of highly symmetric deposition means that the stagnation state is ultimately relatively symmetric, compared to higher and lower beam radii simulations.

Comparison of the growth between the $R_b/R_t = 0.9$ and 1.0 simulations is slightly more nuanced. As is seen in the radially integrated power deposition from Figs. 4.11.c and 4.11.e, the mode-flip at $t \sim 0.5$ ns develops into a full inversion of the deposition asymmetry. Because the beams are wider in the $R_b/R_t = 1.0$ simulation, more light wraps around the target. This leads to stronger caustic fields, more CBET and therefore greater distortion of the inbound field profiles. Thus, the $R_b/R_t = 1.0$ has a less symmetric deposition profile than the narrower beam case, until the next mode-flip occurs at $t \sim 1.2$ ns. This explains why the $\ell = 10$ of the $\rho_{DT}R$ is larger for the $R_b/R_t = 1.0$ simulation up until $t \sim 1.5$ ns.

This remains the case until the wider beam of the $R_b/R_t = 1.0$ case begins to wrap around the target sufficiently to begin the development of an additional mode-flip of the deposition at $t \sim 1.6$ ns. Similar to the $R_b/R_t = 0.8$ case from $\sim 0.5 \rightarrow 1.2$ ns, the development of a flip in the deposition profile acts simply to reduce the deposition asymmetry, since the beam is not sufficiently wide. Therefore, the deposition asymmetry of $R_b/R_t = 1.0$ from $\sim 1.6 \rightarrow 2.1$ ns is reduced compared to the $R_b/R_t = 0.9$ simulation, for which the beams are not wide enough to develop another mode-flip. The reduction of deposition asymmetry is again clearly visible by comparing Figs. 4.11.c ($R_b/R_t = 0.9$) and 4.11.e ($R_b/R_t = 1.0$) from $\sim 1.6 \rightarrow 2.1$ ns. More symmetric deposition during this period reduces the rate of $\rho_{DT}R$ asymmetry growth. This results in a less symmetric $R_b/R_t = 0.9$ stagnation state, as this simulation does experience the same disruption to the deposition asymmetry source. For $R_b/R_t > 1.0$, the beams are wide enough that the mode-flip at $t \sim 1.6$ ns does lead to a full inversion of the deposition asymmetry and therefore the $\ell = 10$ of the $P_{dep}R$ is not reduced in the same manner for even wider beams. Ultimately, this leads to increasing asymmetry of the stagnation state at $R_b/R_t > 1.0$, as is observed in Fig. 4.14.a.

While this analysis does explain the trend observed in Fig. 4.14.a, subtleties exist which are not explicitly addressed, but could offer additional insight to the growth of stagnation state asymmetries in direct-drive implosions. For example, the deposition all occurs below the critical density, which is separated from the ablation surface by a finite stand-off distance. Thus, the energy must be transported from the absorption region to the ablation surface by thermal conduction. This leads to a time-lag for asymmetry in Fig. 4.15.b (deposition) to propagate through to the curves in Fig. 4.15.a (fuel density). Changing coronal plasma conditions during the implosion will also change this time-lag in a non-obvious way. For example, early in the implosions, deposition occurs exclusively in the outermost CH mate-

rial. As time progresses, the CH expands away and the ablation surface changes from CH to the DT layer. Even later, (at $t \sim 2.1$ ns in these simulations), the critical surface moves from CH into DT. All of these changes in coronal and transport layer conditions will impact how conduction propagates asymmetry from the deposition to the fuel, *i.e.* there is additional complexity in mapping from $P_{\text{dep}}R$ asymmetry to $\rho_{\text{DT}}R$ asymmetry. Future work could perhaps investigate this further, by also comparing the heat flux just above the ablation surface, which may further elucidate the important physics.

4.4.3 Comparison of Results to Statistical Modelling

Ultimately, this work was conducted to explain why statistical modelling of OMEGA implosions demonstrates that distinct regimes exist in the R_b/R_t degradation of implosions from 1-D calculations, *i.e.* to explain the shape of the curve in Fig. 4.3. The fit to the experimental scatter plot has a transition at $R_{b/t}^{\text{crit}} = 0.86$, after which, improvement of the yield for increasing R_b/R_t slows. Neither 1-D codes which include CBET models, or 3-D codes without CBET, are able to reproduce this trend, which means that multidimensional effects, seeded by CBET, are considered a likely explanation [210]. The work conducted in this chapter is not directly comparable to these plots because a 1-D tuning process was conducted to ensure that all implosions had the same coupled energy. This was to prevent having to perform a far more in depth optimisation procedure for both the pulse and target simultaneously. Therefore, changes in the 1-D effects of CBET at different R_b/R_t are not included in these simulations (changes to magnitude of deposited energy), only the redistribution of deposited energy due to CBET, *i.e.* how CBET alters asymmetry. The differences in cylindrical and spherical geometry also complicate direct comparison between results due to different coronal plasma profiles and convergence properties. For example, the R_b/R_t values of transitions in behaviour would not be expected to be the same between geometries.

Despite this, the asymmetry and compression metrics in Figs. 4.14.a and 4.14.b, respectively, clearly demonstrate relatively sharp transitions in behaviour. For example, asymmetry and compression both get worse from $R_b/R_t = 0.8 \rightarrow 0.9$, but improve from $0.9 \rightarrow 1.0$. If this behaviour was also observed in spherical geometry simulations, which is due to CBET introducing mode-flips in the deposited power, it could potentially be the cause of the experimental data transition at $R_{b/t}^{\text{crit}} = 0.86$. Explicitly, if CBET acted to degrade deposition symmetry in spherical implosions from $R_b/R_t = 0.86 \rightarrow 1.00$, this would slow the rate of improvement in the yield over clean metric, plotted in Fig. 4.3. While this is not at all a conclusive explanation of the behaviour in the statistical model, results from these simulations do not seem to rule it out. It therefore offers motivation for extending the simulation platform to a spherical geometry and including 1-D effects of CBET in R_b/R_t trends by changing the target and fixing the total incident laser energy. This should conclusively demonstrate that shape degradation due to CBET either is or is not responsible for the transitions.

4.5 Conclusions

4.5.1 Summary of work

This chapter has described the development of a novel, 2-D cylindrical simulation platform to study direct-drive implosion physics, and its use to investigate how altering R_b/R_t changes the stagnation symmetry both with, and without the effects of CBET. The work was conducted to test the hypothesis that changes to stagnation state symmetry due to changing CBET action at different R_b/R_t were responsible for unexplained behaviour in statistical modelling of OMEGA experiments. In order to avoid performing an in-depth optimisation process, where the target and pulse were simultaneously varied, only the incident laser energy was varied, to fix the coupled energy to the target. Therefore, only the changes to deposition symmetry due to CBET were studied at different R_b/R_t . This 1-D tuning procedure yielded eight implosions both with and without CBET effects.

By conducting 2-D simulations with the initial conditions obtained from the tuning, interesting behaviour was observed for both the CBET and no-CBET cases. In both cases, mode-flips of the deposition were observed, due to different effects. For no-CBET, these mode flips occurred due to increasing plasma scale lengths, which shifted maximum deposition from beam angles to the overlap region between beams. When including CBET, this scale-length behaviour was obscured by much more dominant mode-flips, which were induced due to caustic CBET interactions distorting the incident field profiles of beams.

Without CBET changes to deposition uniformity, the stagnation state asymmetry decreased with increasing R_b/R_t , as more beam overlap led to increasingly uniform target illumination. When CBET was included however, increasing CBET with wider beam radii resulted in less and less symmetric deposition, which broadly led to increasing asymmetry with R_b/R_t . The precise behaviour was, however, non-monotonic due to CBET mode-flips increasing the symmetry of illumination at specific R_b/R_t values. This led to distinct transitions regions in the trends of symmetry and compression, plotted against R_b/R_t , which were identified as a potential explanation of the transitions observed in OMEGA statistical modelling. Further work must however be conducted to obtain direct comparison to the statistical model and thus conclusively state whether CBET asymmetries are, or are not, the cause.

4.5.2 Future Work

There are many ways that the work presented in this chapter could be extended, through *e.g.* improving the simulation platform to better reflect reality, or by studying additional physics on the existing platform. One interesting way to extend the analysis of the work would be to investigate how heat transport affects the conversion of deposition asymmetries into hydrodynamic perturbations. As was discussed at the end of Sec. 4.4.2, there is a lag in the conversion due to thermal conduction taking this information inward via a finite transport region. The transport region would also vary throughout the implosion as the hydrodynamic profiles evolve and the material changes from CH to DT. By including analysis of heat flux, the way that this information is propagated could be studied. Furthermore, a model for

non-local heatflow could be incorporated into CHIMERA, such as SNB [110, 111, 112] or FAST-VFP [114]. Due to the strong temperature gradients in the transport layer, non-local effects could significantly affect how this information is propagated to the ablation surface and thus imprints on the hydrodynamics.

The geometry of the platform could also be modified to allow for spherical rather than cylindrical convergence. This would allow for better comparison between simulation results and spherical implosion experiments. Cylindrical geometry leads to greater target convergence for a given incident energy and also less divergence of the outflowing, coronal plasma, thus yielding higher coronal plasma densities for a given ablation velocity. By simulating a thin, equatorial ‘slice’ of a spherical domain, *i.e.* a spherical simulation with bounds $\phi \in [-\pi, \pi]$ and $\theta \in [\lesssim \pi/2, \gtrsim \pi/2]$, a 2-D ray-trace could still be conducted so simulation run-times should be very similar to the cylindrical case. This was not done for the work in this chapter because the SOLAS 2-D ray-trace assumes that the simulation occurs in a plane, and the spherical ‘slice’ is not a plane. Minor development for this specific case could rectify this however and enable similar simulations in spherical geometry.

Similar initial conditions could also be used to study further multidimensional direct-drive physics, without the need to conduct, an expensive, full 3-D ray-trace. For example, laser imprint is a significant issue in direct-drive and means that implosions must be conducted at a high adiabat for optimal yields [229]. This limits the compressibility and therefore maximum performance of Inertial Confinement Fusion (ICF) experiments [230, 231]. Improved understanding of laser imprint requires high resolution simulations, due to the small spatial scale of the intensity modulations. Typically, high resolution, 2-D, $r - z$ simulations with a 3-D laser ray-trace [232], or lower resolution, fully 3-D simulations [233] are conducted to study the effect of imprint. 2-D $r - z$ Rad-Hydro simulations with a 3-D laser ray-trace average the deposition from the 3-D ray-trace around the azimuthal angle, and therefore artificially smooth out the imprint modulation to power deposition. While 3-D simulations at a sufficiently high resolution should resolve the imprint seed properly, these are extremely expensive to perform, particularly for the fixed angular resolution, Eulerian code ASTER which is typically used for these problems. By performing a spherical or cylindrical ‘slice’ simulation with an imprint model, the seed and the growth of imprint induced instabilities could be consistently modelled at high resolutions, providing a complimentary study to the existing literature. This setup would not include out-of-plane stabilisation flows however, and could therefore overestimate growth of these small modes.

5 Simulations of Cross-Beam Energy Transfer for Magnetised Direct-Drive

This chapter describes a set of simulations which were conducted to understand the role of Cross-Beam Energy Transfer (CBET) in magnetised, direct-drive implosions. Magnetised Inertial Confinement Fusion (ICF) is a promising route to achieving higher target gains, due to the reduction of hotspot thermal energy loss at stagnation and additional confinement of the DT fusion alpha products. For direct-drive implosions, magnetisation can significantly alter the coronal plasma conditions, due to the anisotropisation of thermal transport. The Ion Acoustic Wave (IAW) dispersion relation, which mediates CBET interactions, depends upon the background plasma and, therefore significantly altered temperature and density profiles could alter the action of CBET. No direct-drive suitable CBET model had been integrated into a Radiative-Magnetohydrodynamics (Rad-MHD) code, prior to the development of SOLAS. Therefore, the CHIMERA-SOLAS framework has enabled the effect of magnetisation on CBET to be studied for a direct-drive implosion.

The chapter begins with a review of experimental and computational work on magnetised ICF, with a particular focus on magnetised direct-drive. The work presented in this chapter focuses on the study of ‘exploding-pusher’ experiments. These are very different implosions to the typical central hotspot ignition designs, presented in previous chapters, so a short summary of exploding-pushers is also provided. 1-D and 2-D simulation results of unmagnetised exploding-pushers, both with and without the effect of CBET, are presented and demonstrate that CBET does significantly affect these implosions. This is followed by an investigation of how various extended-Magnetohydrodynamics (MHD) terms alter the implosion, including the Nernst effect, the Lorentz force and resistive diffusion of the magnetic field. Results are given of how magnetisation affects the CBET interaction and ultimately how it changes the stagnation shape of the target. The results demonstrate that redistribution of deposited power due to CBET reduced the amplitude of the stagnation asymmetry, which originated from the polar beam configuration used. However, the reduction of asymmetry was consistent for different initial seed magnetic field values, and therefore CBET was not observed to be sufficiently strongly affected by magnetisation to lead to observable signatures in experimental measurements. The chapter concludes with a summary of the work conducted, and suggestions of additional experimental configurations, which may leave a more significant signature of magnetisation altering CBET.




Figure 5.2: Simulated fusion yields versus effective residual fuel kinetic energy under imposed low-mode radiation flux perturbations for imposed fields in the range $B_0 = 0 \rightarrow 70$ T. The plot demonstrates, that with increasing departure from ideal compression (moving to the right on the x axis), magnetisation can enable the onset of the ignition. Reproduced with permission from Ref. [237].

and $\rho \sim 50 \text{ g cm}^{-3}$, a field strength $|\mathbf{B}| \sim 2.5 \text{ kT}$ is required to obtain $\omega\tau \sim 1$ [94]. This field strength cannot be produced directly, but it is possible to produce a smaller field which, assuming frozen in magnetic field and a spherical compression, is amplified by the square of the convergence,

$$|\mathbf{B}_1| = |\mathbf{B}_0| \left(\frac{R_0}{R_1} \right)^2, \quad (5.1)$$

where $|\mathbf{B}_0|$ and $|\mathbf{B}_1|$ are initial and final magnetic fields, respectively and R_0 and R_1 are initial and final radii, respectively. Laboratory magnetic fields can be produced from pulsed power coils with field strength $|\mathbf{B}| \sim \mathcal{O}(50) \text{ T}$ [236], so even moderate convergence-ratio targets ($R_0/R_1 \sim 10$) are able to produce strongly magnetised core plasma.

Fig. 5.2 plots results of magnetised indirect-drive simulations, for a target on the threshold of ignition [237]. Increasing magnitudes of drive perturbation were applied to the drive (moving to the right on the x -axis). At some threshold level of perturbation the target is prevented from achieving ignition, which is visible as the steep decrease in yield at a specific level of perturbation for each curve. The results demonstrate that when an initial magnetic field was applied to the target it more robustly ignited with increasing field strength, due to reduced conduction losses. This simulation work, prior to the achievement of ignition on the NIF [53], motivated the development of a magnetised ICF campaign at Lawrence Livermore National Laboratory (LLNL) [238].

The CHIMERA code has been used to study a wide array of physics relevant to magnetised ICF. Simulation work has shown that magnetisation can alter instability growth of magnetised laser fusion implosions. While in the deceleration phase, magnetic tension can reduce low-mode perturbation growth [239], and magnetisation of directly-driven targets inhibits heatflow in the plasma corona, thus limiting thermal stabilisation of short wavelength




Figure 5.3: a) A 1.1 keV T_i increase was achieved by adding a 26 T B_0 field to a D_2 gas capsule implosion on the NIF. Also shown in the plot are the simulation results. b)→d) Equatorial shapes of the implosions. Reproduced with permission from Ref. [241].

modes from laser imprint [240]. Recent work has also demonstrated that magnetisation of high-yield, indirect-drive targets must be carefully optimised, in order to avoid significant degradation to the implosion shape, due to anisotropic thermal conduction and inhibition of burn propagation from α magnetisation [94].

5.1.2 Experimental Studies of Magnetised-ICF

Indirect-drive experiments have been conducted on the NIF to demonstrate the efficacy of magnetised targets in reducing thermal conduction losses in the hotspot. Non-cryogenic, deuterium-filled capsules were deployed with initial field strengths up to 26 T [241]. Results from this experimental campaign are shown in Fig. 5.3.a. The magnetised targets demonstrated significantly enhanced ion temperatures and neutron yields and work is underway to explore non-uniform field configurations to further enhance the benefits of magnetisation [242]. Fig. 5.3.b, 5.3.c and 5.3.d plot x-ray images at stagnation of different experiments, showing that a shape-tuning process had to be conducted in order to optimise the sphericity of the target, partially due to the field leading to anisotropic thermal conduction.

Magnetisation of direct-drive targets has been investigated by experiments on the OMEGA laser facility for several years. Initial OMEGA experiments focussed on verification of magnetic flux compression, by applying an initial seed field along the axis of a cylinder that was imploded via laser irradiation [244]. The magnetised implosions validated predictions of flux compression and demonstrated enhanced neutron yields and core ion temperatures over unmagnetised implosions. Spherical targets were subsequently fielded, which also resulted in increased stagnation temperatures and yield compared to unmagnetised targets [245,




Figure 5.4: X-ray self-emission images of (left) an unmagnetised and (right) a magnetised implosion. The average radius of the marked contour (corresponding to 40% of peak intensity), and the oblateness parameter R_a/R_b (ratio of major-to-minor axis) are listed below each image. The polar laser-drive is indicated by the white arrows, and the axis of the initial magnetic field by the black dashed line on the right. Applying an initial magnetic field demonstrated increased oblateness of the implosion. Adapted with permission from Ref. [243].

246]. No noticeable degradation to the implosion shape or performance was observed in these experiments, which was assumed to be due to the high ratio of plasma pressure to magnetic pressure, $\beta \gg 1$.

The most recent experimental, magnetised direct-drive work has focussed on exploring higher initial seed field values ($|\mathbf{B}_0| \sim 50 \text{ T}$ compared to $|\mathbf{B}_0| \sim 8 \text{ T}$), to understand the saturation of performance with increasing field. A shock-driven, exploding-pusher target configuration was used for these experiments, in order to create high ion temperatures and thus create a platform to study magnetised ions. Exploding-pushers are significantly different implosions compared to hotspot ignition targets discussed in previous chapters and shall be described in detail in Sec. 5.1.4. Creating these strong, 50 T fields at the target necessitated reducing the radius of the equatorial field coil compared to previous experiments. Therefore, a 40-beam configuration had to be used, without the 20 equatorial beams, leading to a pole-heavy laser drive. The high fields of these implosions led to strongly magnetised coronal electrons, $\omega_e \tau_e \sim 50$, resulting in strongly anisotropic thermal conduction $\kappa_{\perp,e}/\kappa_{\parallel,e} \sim 10^{-4}$. This is compared to previous experiments which produced $\omega_e \tau_e \sim 1$ and therefore $\kappa_{\perp,e}/\kappa_{\parallel,e} \sim 1/3$. For direct-drive on OMEGA, laser deposition is transported to the ablation surface by electron thermal conduction, thus large electron Hall parameters led to an effective anisotropisation of the implosion drive.

Fig. 5.4 shows x-ray self-emission images of an unmagnetised (left) and magnetised (right) target with an initial $|\mathbf{B}_0| = 50 \text{ T}$ seed field. The strongly magnetised coronal electrons led to decreased drive $\perp \mathbf{B}_0$, markedly increasing the oblateness of the diagnostic image compared to the unmagnetised target. An ion magnetisation of $\omega_i \tau_i \sim 7$ was also reported. Previous Rad-MHD modelling of these experiments, using the CHIMERA code, did not include the effects of CBET. The development of SOLAS, and particularly the CBET model, motivated further computational study of these experiments, to explore whether CBET played a significant

role in dictating the shape of these implosions. This is because CBET is known to markedly compensate global, $\ell = 1$ asymmetries [247, 67], therefore the anisotropy introduced from magnetisation could affect the action of CBET.

5.1.3 Magnetised Laser-Plasma Instabilities

Sec. 5.4 aims to understand how magnetisation of a direct-drive implosion anisotropically changes the hydrodynamics, and how these altered coronal plasma conditions modify the calculated CBET gains, discussed in Sec. 3.5.1. For example, magnetisation restricts thermal conduction and therefore enhances coronal electron temperatures along the initial field axis. Approximately, the fluid CBET gain, $\gamma_{ij} \propto T_e^{-1}$, therefore anisotropic changes to T_e could result in reduced CBET gains around the target and therefore change CBET scattering compared to implosions without an applied field. This modification to CBET via the altered hydrodynamic profiles is called the *indirect* effect of magnetisation on LPIs.

Magnetisation can, however, also *directly* modify scattering from LPIs, in a number of ways. For ICF conditions, when the field strength is sufficiently high, electron cyclotron motion can become comparable to plasma wave frequencies, and therefore alter the dispersion relation of the mediating plasma wave in LPIs. In underdense, ICF-relevant plasma ($n_e \sim 10^{20} \text{ cm}^{-3}$ and $T \sim 2 \text{ keV}$), the IAW, which mediates Stimulated Brillouin Scattering (SBS) and CBET, is significantly modified when $|\mathbf{B}| \sim 100 \text{ T}$ and the Electron Plasma Wave (EPW), which mediates Stimulated Raman Scattering (SRS) and Two Plasmon Decay (TPD), when $|\mathbf{B}| \sim 1000 \text{ T}$ [248]. Additionally, the (predominantly collisionless) damping of plasma waves can also be modified, because cyclotron motion of particles can affect their trapping in plasma waves [248]. Significant theoretical progress has been made in this field in recent years by Shi et al., who derived analytic formulae for three-wave coupling in the presence of a magnetic field [249, 250]. This was a challenging problem, due to the lack of simple geometries for the interaction, when a field is applied to a plasma with an arbitrary direction.

The simulation results here neglect this direct affect of magnetisation on CBET, partially because the theory is not yet deemed to be significantly mature to implement within a reduced, ray-based model. Additionally, coronal magnetic field strengths of $|\mathbf{B}| \lesssim 50 \text{ T}$ were observed in the underdense coronal plasma so significant modifications to the IAW dispersion relation were not expected. It is noted, however, that altered damping of the waves from magnetisation may affect the results, but the focus of the study presented here was predominantly to explore how magnetisation might indirectly affect CBET.

5.1.4 The Exploding-Pusher Configuration

Exploding-pushers are considered to be a highly reproducible platform, robust to instabilities and capable of producing large neutron yields. Although historically it had a slightly different meaning [127], the term ‘exploding-pusher’ is now typically used for low convergence, thin-shell targets [251]. When irradiated with significant intensity, frequency-tripled laser

light¹, the thin shell rapidly heats and then explosively ablates, driving a strong shock radially inward, ahead of the majority of in-falling ablated material. This shock strongly heats the ions as it propagates through gas-fill to large, fusion-relevant temperatures. After rebounding from the axis, the shock recompresses the infalling exploded shell material, resulting in sufficient density for a significant number of fusion reactions.

Directly-driven exploding-pusher targets have the largest direct-drive fusion yields recorded on the NIF, resulting in $E_{\text{fusion}} \sim 30$ kJ [252]. However, they are not suitable for high gain designs, as the low areal densities of the target are insufficient to confine α particles and thus enable burn propagation. A variety of interesting physics may be studied using the platform due to the significant ion temperatures that can be achieved, such as equilibration between electrons and ions [253] and kinetic effects in the ion population [254]. The strong shock is also highly kinetic, and thus accurate comparison to experimentally measurable variables, such as yields and ion temperatures, is expected to be difficult for Radiative-Hydrodynamics (Rad-Hydro) codes, which neglect kinetic effects [255]. However, much of the key dynamics and results can be studied more qualitatively.

5.2 Cross-Beam Energy Transfer in Unmagnetised Exploding-Pushers

This section presents simulation results, which demonstrate the effect of CBET in exploding-pushers on OMEGA. Both 1-D and 2-D CHIMERA-SOLAS simulations of 40-beam, pole-heavy drive, exploding-pusher experiments are presented, with a focus on how CBET acts to change the implosion. The 1-D results demonstrate that CBET significantly reduced the coupled laser energy to the implosion from 81.1% to 69.7%. Simulations conducted in 2-D, with a full 3-D ray-trace and CBET model, clearly demonstrated that the polar drive configuration led to an oblate implosion.

5.2.1 Simulation Configuration

The simulations conducted for this chapter aimed to study experimental configurations similar to the results from Bose *et al.*, discussed in Sec. 5.1.2 [243]. Experimental data from these magnetised exploding-pushers, demonstrated a clear amplification of the mode-2 due to magnetisation, which could have affected CBET scattering. The beam configuration, capsule initial conditions and pulse shape used for all simulations in this chapter are shown in Fig. 5.5. The exact pulse shape and target specifications were from a set of follow-up experiments to Bose *et al.* and were provided by C. W. Chang and J. Frenje from Massachusetts Institute of Technology (MIT) [256]. 40 beams from the OMEGA laser delivered a total of 17.7 kJ laser energy to a $2.5 \mu\text{m}$ thick, glass $(\text{SiO}_2)^2$ capsule, filled with room temperature and pressure D_2 . All experiments removed the 20 equatorial beams from the drive, because

¹When frequency-tripled light is not used, suprathermal electrons, rather than ablation, are the dominant driver of the strong shock [252].

²By ion number density, the material comprised 1/3 Si and 2/3 O.




Figure 5.5: The initial conditions used for all simulations presented in this chapter. Panel a) plots the D_2 filled, glass shell capsule and direction of the initial magnetic field. An example field coil (illustrative and not included in simulations) is also shown, the presence of which necessitated the polar laser drive in experiments. Panel b) plots the laser pulse shape used, which had a total of 17.7 kJ laser energy.

the presence of the small radius field coil in experiments precluded them. Magnetic fields of strength $B_{z0} = 0, 25$ and 50 T were applied along the z-axis of the configuration for different simulations.

An explicit $P_{1/3}$ radiation transport algorithm was used for all simulations, using tabulated opacities and emissivities from the SPK code [206]. A tabulated Sesame equation of state was employed for each material [205]. The CHIMERA extended-MHD package was used, which included resistive diffusion of the magnetic field, Lorentz force of the magnetic field on the hydrodynamics, Nernst-advection of the field down temperature gradients and anisotropic thermal conduction [257]. Thermal conduction was solved using a Super-time-stepping, semi-implicit algorithm [258], with flux-limited, Spitzer conductivities [207]. An electron flux limiter of $f_{lim,e} = 0.15$ was set for all simulations. Exploding-pushers exhibit high coronal temperatures due to the high-Z ablator and large temperatures for the strong imploding shock, therefore simulation yields and bangtimes are highly sensitive to thermal conduction and the choice of flux limiter. Although not the focus of this work, if attempting to match simulation result to experimental yields, reducing the flux limiter value would strongly inhibit the drive and therefore reduce simulation bangtimes. The flux limiter would therefore be the primary parameter to vary, in order to obtain better match the integrated results to experimental values. In order to improve the speed of the thermal conduction algorithm, which severely limited the stable simulation timesteps on vacuum-plasma interfaces, an artificial material was placed outside the capsule, which had an enforced ionisation state of $Z = 0$. This had minimal impact on the implosion results, but significantly sped up simulation run-times. The presence of the material did create large viscous heating of the ions as the coronal plasma expanded into it, which is visible in results throughout the chapter. However, this viscously heated layer was well separated from the region of interest (near critical density) for the majority of the implosion, and therefore did not impact upon the results.

All simulations used a spherical-polar CHIMERA mesh, with a fixed radial resolution of $0.5 \mu\text{m}$, over the full 4π str. 1-D simulations were run from beginning to end in spherical with $r \in [0, 1400] \mu\text{m}$. 2-D calculations assumed azimuthal symmetry of the hydrodynamics and conducted an initial ‘drive-phase’ in spherical coordinates, with $r \in [80, 1400] \mu\text{m}$, where the central cutout region was removed to avoid taking small radiation transport timesteps due to the small cell faces close to the axis. The 2-D drive-phase grid used 120 cells in the polar direction. When the shock reached this cutout region, hydrodynamic variables were trilinearly interpolated onto a 2-D cylindrical (r_{cyl}, z) mesh for the ‘stagnation-phase’, with fixed resolution, $\Delta r_{\text{cyl}} = \Delta z = 1 \mu\text{m}$, and simulation bounds $r_{\text{cyl}} \in [0, 1200] \mu\text{m}$ and $z \in [-1200, 1200] \mu\text{m}$. This grid contained no singularities, unlike a 2-D spherical grid without a cutout, and therefore excessively small radiation transport timesteps were not an issue.

All simulations used a 3-D laser ray-trace with a variety of CBET treatments. CBET was fully included for some simulations and neglected for others. Alternatively, partial-CBET simulations were conducted to explore the effect that CBET *spatial redistribution* of deposited power had on implosions. These simulations were conducted without the CBET model, but were forced to deposit the same magnitude of power that was absorbed from the equivalent CBET simulation. Explicitly, a no-CBET ray-trace was conducted, where the power of all rays was normalised to unity. When this ray-trace was complete, it read in the absorbed power from the (previously conducted) CBET simulation at that time, and multiplied the (normalised) deposited power by this value. This created a hydrodynamically similar implosion to the CBET ray-trace, and via comparison of these two simulations, the impact of CBET redistribution-of-deposition upon the hydrodynamics was studied. Temporally-and-spatially integrated results from all simulations are presented in Tab. 5.1, and are referred to when relevant throughout the chapter.

Several metrics have been included in the table, which are explicitly defined here. The metrics are commonly used to compare implosions in ICF, because, apart from the oblateness parameter in the last column, they can all be directly computed from experimental neutron spectra [259]. As previously stated, bangtime is defined as the time of peak neutron production, which for this simulation, assumes only deuterium-deuterium reactions contribute,

$$t_b = \operatorname{argmax}_t \left(\int Y_{DD}(\mathbf{x}, t) dV \right), \quad (5.2)$$

where Y_{DD} is the deuterium-deuterium neutron production rate, which is calculated across the CHIMERA computational grid at all locations and times throughout the simulation, using Bosch-Hale fits to the reactivity [260]. The total neutron yield is the spatially and temporally integrated neutron production,

$$Y_n = \int \int Y_{DD}(\mathbf{x}, t) dV dt. \quad (5.3)$$

The burn-width, Δ_b is the full width half maximum of the spatially integrated neutron production. One interpretation of this diagnostic, for this very two-dimensional configuration, is that a highly oblate implosion will have less temporally and spatially localised conver-

Table 5.1: Spatially and temporally integrated results from all simulations. In the CBET column, ‘~’ indicates partial-CBET, where CBET reduced absorption magnitude, but *not* spatial location, of deposition.

Run	Dim.	CBET	Note	B_{z0} (T)	t_b (ns)	$\langle T_i \rangle$ (keV)	Y_n ($\times 10^{10}$)	Δ_b (ps)	$\left. \frac{R_{\text{equator}}}{R_{\text{pole}}} \right _{t=t_b}$
1	1-D	Off	-	0	0.69	14.66	11.62	87	$1.00^{+0.00}_{-0.00}$
2	1-D	On	-	0	0.78	12.37	10.57	78	$1.00^{+0.00}_{-0.00}$
3	2-D	Off	-	0	0.71	8.44	6.20	148	$2.96^{+0.20}_{-0.19}$
4	2-D	~	-	0	0.75	7.61	5.23	153	$3.26^{+0.25}_{-0.23}$
5	2-D	On	-	0	0.75	7.77	5.46	148	$3.23^{+0.25}_{-0.23}$
6	2-D	Off	-	25	0.74	7.26	4.73	130	$3.80^{+0.41}_{-0.33}$
7	2-D	~	-	25	0.78	6.58	4.14	125	$4.55^{+0.50}_{-0.43}$
8	2-D	On	-	25	0.78	6.72	4.44	123	$4.32^{+0.47}_{-0.41}$
9	2-D	Off	-	50	0.73	6.82	3.73	134	$4.40^{+0.43}_{-0.38}$
10	2-D	~	-	50	0.78	6.30	3.30	130	$4.92^{+0.56}_{-0.48}$
11	2-D	On	-	50	0.78	6.37	3.52	129	$4.79^{+0.55}_{-0.47}$
12	2-D	Off	No Aniso.	25	0.81	6.40	5.02	118	$3.14^{+0.68}_{-0.50}$
13	2-D	Off	No Lor.	25	0.74	7.29	4.77	131	$3.80^{+0.41}_{-0.33}$
14	2-D	Off	No Nern.	25	0.73	7.41	4.85	130	$3.74^{+0.38}_{-0.32}$
15	2-D	Off	No Resis.	25	0.73	7.07	4.59	132	$3.84^{+0.42}_{-0.33}$

gence. Therefore, neutron production for less symmetric implosions will likely happen over a longer timescale, because thermonuclear conditions are produced at different times in different locations. The burn-averaged ion temperature, $\langle T_i \rangle$, is defined as the ion temperature, weighted by temporally and spatially resolved neutron production,

$$\langle T_i \rangle = \frac{\int \int T_i(\mathbf{x}, t) Y_{DD}(\mathbf{x}, t) dV dt}{Y_n}. \quad (5.4)$$

It is an integrated metric, that summarises the average temperature of the regions which are key in producing fusion-yield. Finally, the oblateness parameter in the final column of the table, $(R_{\text{equator}}/R_{\text{pole}})|_{t=t_b}$, was obtained by fitting an ellipse (with axes orientated along $\hat{\mathbf{x}}$ and $\hat{\mathbf{z}}$), to the spherical radius of maximum density at bangtime. The fitting procedure also returned asymmetric error bars, which are presented alongside the result.

5.2.2 1-D Simulations

In order to demonstrate how CBET typically affects exploding-pushers, two 1-D simulations were conducted, both with and without CBET included. The no-CBET and CBET simulation results are provided in Tab. 5.1, labelled as run 1 and run 2, respectively. Note that because the simulations were 1-D and therefore spherically symmetric, the ratio of equatorial radius




Figure 5.6: 1-D simulation results both without (left) and with (right) CBET. The top row plots the incident and absorbed energy from the simulation on the left axis and the radius of maximum density on the right. In order, the next rows plot n_e , T_e , T_i and V_r . The full-width half-maximum times of the D₂ yield are plotted as dotted vertical black lines on all panels.

to polar radius at bangtime was exactly unity. Both of these simulations used the beam configuration from Fig. 5.5, which unlike the hydrodynamics of the simulation, is not spherically symmetric. As described in Sec. 3.3.2, a sparse 3-D SOLAS mesh was created on which to compute the correct beam interaction in 3-D. For the 1-D simulations, 23 cells were used for the underlying spherical polar mesh in the θ direction and 54 cells in the ϕ direction, which were then combined to create a semi-structured SOLAS mesh, with approximately equal area across each spherical surface.

Absorbed laser power as a function of time from these two simulations along with streak plots of hydrodynamic variables are plotted in Fig. 5.6. The left column of panels is from the no-CBET simulation and the right is from the CBET simulation. The top row demonstrates that CBET reduced the absorbed power by $\sim 15\%$ during the implosion phase, up to about $t \sim 0.6$ ns. Total absorbed laser energy was reduced from 81.1% to 69.7%, from the

no-CBET to the CBET simulation. This reduced deposition led to a slower and weaker shock being driven ahead of the imploding SiO_2 material, which is most clearly visible in the fourth panel, plotting T_i on a log scale. CBET reduced the speed of the shock, such that it hit the axis for the with-CBET simulation at $t \sim 0.73$ ns compared to $t \sim 0.65$ ns without CBET. Bang-time occurred for both simulations after this ‘shock-flash’, as the rebounding deuterium fuel compressed against the in-falling shell, creating thermonuclear densities. Greater densities are observed for the with-CBET simulation at bangtime, because the in-falling shell collides the rebounding shock at a smaller radius, *i.e.* the shock timing of the implosion is better. This somewhat compensated for the lower T_i values due to the weaker shock, but overall, the no-CBET simulation still produced a $\sim 10\%$ higher neutron yield. Note, that the effect of the fictitious $Z = 0$ material, that was placed outside the capsule to speed up the super-time-stepping thermal conduction routine, is visible from the large T_i on the outer radius boundary of the expanding coronal plasma, which left the plot bounds at $t \sim 0.2$ ns. Although not presented here, 1-D simulations with vacuum outside the glass ablator initially, showed minimal difference in bangtime hydrodynamic conditions and therefore integrated diagnostics. Thus, the substantially faster configuration with the fictitious ‘vacuum’ material was used for all 2-D simulations.

Compared to the streak plot of a more conventional, hotspot ignition implosion in Fig. 4.6, clear differences can be seen. Firstly, the simulations presented in Fig. 4.6 did not include radiation transport, which is the origin of the preheat ahead of the shock³ in Fig. 5.6. Secondly, the hotspot ignition design maintained a cold dense shell throughout the implosion phase, whereas the initially thin shell of the exploding-pusher simulation was approximately volumetrically heated. Therefore, relatively little mass was left in the glass shell when the rebounding shock collided with it. Finally, the T_i increase in the exploding-pusher design was predominantly from the spherical convergence of the strong shock, compared to the stagnation heating in the hotspot design, which was localised to inside the decelerating shell when it compressed on axis, at $t \sim 2.5$ ns in Fig. 4.6. The lack of the cold dense shell is the reason that the exploding-pusher design cannot scale to high yield, because they have insufficient areal density to confine α particles and thus sustain a burn wave.

5.2.3 2-D Simulations

Simulations of this target configuration were also conducted in 2-D to explore how the non-uniformity of the laser-drive affected the implosion, both with and without CBET. 2-D simulations used a 3-D SOLAS mesh to resolve the beam overlap pattern, with 58 cells in the azimuthal direction. Fig. 5.7 provides plots, which illustrate the progression of the simulation during the drive-phase. The left column plots T_e , T_i and n_e from the 2-D, $B_{z0} = 0$ T simulation without CBET. Integrated results from this simulation are given in the row labelled run 3 in Tab. 5.1. The right column plots the same for the equivalent simulation, but with the full effect of CBET included, which corresponds to the row labelled run 5 in Tab. 5.1. Note that the viscous ion-heating of the fictitious material (placed outside the glass shell initially) is

³*i.e.* The temperature increase (visible in T_e and T_i plots), which hit the axis at $t \sim 0.2$ ns.




Figure 5.7: n_e (right-side), T_e (top-left-side) and T_i (bottom-left-side) plots from the 2-D, 0 T simulations without (left-column) and with (right-column) CBET at a variety of in-flight times. The decreased deposited energy due to CBET, results in lower coronal electron temperatures and therefore a slower, weaker shock being driven, which is especially evident at later times.

again visible as large ion temperature in the layer immediately outside the coronal plasma expansion.

Particularly at the later times plotted in Fig. 5.7, the decreased, coronal T_e of the CBET

compared to the no-CBET simulation is visible. The lower T_e was due to of reduced absorption, because of CBET scattering, and led to the shock from the no-CBET simulation imploding significantly faster than the CBET shock. Thus, the no-CBET implosion was more oblate during the drive-phase, because the shock travelled primarily along the $\hat{z}$ -axis, due to the beam-geometry. However, when including CBET, the velocity of the shock was reduced more than the velocity of the imploding portion of the glass shell, *i.e.* $|V_{r,\text{shock}} - V_{r,\text{shell}}|$ was larger for the no-CBET simulation than the CBET simulation. Therefore, the no-CBET shock, after rebounding from the axis, collided with the shell and underwent maximal neutron production at a larger radius than for the CBET result. Thus, the oblateness parameter at bangtime in Tab. 5.1, $(R_{\text{equator}}/R_{\text{pole}})|_{t=t_b}$, is significantly larger for the CBET calculation.

Both 2-D simulations exhibited substantially reduced burn-averaged ion temperatures, compared to their 1-D equivalents. This was primarily due to the strong shock travelling mainly along $\hat{z}$, rather than radially inward as for the 1-D calculations. Therefore, the convergence of the shock was reduced, leading to lower temperatures, and therefore neutron yields. This interpretation is also corroborated by the increased Δ_b of all 2-D calculations compared to 1-D. The burn-width was larger, because thermonuclear conditions were produced at different times throughout the hot fuel, as the rebounding D₂ compressed against the infalling shell material, compared to the 1-D where it was spherically symmetric.

To summarise results from the 1-D and 2-D unmagnetised implosions, it has been shown that CBET substantially reduced the deposited power for these exploding-pusher calculations. 1-D spherical calculations, which averaged the deposited power across all angles, demonstrated that this substantially reduced deposition led to a weaker shock being driven, which reduced thermonuclear yield and delayed the bangtime. When 2-D effects were included, which better reflected the geometry of the laser-drive, this resulted in an oblate implosion, which reduced the convergence of the shock, and therefore the ion temperatures. The reduced drive in the CBET calculations also slowed the shock speed more than the infalling material, so the oblateness at bangtime (when the rebounding shock collided with the in-falling material) was greater when including CBET.

5.3 Magnetisation in Exploding-Pusher Implosions

This section presents the effect that various extended-MHD terms had on the magnetised, 2-D exploding-pusher simulations. Simulations were conducted with an initial field strength $B_{z0} = 25\text{T}$ and particular terms turned off, to deduce what the important physical processes were. The origin of the field structure is presented, which demonstrates that in the plasma corona, field lines are mainly radial due to the field being frozen in to highly conductive, radially outflowing plasma. Anisotropic thermal conduction in this highly magnetised coronal plasma acted to keep heat localised to the polar regions, which enhanced the drive on the pole relative to the waist. The results demonstrate that resistive diffusion and the Lorentz force have very little impact on the implosion physics, due to the bulk of the plasma being highly conductive and high- β , respectively. Nernst-advection of the magnetic field acted to significantly redistribute the field in the low Hall parameter, equatorial region of the capsule,




Figure 5.8: The development of the hydrodynamic variables and magnetic field structure from the $B_{z0} = 25$ T, CBET simulation. Panels a), b) and c) plot T_e (top-left), T_i (bottom-left), n_e (right) and $\mathbf{B}$ (streamlines) at $t = 0.15$, 0.30 and 0.45 ns respectively. The approximately radially outward flowing, hot (and therefore highly conductive) ablating plasma pulled the magnetic field with it, resulting in radial $\mathbf{B}$ field lines, which were weaker at the capsule equator. Panels d), e) and f) plot n_e lineouts along the pole ($\theta = 0^\circ$) and equator ($\theta = 90^\circ$). Panels g), h) and i) plot equivalent T_e and T_i lineouts. These all show that the increased polar temperatures, partially due to beam geometry and partially due to magnetisation, led to preferential ablation along the pole.

which formed a ‘divot’ in density on the capsule waist. This divot was however well separated from the region where burn was important, and thus had minimal impact on integrated neutron diagnostics.

5.3.1 In-Flight Field Structure

Initially, the development of the coronal field structure from the $B_{z0} = 25$ T, no-CBET simulation (labelled as run 5 in Tab. 5.1) is presented. The top row of Fig. 5.8, plots the drive phase hydrodynamic profiles, overlaid with streamlines of the magnetic field, coloured by its magnitude, at three different times. Lineouts of n_e and the temperatures are plotted in




Figure 5.9: Plasma β (left-side) and Magnetic Reynolds Number, R_m , (right-side) at various in-flight times, throughout the $B_{z0} = 25$ T, no-CBET simulation. Contours of the $n_e = n_{cr}/10$ are plotted on all panels as dashed green lines to indicate the bounding region, which contained a significant amount of plasma. Broadly, the β and R_m values are $\gg 1$ in all regions with an appreciable amount of material, which demonstrates that the Lorentz force and resistive diffusion should have minimal effect on the implosion dynamics.

the middle and bottom rows respectively, along both the polar and equatorial directions. The hot coronal plasma was highly conductive, shown explicitly in Sec. 5.3.2, and thus the field remained frozen in to the plasma. As the coronal plasma expanded outward therefore, it dragged the field lines with it, leading to $\mathbf{B} \sim \pm |\mathbf{B}| \hat{\mathbf{r}}$ in this, laser-heated region. The geometric stretching of the field lines at the target poles was less significant than at the equator, and therefore the coronal field strengths were smallest at the target equator and highest on the poles. As the target began to implode, the field compressed on the interior edge of the dense shell, resulting in non-radial field lines and an increase in field strength. This effect is most clearly visible at $t = 0.45$ ns, in Fig. 5.8.c in the vicinity of the dense shell material region.

The lineouts clearly demonstrate that the preferential heating of the target on the pole led to faster ablation of the shell along this direction. This led to a much stronger shock along the pole, which is seen most clearly by the discrepancy in ion temperature between the pole and equator at $t = 0.45$ ns in Fig. 5.8.i. Increased polar electron temperature is partially due to the pole heavy drive from the beam geometry, and also due to anisotropic thermal conduction. The field structure plotted in Fig. 5.8 inhibited equilibration of temperature via thermal conduction in the polar direction, which increased the temperature asymmetry compared to the unmagnetised simulation. Anisotropic thermal conduction and its impact in this configuration is discussed in more detail in Sec. 5.3.3.

5.3.2 Resistive Diffusion and the Lorentz Force

As was stated in the above Sec. 5.3.1, the coronal plasma was highly conductive for all simulations, due to the combination of high temperature and relatively low density. The right side of each panel Fig. 5.9 plots the Magnetic Reynolds Number, R_m , at three times during the implosion of the $B_{z0} = 25$ T, no-CBET simulation, labelled as run 6 in Tab. 5.1. As was stated in Sec. 2.2.4, R_m is the ratio of magnetic field advection to diffusion. Conservatively,

a $1\ \mu\text{m}$ length scale was used to calculate this value, which was approximately the smallest length scale observed throughout the implosion across all times. Thus, for the majority of the simulation, R_m is likely underestimated, because $R_m \propto L$. The green dashed lines are contours of $n_e = n_{\text{cr}}/10$, which are included in the plots to illustrate the regions of the simulation that contained the majority of the plasma material. Although Fig. 5.9.a shows that early in the simulation, the gas fill of the capsule had a small value of R_m , and therefore resistive diffusion dominated over frozen in flow, the dynamics of the implosion had not reached this region yet, *i.e.* it is not bounded by the n_e contours. The radiative preheat also raised the temperature, and therefore R_m value in this region, before the time plotted in Fig. 5.9.b. Therefore, the majority of the plasma material had a high value of R_m throughout the implosion, mainly due to the high temperatures, demonstrating that frozen-in-flow dominated over resistive diffusion. This is corroborated by comparing run 6 and run 15 in Tab. 5.1, which did not include the effects of resistive diffusion on the field transport. There is minimal difference across all metrics between these two rows, which shows that the global dynamics of both simulations were very similar.

The left side of the panels in Fig. 5.9 plots the plasma β at the same times for the same simulation. As was stated in Sec. 2.2.4, this is the ratio of thermal plasma pressure to magnetic pressure and thus describes whether thermal forces or the Lorentz force, is likely to predominantly influence the plasma dynamics. Just as was seen for R_m , the value of this parameter was generally high in all regions with appreciable plasma density throughout the implosion, so the Lorentz force was assumed to minimally influence the plasma dynamics. An additional simulation was performed, in which the effect Lorentz force was not included on the plasma, labelled by run 13 in Tab. 5.1. Again, there is minimal difference in integrated metrics between this and run 6 (the comparable run which included the Lorentz force), which verifies this analysis.

5.3.3 Anisotropic Thermal Conduction

In order to understand the effect of anisotropic thermal conduction, an additional simulation was performed, which had the same setup as the $B_{z0} = 25\ \text{T}$, no-CBET simulation, but thermal conduction was isotropically suppressed by the local magnetic field strength, regardless of its orientation. The anisotropic and isotropic simulations are labelled by run 6 and 12 respectively in Tab. 5.1. Explicitly, the parallel conductivity was forced to take the value of the perpendicular conductivity, $\kappa_{\parallel} = \kappa_{\perp}$. Comparison of this isotropically suppressed conduction simulation, with the $B_{z0} = 25\ \text{T}$, no-CBET simulation, elucidated the role of anisotropic conductivities relative to the orientation of the field. Fig. 5.10.a plots the ion (left-side) and electron (right-side) Hall parameters from the $B_{z0} = 25\ \text{T}$, no-CBET, anisotropic thermal conduction simulation at $t = 0.5\ \text{ns}$. It can be seen that the polar region of the corona was significantly magnetised, with $1 < \omega_e \tau_e < 10$, therefore thermal conduction in this region was significantly suppressed perpendicular to the field in the anisotropic simulation, which was approximately the polar direction, as is seen in Fig. 5.8.c. Ion Hall parameters peaked at values of $\omega_i \tau_i \sim 10^{-2}$ in-flight, thus ion transport was not expected to




Figure 5.10: In-flight a) and bangtime c) Hall parameters, from the $B_{z0} = 25$ T, no-CBET simulation. Panel b) plots the T_e from the isotropically magnetised simulation (left-side) and anisotropic conduction simulation (right-side) in-flight. Panel d) plots the same, but at bangtime. The electron Hall parameter was > 1 at the poles due to high magnetic fields and temperatures, which led to significantly restricted thermal conduction from magnetised transport. Isotropically magnetised conduction, therefore resulted in a markedly different bangtime morphology, as is shown in panels b) and d). Ion Hall parameters peaked at bangtime, when values reached about $\omega_i \tau_i \sim 0.1$.

be significantly affected by the field.

The electron temperature profiles at the same time are plotted in Fig. 5.10.b for the isotropically magnetised conduction simulation (left-side) and anisotropic simulation (right-side). The T_e profile of the anisotropic case was significantly more oblate than the isotropic simulation and the polar regions of the isotropic region were also significantly hotter. This was due to decreased transport of deposited laser energy in the isotropic to the ablation region, because heat could not freely stream radially inward along the field lines when conductivity was isotropically suppressed by the field. Isotropic magnetisation of electron transport thus led to reduced ablation at the poles and a larger critical radius along $\pm \hat{z}$, which gave the implosion and shock a rounder shape.

Bangtimes temperature profiles, plotted in Fig. 5.10.d, corroborate this analysis, by demonstrating that the isotropic simulation had a significantly rounder shape. The oblateness parameter in Tab. 5.1, obtained by fitting an ellipse to the radius of maximum density at bangtime, also demonstrates a less oblate profile. By isotropically magnetising the conduction, $(R_{\text{equator}}/R_{\text{pole}})|_{t=t_b}$ was reduced from $3.80 \rightarrow 3.14$. The reduced ablation coupled less energy to the implosion, evident in the reduced bangtime from $t_b = 0.74 \rightarrow 0.81$ ns, going from




Figure 5.11: Panel a) plots in-flight electron density profiles from the $B_{z0} = 25 \text{ T}$, no-CBET simulations with (left-side) and without (right-side) Nernst-advection of the magnetic field. Panel b) plots magnetic field magnitude from the Nernst (left-side) and no-Nernst (right-side) simulation along with Nernst-advection velocity streamlines from the Nernst-on simulation. Advection of the field is important in the low Hall parameter equatorial region, pulling $\mathbf{B}$ down ∇T_e , into the dense wall. Altered field at the equator impacts on the magnetised thermal conduction, which ultimately imprints on the density, as is seen in panel a).

anisotropic to isotropic thermal conduction. Yield still increased for the isotropic simulation from $4.72 \rightarrow 5.02 \times 10^{10}$, which shows that integrated yield metrics are a fine balance of many factors, including coupled energy and implosion shape.

The bangtime Hall parameters are plotted in Fig. 5.10.c, and show highly magnetised electrons and moderately magnetised ions, with a maximum $\omega_i \tau_i \sim 10^{-1}$. Note that this is also for an initial field strength $B_{z0} = 25 \text{ T}$ simulation. For the $B_{z0} = 50 \text{ T}$, with-CBET simulation, which included CBET fully and is therefore the most realistic simulation to compare with experiment at this field strength, the maximum ion Hall parameter was also $\omega_i \tau_i \sim 10^{-1}$. This is lower than the reported value of $\omega_i \tau_i \sim 1$ in the original paper [243]. However, improved sphericity of the converging shock, which could change slightly for example by changing $f_{\text{lim},e}$, would markedly boost this value. The Hall parameter, $\omega \tau \propto T^{3/2} |\mathbf{B}|$, so a less oblate shock would increase the value this by both better flux compression of the field, and higher stagnation temperatures, due to greater convergence of the shock.

5.3.4 The Nernst Effect

The final extended-MHD term that was independently examined was the Nernst effect. In CHIMERA, Nernst is implemented as an advection of $\mathbf{B}$ down ∇T_e [257]. It is a collisional term, and therefore dominates in plasma with low values of $\omega_e \tau_e$. The plots of Hall parameter in Fig. 5.10 demonstrate that the coronal plasma has low Hall parameters at the equator,

which suggests that the Nernst effect would be most important at the equator of the capsule and have minimal effect on the field profiles near the poles. An additional simulation, with $B_{z0} = 25$ T and no-CBET, was conducted without Nernst-advection of magnetic field to understand its impact upon the field and plasma dynamics. The simulation without the Nernst effect is labelled in Tab. 5.1 as run 14.

Fig. 5.11.b plots the magnetic field strength from equivalent simulations with-Nernst (left-side) and without-Nernst (right-side) at $t = 0.60$ ns. The simulation including Nernst (run 6 in Tab. 5.1), also plots the Nernst-advection velocity as streamlines, coloured by the speed. The largest advection speeds are $\sim 10^6$ ms $^{-1}$, which is comparable to the coronal fluid ablation speed, demonstrating that Nernst-advection is non-negligible in dictating the magnetic field structure. As expected, the equator of the target had the largest Nernst speeds, due to the low Hall parameter, and it acted to push field from the coronal plasma onto the interior edge of the shell. This created a localised field ‘pile-up’ in this region for the Nernst simulation, compared to the simulation on the right.

Fig. 5.11.a plots the electron density at the same time for the Nernst (left-side) and no-Nernst (right-side) simulations alongside magnetic field streamlines from the same simulations on the top-half of the plot. The streamlines show that the Nernst effect caused the coronal field lines at the equator to be radial, by pushing B_z down the mostly radial temperature gradient, towards the shell. It is visible that the altered magnetic field structure impacted on the plasma density indirectly, creating an equatorial bump in the shell density, which is not present in the right-side plot without Nernst included. This bump only appeared in simulations with both Nernst and anisotropic conduction, which shows that this altered field profile changed the anisotropic thermal conduction at the equator, which then impacted upon the density. The effective re-orientation of $\mathbf{B}$ at the equator meant that thermal conduction, which transported laser energy from the deposition surface to the ablation surface, experienced the higher, parallel thermal conductivity $\kappa_{\parallel}$, rather than $\kappa_{\perp}$. This locally increased the drive along $z \sim 0$, leading to the bump at the equator.

Comparison of run 6 (Nernst) and run 14 (no-Nernst) in Tab. 5.1 shows minimal difference in the integrated metrics, despite the altered dynamics. This is mostly because the changes were localised toward the equator. This material was thus well separated from the polar region, which experienced the strongest shock convergence and therefore greatest temperatures and neutron production. It is expected however, that experiments with symmetric laser illumination, and therefore a less oblate implosion, may be more influenced by the Nernst effect, because more neutron production would occur near the equator of the shell. Additionally, it is worth noting that this bump is likely to only be observed in magnetised direct-drive implosion and not indirect-drive. This is because the main action of Nernst in these simulations was to change the field and therefore anisotropic transport in the conduction zone, which is the layer in direct-drive implosions that transports energy from the laser absorption region to the ablation surface. However, for indirect-drive, x-rays penetrate effectively all the way through to the ablation surface, so there is no conduction zone. Thus, the altered conduction zone field profile would have a much less significant impact on the density for indirect-drive.




Figure 5.12: Comparison of n_e and T_e profiles from the $B_{z0} = 0$ (panel a) left-side) and 25 T (panel a) right-side), with-CBET simulations, both at $t = 0.5$ ns. Panel a) also plots streamlines of $\mathbf{B}$ for the $B_{z0} = 25$ T simulation, coloured by the field magnitude. Panels b) and c) plot n_e and T_e lineouts respectively, along both the pole and equator. It is evident from these lineouts that magnetisation anisotropically affects hydrodynamic variables, which are used to calculate the CBET gain. Therefore, it is anticipated that magnetisation could anisotropically affect the CBET scattering volume.

5.4 The Effect of Magnetisation on Cross-Beam Energy Transfer and Stagnation

This section presents results on how the magnetisation of the corona affects both CBET scattering and the stagnation shape of the implosion. Fig. 5.12.a plots T_e and n_e from the $B_{z0} = 0$ (left-side) and 25 T (right-side), with-CBET simulations at $t = 0.5$ ns. Fig. 5.12.b and Fig. 5.12.c plot lineouts of n_e and T_e respectively, along both the poles and equator for both simulations. As was shown in Sec. 5.3, anisotropic thermal conduction is the dominant effect of magnetisation in these implosions. The equatorial lineouts are similar, due to low Hall parameters in this region and thus minimal impact of magnetisation upon the transport. However, the polar lineouts are significantly different in the magnetised and unmagnetised cases. Magnetisation has amplified the $\ell = 2$ asymmetry of the coronal plasma n_e and T_e profiles, which are used to calculate the CBET gain. It was therefore hypothesised that the amplified asymmetry in coronal profiles due to magnetisation, would leave a signature upon the CBET and therefore implosion dynamics, because CBET is known to interact strongly with low mode, coronal asymmetries [247].

In this chapter, ‘partial-CBET’ simulations shall be presented, for which the coupled power as a function of time was kept the same as the equivalent with-CBET simulation, *i.e.* so CBET only acted to reduce the magnitude of the deposited power, rather than redistribute it around the target. It was hoped that, by comparison with the full-CBET simulations, this would demonstrate that CBET changed the shape of the implosion differently, for different levels of seed magnetic field. The results presented in this section demonstrate that the $\ell = 2$ of the long-wavelength, density perturbation was slightly reduced by CBET, consistent with exist-




Figure 5.13: Various CBET diagnostics used in the analysis presented in this section. All plots are from the $B_{z0} = 0$ T, with-CBET simulation at $t = 0.30$ ns. Panel a) plots the instantaneous deposition with (right-side) and without (left-side) the effect of CBET. The ‘CBET-deficit’, ΔP_{CBET} , which is the no-CBET deposition, subtracted from the with-CBET deposition, is plotted in panel b). Panel c) plots the radially integrated ΔP_{CBET} , as a function of polar angle ($\theta = 0, 180^\circ$ are the poles). The ‘CBET-scattering’, is plotted in panel d), and the radial integral is plotted in panel e). At this time, it is evident from panel e) that more CBET occurred at the capsule poles, resulting in a CBET-reduction of deposition near the poles, seen in panel c). The critical surface is shown in all 2-D panels by a dashed magenta line.

ing literature on how CBET mitigates $\ell = 1$ asymmetries [67]. Results also showed that the increasingly anisotropic coronal plasma profiles for increasing seed magnetic field strengths did lead to changes in the CBET scattering. However, this was too small an effect to lead to experimentally observable changes in density and temperature. It is hypothesised that the polar beam geometry, the shock driven characteristics of the exploding-pusher implosions and the high Z shell (which results in increased coronal temperatures, and therefore reduced CBET [152]), minimised the impact on CBET scattering. Suggestions for an alternative implosion are provided, which would increase the likelihood of an experimentally observable signature.

5.4.1 Analysis and Key Definitions

Firstly, key quantities used for the analysis of this section shall be introduced. The key aim of this section was to discover if CBET acted differently in the magnetised coronal hydrodynamic profiles, compared to the unmagnetised profiles. Particular focus was dedicated

to discovering if this affected the implosion shape, because this shape change would not be captured by a laser-MHD model that did not have a full-CBET capability, but rather reduced the incident laser energy to compensate for CBET. Two quantities were used to understand how CBET redistributed the deposited power. The first was the ‘CBET-deficit’, which is defined as the difference in deposited power computed by a CBET and no-CBET ray-trace through the *same* hydrodynamic profiles,

$$\Delta P_{\text{dep}}(\mathbf{x}, t) = P_{\text{dep}}^{\text{CBET}}(\mathbf{x}, t) - P_{\text{dep}}^{\text{no-CBET}}(\mathbf{x}, t), \quad (5.5)$$

where $P_{\text{dep}}^{\text{CBET}}$ is the deposition from the calculation including CBET and $P_{\text{dep}}^{\text{no-CBET}}(\mathbf{x}, t)$ is without CBET. As was described in Sec. 3.5.5, a CBET calculation in SOLAS always begins with a ray-trace, which does not include any CBET effects (the field-reconstruction ray-trace), so $P_{\text{dep}}^{\text{no-CBET}}(\mathbf{x}, t)$ is the deposited power from this ray-trace. $P_{\text{dep}}^{\text{CBET}}(\mathbf{x}, t)$ is the value after the pump depletion and energy conservation iterations, *i.e.* after CBET has been fully accounted for.

Fig. 5.13.a plots the deposited power without (left-side) and with-CBET (right-side), from the $B_{z0} = 0$ T, with-CBET simulation at $t = 0.30$ ns. The corresponding CBET-deficit, which is the difference between these plots, is shown in Fig. 5.13.b. Note that the value is almost all negative, because for direct-drive configurations, CBET mainly acts to reduce the intensity of light near the critical surface where Inverse-Bremsstrahlung (Inv-Brem) predominantly occurs, and therefore it reduces the deposited power. There are small ($\sim 1\%$) increases in deposition at radii outside the peak scattering surface, near Mach-1, but this increase is not visible to the naked eye when the maximum colourbar scale is increased above zero. It can be seen from the CBET-deficit, that at this time, CBET decreased the absorption at the poles more than at the equator of the capsule. This can be understood by considering the geometry of the pole heavy drive. The beam geometry meant that the intensity was greater in the polar coronal plasma, where many beams overlapped each other. Therefore, inbound laser sheets encountered stronger CBET resonances, because more reflected sheets were present in this region, thus their energy was depleted more than at the waist. The radial integral of the CBET-deficit is plotted in Fig. 5.13.c. This conveniently illustrates the previous point, that the reduction in absorption due to CBET (specifically at the time plotted, for this simulation configuration), was greater near the poles than the equator. It is also worth noting that if CBET acted symmetrically to simply reduce the absorption rather than redistribute it, then Fig. 5.13.c would be uniform in θ . This is the case for the ‘partial-CBET’ simulations, labelled by ‘ $\sim$ ’ under the CBET column in Tab. 5.1, where the magnitude of the power was reduced to account for CBET, but the redistribution of power was not accounted for.

The second metric was the CBET scattering, G_{CBET} , defined in Sec. 3.6.1, but restated here for convenience,

$$g_{\text{CBET}}(\mathbf{x}, t) = \sum_i^{\text{rays}} |\Delta P_{i,\text{CBET}}(\mathbf{x}, t)|, \quad (5.6)$$

where the time dependence has now been made explicit and $\Delta P_{i,\text{CBET}}$ is the power change of a ray due to CBET. This quantity shows where the majority of CBET power change occurred,




Figure 5.14: The radially integrated CBET-deficit and CBET-scattering, plotted as a function of angle and time for the $B_{z0} = 0$ and 50 T, with-CBET simulations. Panel b) plots the CBET-deficit from the 50 T (top-half) and 0 T (bottom-half) simulations. Lineouts in θ at $t = 0.25$ (black), 0.50 (dark-green) and 0.75 ns (light-green) are plotted in panel a). The same results, but for CBET-scattering are plotted in panels d) and c). It is evident from these plots that for both simulations, as time passed and the poles of the capsule fell in faster than the equator, the region where CBET mostly occurred, shifted in angle around the capsule. Differences are visible in all plots between the $B_{z0} = 0$ and 50 T simulations, indicating that magnetisation affected CBET indirectly, via the altered hydrodynamics.

across the simulation domain, at a given time. Fig. 5.13.d plots g_{CBET} on a log scale, from the same simulation and at the same time as Fig. 5.13.a. The radial integral of this plot is shown in Fig. 5.13.e, as a function of angle. It is evident by comparison of Fig. 5.13.c and Fig. 5.13.e that the angles of maximal CBET scattering align with the most significant deficits, and vice-versa. It should be noted that there is some noise visible in both CBET scattering plots near the poles. This was presumed to be due to small cells and therefore poorer ray-per-cell statistics in this region. This noise was, however, localised to a few polar-cells, and therefore not thought to have a significant impact on the overall dynamics.

5.4.2 Spatial Change of CBET and Deposition from Magnetisation

This section shall explore differences in the CBET-deficit and CBET-scattering between magnetised and unmagnetised implosions. As was illustrated in Fig. 5.4, magnetisation significantly amplified coronal plasma asymmetries in direct-drive implosions. Thus, an effect on the CBET gain calculations was expected. How changes to CBET affect the hydrodynamics

and stagnation shape of the target, is discussed in the subsequent Sec. 5.4.3.

Fig. 5.14.b plots the CBET-deficit as a function of time (x axis) and polar angle (y axis) for the with-CBET, $B_{z0} = 50$ T (top-half) and $B_{z0} = 0$ T (bottom-half) simulations, respectively. Note that the colour scale has a maximum of zero, and is most saturated at negative values, because the radially integrated CBET-deficit was entirely negative. Lineouts in θ at three times, $t = 0.25, 0.50$ and 0.75 ns (times indicated by dashed lines on Fig. 5.14.b), are plotted in Fig. 5.14.a, where the x axis of this plot share the colour scale from Fig. 5.14.b. Broadly, both the top half and the bottom half of Fig. 5.14.b appear to be qualitatively similar, indicating that magnetisation of coronal profiles did not play a dominant role in altering CBET. The simulations featured a large CBET-deficit of deposition near the poles between $t \sim 0.2 \rightarrow 0.3$ ns. As the implosion progressed, the asymmetry of the deficit reduced, which is visible as the reducing saturation of the plot, with increasing time. This was, however, also partially due to the lengthening of the plasma corona, which increased symmetry primarily through increasing deposition at larger radii. Large radius deposition is less efficient at coupling to ablation, compared to if the energy were deposited close to critical, due to thermal conduction having to transport it further. An improved metric, which appropriately weights deposition by distance from the critical surface, may further elucidate the role of CBET. Both plots show that the peak of the CBET-deficit shifted to the equator as the simulations progressed.

The CBET-scattering streak plot and lineouts in Fig. 5.14.d and 5.14.c, respectively, show similar behaviour. Before $t \sim 0.2$ ns, the level of CBET-scattering was low, primarily due to the low incident power, and therefore laser field strengths. Maximal scattering occurred near the poles, at early times in both the magnetised and unmagnetised simulations, which then moved more to the equator later on. The movement of the scattering and deficit to the equator was due to the asymmetric convergence of the target. As can be seen from the time resolved n_e profiles from the with-CBET simulation in Fig. 5.7, the coronal plasma was relatively round up until $t \sim 0.35$ ns, but then began to implode more quickly along the $\pm \hat{z}$ axis than the equator. When this occurred, the trajectory of light changed asymmetrically, leading to a change of polar-angle where CBET predominates.

Firstly, the expected change to the bangtime profiles from the similar, qualitative behaviour shall be described. Bangtimes for these implosions were at $t \sim 0.75$ ns, so the late time behaviour shortly before and after this, was unimportant in dictating the bangtime shapes. The most significant ΔP_{CBET} values occurred between $t \sim 0.2 \rightarrow 0.3$ ns at the poles, for both the magnetised and unmagnetised simulation. This CBET-induced power redistribution would have reduced the strength of the shock launch along the poles, compared to the equator. Therefore, simulations with an equal magnitude of power deposition, but no spatial relocation of deposition (*i.e.* the partial-CBET simulations), were expected to exhibit a larger $\ell = 2$ asymmetry, because CBET appeared to reduce the asymmetry of the drive.

Finally, the change to the CBET-deficit due to magnetisation of the coronal plasma shall be explicitly examined. If the altered coronal plasma profiles due to magnetisation had played a dominant role in changing CBET, then the top and bottom plots would look markedly different. The similar qualitative behaviour seems to suggest therefore that this is not the

dominant role in dictating implosion symmetry. However, minor differences are observable, for example the magnetised simulation exhibited more early-time CBET-scattering and a slightly larger CBET-deficit at the poles, which also lasted for a slightly longer period of time. This suggests that the shock launch in the $B_{z0} = 50$ T case would be slightly weaker than the $B_{z0} = 0$ T case. A significant reason for the relatively small changes, is that the pole-heavy drive obscures these, more subtle, changes to implosion dynamics. Explicitly, an unmagnetised, symmetric drive beam-configuration should not display an $\ell = 2$ asymmetry in either of the diagnostics, and should thus result in a mostly symmetric bangtime profile. Changes to the shape would therefore be much less obscured compared to these simulations, which feature a global asymmetry imposed by the 40-beam drive.

5.4.3 Stagnation Profiles

The bangtime n_e , T_e and $\mathbf{B}$ profiles from the 2-D simulations with all magnetisations (columns) and all CBET treatments (rows) are plotted in Fig. 5.15. The top row of figures correspond to full CBET treatment and the bottom row, to simulations where no effect of CBET is included. As a reminder, to conduct the partial-CBET simulations in the middle row (labelled ‘~CBET’ for shorthand), the deposited power as a function of time was loaded in from the (already completed) full-CBET simulation with the same initial B_{z0} value. No-CBET ray-traces were then conducted with incident power normalised to unity, and the deposition was multiplied by the absorbed power from the full-CBET simulation at the same time. Thus, full-CBET and ~CBET simulations with equivalent B_{z0} have identical power deposition as a function of time. The only difference between them is that the full-CBET simulations, on the top row, include redistribution of deposition around the corona.

Rows labelled run 2 $\rightarrow$ 11 in Tab. 5.1 provide the integrated metrics for each of these simulations. The first thing to note about these results, is that all implosions were oblate, *i.e.* the radius of peak density was significantly larger in the equatorial direction than polar. Magnetisation also evidently flattened the implosion, regardless of CBET treatment. No-CBET simulations (bottom row) were also significantly less oblate than the top two rows, which both included reduced absorption due to CBET. To remind the reader, it was due to the larger $|V_{r,\text{shock}} - V_{r,\text{shell}}|$ for the no-CBET simulation, compared to when CBET was included, as was described in 5.2.3. Thus, the collision between the rebounding shock and in-falling fuel material happened at larger radii, resulting in a rounder bangtime shape.

The stagnation field structure for the magnetised simulations exhibited field lines at $z \sim 0$ μm which point along $\pm\hat{\mathbf{x}}$. This was due to the oblate structure of the initial strong shock, which moved mainly along $\pm\hat{\mathbf{z}}$ for a wide array of x values. When the shocks from the two poles met at $z = 0$ μm , a thin layer was superheated to extreme temperatures of several 10keV. This thin layer of highly conductive plasma, reoriented the un-shocked $+\hat{\mathbf{z}}$ field to point in the $\pm\hat{\mathbf{x}}$ directions. The large compressed field strengths also increased the peak temperatures significantly, by insulating thermal losses. Note however, that this increase was highly anisotropic and was localised to lower density regions of the hotspot, therefore did not directly translate to increased burn-averaged temperatures in Tab. 5.1. For simula-




Figure 5.15: n_e , T_e and $\mathbf{B}$ bangtime profiles for different initial magnetisation (columns) and CBET treatments (rows). Panels a), b) and c) are from the full CBET simulations, which show that magnetisation increased the oblateness of the stagnation profile. Panels d), e) and f) are from the simulations where only the CBET effect on the magnitude, but not spatial location, of deposition was included. These simulations have identical coupled energy to the top row, and therefore have the same bangtimes. Panels g), h) and i) are from the no-CBET simulations, which had earlier bangtimes and increased temperatures, due to the greater coupled energy.

tions which exhibited larger hotspot temperatures, the higher pressures appeared to develop some form of magneto-Rayleigh–Taylor Instability (RTI) at the equator of the capsule, most evidently visible in T_e at $x, z \sim [-150, 0] \mu\text{m}$, although the exact dynamics of this have not been studied in detail.

If CBET was anisotropically affected by magnetisation, this would result in spatial differences in deposition location, between the full CBET (top row) and partial-CBET simulations (middle row), and therefore the bangtime profiles would be different. Additionally, higher




Figure 5.16: Difference in n_e and T_e bangtime profiles, between the full-CBET and partial-CBET a) $B_{z0} = 0$, b) $B_{z0} = 25$ and c) $B_{z0} = 50$ T simulations. The difference in variable ν , is defined as $\Delta\nu = \nu_{\text{full-CB}} - \nu_{\text{partial-CB}}$, so higher colour scale values represent regions with increased ν for the full-CBET calculation. These results primarily show that CBET slightly reduced the bangtime equatorial radius, and thus reduced the oblateness, compared to simulations where the spatial redistribution of power due to CBET was neglected.

initial magnetisation increased the coronal plasma anisotropy. Therefore, the difference between the top two rows would have changed as B_{z0} increased, *i.e.* from column to column. Evidently, the top rows appear very similar to each other, which illustrates that anisotropic changes to CBET from magnetised hydrodynamic profiles was a small effect compared to both the polar nature of the drive and the anisotropic drive due to magnetised transport. Small changes are present however, which is far more clearly displayed in Fig. 5.16. This figure plots the difference in n_e and T_e between the full-CBET (top row of 5.15) and partial-CBET (middle row of 5.15) bangtime profiles. As can be seen, differences became more significant at larger magnetisations, demonstrating that anisotropic changes of CBET scattering do impact upon the hydrodynamics. The main difference can be seen at the equator of the n_e profile, of both magnetised simulations, where the full-CBET maximum density occurs at a lower x value. Thus, CBET redistribution of deposition in magnetised coronal profiles marginally reduced the oblateness of the bangtime density. This aligns with the expected behaviour, described at the end of Sec. 5.4.2. Restating here, the early, localised CBET-deficit reduced the strength of the polar shock launch compared to the equatorial shock. This improved the symmetry of the shock and therefore the stagnation state. It was also observed that this effect seemed to become marginally more significant in magnetised coronal profiles. By comparing Fig. 5.16.a with Fig. 5.16.b and 5.16.c, the differences are slightly more significant, which supports this analysis.

The oblateness parameters from Tab. 5.1 are plotted for these implosions as a function of initial field strength for different CBET treatments. These values were obtained by fitting an ellipse to the contour of maximum density at bangtime. The fitting procedure returned asymmetric error bars, which are also plotted. As was observed by examining the profiles in Fig. 5.15 directly, the no-CBET results are all significantly rounder, compared to simulations for which the effect of CBET on deposition magnitude is included. This was due to




Figure 5.17: The oblateness of bangtime density profiles for different magnetisation and CBET effects. All values and errors were obtained by fitting an ellipse (with axes orientated along $\hat{x}$ and $\hat{z}$), to the radius of maximum density. No-CBET simulations were consistently more round, because the initial shock was stronger, and therefore travelled ahead of the pusher material more quickly than the CBET equivalent. Thus, after rebounding off the axis, it met the infalling mass and produced thermonuclear conditions at a larger radius. As was seen in Fig. 5.16, throughout the entire implosion, when the effect of CBET on spatial location of deposition was included, it acted to slightly move energy from the pole to the waist and thus marginally reduced the oblateness of the stagnation state. Note that the dashed lines are not fits, they are simply included to guide the eye.

the faster shock launched with greater early time deposition. The partial-CBET simulations, which do not include redistribution of deposition are all marginally more oblate than the full-CBET simulations, with the caveat this cannot be stated conclusively, because of the error bar magnitudes. Simply stated, CBET acts to slightly reduce the oblateness of the stagnation state by effectively redistributing a small quantity of power from the poles of the capsule to the equator. There is no evidence from this plot, however, that increasing magnetisation, anisotropically alters CBET sufficiently to leave an experimental signature. If, for example, the difference between the full- and partial-CBET curves had diverged (or converged) with increasing magnetisation, this would have suggested that the increasingly anisotropic hydrodynamic profiles from magnetisation, significantly impacted upon the CBET scattering.

In retrospect, however, it was noted that the configuration used for these simulations likely minimised any observable effect in several ways. Firstly, as was discussed at the end of Sec. 5.4.2, the polar beam geometry obscured the less significant CBET anisotropy from magnetisation. Secondly, the shock driven nature of the exploding-pusher configuration, could have minimised experimentally observable signatures. The shape of the stagnation state was predominantly dictated by the shock launch geometry from the shell explosion, which was mostly sensitive to the early laser-energy deposition. A more conventional, hotspot ignition target would have maintained a dense shell throughout the implosion, which would have developed density modulations due to CBET-induced deposition asymmetries, as was seen in the previous chapter, for example in Fig. 4.13. Experimentally, a mode $\ell = 2$ is long

wavelength enough to be inferred from neutron diagnostics, and would thus leave an experimentally observable signature in the neutron spectrum [261, 262]. Finally, higher Z ablator materials such as SiO_2 , exhibit higher coronal plasma temperatures than CH shells, due to increased Inv-Brem efficiency. CBET becomes less significant for direct-drive, as the coronal temperatures increase [152], therefore an experimental setup with a CH outer material may have lead to more CBET scattering overall, and therefore a more easily observable change to hydrodynamic profiles from magnetisation.

5.5 Conclusions

5.5.1 Summary of Work

This chapter described simulations of directly-driven, magnetised exploding-pusher implosions. For the first time, an in-line model for CBET was included in a magnetised direct-drive implosion simulation. Simulations were performed in order to study how asymmetry in coronal plasma profiles from magnetisation may affect the action of CBET. 1-D simulations of the exploding-pusher configuration without magnetisation demonstrated that CBET was dynamically significant to the implosion, reducing the coupled laser energy to the target by $10 \rightarrow 15\%$. The main effect of this reduced absorption was too slow the speed of the strong shock, which led to bangtime occurring at greater target convergence after the rebounding shock compressed against the in-falling material. When extended to 2-D, the simulations demonstrated that the polar nature of the drive was highly significant to the symmetry of the implosion, leading to a bangtime density profile with an aspect ratio $R_{\text{eq.}}/R_{\text{pole}} \sim 3$. The slowing of the shock due to the CBET reduction of absorption, led to bangtime occurring at a larger radius, and thus a less oblate profile was obtained, because the shock travelled mainly along the $\pm \hat{z}$ -axis. Simulations which included the reduced absorption due to CBET, but not spatial redistribution of energy, showed that CBET power redistribution marginally reduced the oblateness of the implosion, consistent with how it acts to reduce $\ell = 1$ asymmetries [247].

Simulations were also conducted which included an initial magnetic field along the $+\hat{z}$ -axis. The key physical processes which were dynamically significant to the implosion dynamics were identified by conducting additional simulations which turned off the effect of each process in turn. Anisotropic thermal conduction was shown to be highly significant. It acted to insulate the deposited energy at the poles of the equator, preventing thermal equilibration in the polar direction and thus led to an increased asymmetry of the drive. The generally high temperatures of the target led to large R_m and plasma β values, therefore the Lorentz force and resistive diffusion had minimal impact on the implosion. Nernst-advection of the magnetic field effectively reoriented the coronal field structure at the low-Hall parameter, equator of the target. This field reorientation altered the anisotropic thermal conduction, leading to a ‘bump’ in the density at the target waist. Although this equatorial-bump had minimal impact on integrated neutron diagnostics, because it was well removed from the polar-regions, which produced the most thermonuclear reactions, it may signifi-

cantly impact yields for symmetrically driven targets.

The impact of the magnetised coronal plasma profiles on CBET scattering of the laser energy was also studied. By examining how CBET reduced absorption instantaneously through both the magnetised and unmagnetised coronas, the dominant effect of CBET was observed to be an early reduction of deposition at the capsule poles, regardless of seed field strength. This led to a marginally more symmetric drive, reducing the oblateness of the bangtime profiles. Increasing the level of the initial seed magnetic field appeared to slightly exaggerate this early, polar reduction of deposition, although the effect was minimal and did not significantly impact upon the hydrodynamics. Therefore, increased anisotropy of the coronal field profiles was not observed to significantly alter CBET scattering for these pole-heavy drive, exploding-pusher targets.

5.5.2 Future Work

Future alterations to the simulation configuration shall now be suggested, which may lead to a more obvious impact of magnetisation on CBET. The simulations conducted for this work all used a 40-beam OMEGA configuration, where the 20 equatorial beams were removed from the drive, due to the presence of a magnetic field coil in experiments. Simulations demonstrated that this was the dominant effect, which dictated the symmetry of the bangtime state. Magnetisation of thermal conduction also played a significant role, increasing the oblateness of the drive. The polar drive largely obscured the much more subtle changes to CBET from the altered hydrodynamic profiles, shown by comparing the top and bottom halves of panels in Fig. 5.14. A configuration with a more symmetric, 60-beam drive, would not have an underlying $\ell = 2$ anisotropy without magnetisation, which may make the changes to CBET in magnetised profiles more pronounced.

Secondly, alterations to the target could be made, which may lead to both increased diagnostic signatures of an anisotropic drive and a larger impact of CBET overall. Firstly, simulations of a more conventional, central hotspot ignition target could be conducted, which maintain a dense shell throughout the implosion. Asymmetry of the deposition seeds density perturbations in the shell throughout the duration of the implosion. This is in contrast to the exploding-pusher targets presented here, where the bangtime symmetry is predominantly dictated by the shape of the early, shock launch. Additionally, using an ablator material with a lower ionisation state would lower the coronal plasma temperatures significantly and thus enhance the levels of CBET [152]. More subtle changes to CBET scattering from magnetisation may thus have a larger impact on the implosion, due to more CBET occurring overall.

6 Conclusions

The work presented in this manuscript details the development, validation and use of the SOLAS laser ray-trace and Cross-Beam Energy Transfer (CBET) model, along with its coupling to the CHIMERA Radiative-Magnetohydrodynamics (Rad-MHD) code. Prior to the development of SOLAS, direct-drive simulations with the CHIMERA code had limited predictive capability. A simple 1-D ray-trace was used, which did not include any mechanisms for loss of coupling. An artificial, ad-hoc multiplier had to be employed to match experimental results, alongside a multi-variate tuning procedure of both the deposition multiplier and electron flux-limiter. The deposition multiplier simultaneously accounted for refractive loss of coupling, loss of coupling due to CBET and an incorrect deposition profile, which increased the fraction of energy deposited close to critical. The extensive validation problems presented in Chap. 3 demonstrate that CHIMERA-SOLAS simulations are now capable of correctly predicting the coupled energy to the target, without the necessity for a tuning procedure. A review of the model capabilities and code validation is first provided, before summaries of the simulations conducted for this thesis are presented. Potential improvements to SOLAS, CHIMERA and the simulations are also suggested. The thesis is concluded with some closing remarks about CBET modelling more generally.

6.1 The SOLAS Ray-Trace and CBET Model

The primary outcome of this thesis was the development of a laser ray-tracing and CBET module for the CHIMERA code. This module was mainly intended for use in direct-drive simulations, but it would also allow for more accurate modelling of other High Energy Density Physics (HEDP) experiments, for example indirect-drive Inertial Confinement Fusion (ICF) and laser-driven gas jet experiments.

6.1.1 Model Development and Validation

As described in Chap. 3, a custom ray-trace grid was implemented, which combines cells from an underlying Eulerian grid, leading to a resolution that is more suitable for laser length scales. The grid resolution is also adaptive, such that areas close to the critical surface can be accurately modelled, where refraction and deposition occur on shorter length scales. The grid also allows for approximately equal area cells in multidimensional spherical simulations, which is crucial for avoiding grid artefacts in the deposition when modelling CBET. An adaptive RK-45 algorithm was employed to evolve the ray trajectories and Inverse-Bremsstrahlung (Inv-Brem) deposition. This efficiently integrates the path of rays, while maintaining an acceptable degree of accuracy. Validation problems were presented which compared the solver to analytic solutions. The SOLAS results were all in excellent agreement with the problem solutions.

Accurate CBET modelling, particularly for direct-drive simulations, depends on an accurate and low-noise field reconstruction algorithm. Methods which rely on ray statistics to obtain the intensity were tested, and although not presented in the thesis, they were found to lead to significant inaccuracies and spurious laser-grid artefacts in the deposition. Therefore, a ‘neighbour-ray’ algorithm was implemented following Ref. [2], in which each ‘main-ray’ is co-traced alongside a triangle of neighbour-rays. This sweeps out an area of the main-ray along its phase-front. The area provides an estimate for the ray amplitude, and therefore field strength, or equivalently intensity [263]. An estimate for the ray amplitude also allows for a proper treatment of the field in the vicinity of caustics, where the WKB assumption used for ray-tracing breaks down, and therefore ray-based field reconstruction diverges. Field Limiter (FL) and Etalon Integral (EI) algorithms were implemented to limit the ray amplitude and hence field strength in the vicinity of caustics, which is necessary for accurate CBET modelling in direct-drive. A memory efficient, domain-balanced grid approach with dynamic memory structure was implemented to store the laser fields from many beams. Numerous field reconstruction test problems were conducted, which compared the SOLAS fields to results from a wave-based code. These results demonstrated that the solver correctly reproduced the field in direct-drive like plasma profiles. It was thus deemed suitable for inclusion in a ray-based CBET model.

The CBET algorithm implemented in CHIMERA saturates via pump depletion, therefore the field of the lasers is solved for iteratively, while accounting for CBET in the updating fields. The linear fluid and kinetic CBET gains were both implemented. The kinetic formulation enables accurate computation of the gain for multiple ion species plasmas, which

are commonly used for the ablator of direct-drive targets. A memory efficient, caching of ray-trajectories and CBET gains was implemented via a linked list, to avoid excessive re-computation costs for the pump-depletion iterations. A Caustic Gain Truncation (CGT) algorithm [173] was implemented to improve the accuracy and energy conservation of the model at the reduced resolutions necessary for full, 3-D simulations. Additional energy conservation iterations were also implemented, which artificially modified the gains in the caustic region to enforce energy conservation. Although this is an ad-hoc approach, it was observed to improve the accuracy of results compared to wave-based solvers. Numerous CBET test problems were conducted, both with and without laser caustics. The SOLAS results were found to compare favourably to both wave- and ray-based models.

Finally, the model was employed, integrated with CHIMERA, to model OMEGA direct-drive implosions. The OMEGA shot 89224 was simulated both with and without CBET [203]. The results demonstrated that CBET significantly reduced the deposited power and hence the bangtime of the implosion. The bangtime, which is effectively a proxy for the energy coupled to the target, was in excellent agreement with the LILAC result for the same initial conditions. A 3-D post-process of spherically symmetric LILAC hydrodynamic data was also conducted [67]. The deposition asymmetry from SOLAS both with and without CBET were in excellent agreement with results from the IFRIIT and BEAMCROSSER codes [135]. These results demonstrate that CHIMERA-SOLAS simulations are capable of accurately modelling direct-drive experiments, including the effect of CBET on absorption-reduction and asymmetry.

6.1.2 Potential SOLAS Enhancements

Other than CHIMERA-SOLAS, only a single code combination, ASTER-IFRIIT, is capable of accurately modelling CBET for direct-drive in 3-D. The effect of CBET on stagnation asymmetry is an inherently 3-D problem, due to the lack of rotational symmetry about a single axis on the OMEGA laser system. Therefore, even without further improvements, the model could be extensively utilised to study 3-D CBET physics and how it impacts direct-drive implosions on both OMEGA and the National Ignition Facility (NIF). Numerous improvements could be made to the model to make it capable of accurately modelling a wider array of ICF physics.

An initial improvement would be to include the Langdon effect on both absorption [126] and CBET [138]. Langdon is the name given to the preferential heating of the colder electron population in the distribution function by Inv-Brem [124]. If the heating rate is faster than the thermalising collisions, then this distorts the distribution function away from Maxwellian, reducing the fraction of colder electrons, and therefore reducing absorption. This is a significant effect for direct-drive ICF implosions, for which it has been shown to reduce absorption rates by $\mathcal{O}(10\%)$ [126]. This was not implemented, mainly because the focus of the work was to understand how CBET affected implosions rather than Langdon. It should now be straightforward to include the absorption reduction factor, which relies mainly on an accurate electric field reconstruction. Additionally, Langdon distribution functions alter the

susceptibility of the plasma. This changes the Ion Acoustic Wave (IAW) dispersion relation, hence CBET gains are also affected. This can also be accounted for in a ray-based model by modifying the gain appropriately, which could also be easily implemented in the model as it stands [138]. One interesting area of investigation with an implementation of the Langdon reduction of CBET, would be to investigate whether this is sufficient to remove the need for a gain clamp in indirect-drive CBET models [44]. Intensities are large in the Laser Entrance Hole (LEH) of indirect-drive configurations, which could lead to significant gain reductions, and therefore make CBET modelling more predictive for these experiments. Indirect-drive simulations have been conducted using SOLAS [264], although these required gain limiters for the pump-depletion iterations to converge.

Additionally, a model for polarised CBET could be implemented into SOLAS. The ponderomotive drift of particles from which CBET grows only occurs due to the interaction of parallel polarised beams. While the OMEGA lasers are approximately randomly polarised, which is simple to account for in CBET models, polarisation smoothing optics actually split each beam into two, spatially offset, orthogonally polarised sub-beams. The CBET interaction of these sub-beams leads to a persistent $\ell = 1$ asymmetry on OMEGA [136, 132, 133]. By accounting for the polarisation of the light in the CBET model, this $\ell = 1$ could be recreated and studied computationally with CHIMERA-SOLAS simulations.

A further suggestion for an improvement to SOLAS would be to include the capability to model laser bandwidth in CBET interactions [153]. Broadband lasers are assumed to be the best route on future facilities to eliminating CBET. The CBET interaction is time-dependent however, and therefore understanding the bandwidth required to mitigate it requires inline modelling of the interaction [152]. Ray-based implementations of CBET models which account for bandwidth exist and could be included in SOLAS.

A final suggestion would be to include models for additional Laser-Plasma Instabilities (LPIs). On OMEGA, CBET is by far the most significant instability in terms of energy loss, due to the high field strength of the daughter wave. However, on larger facilities with longer scale length plasmas, other convective LPIs would likely become dynamically significant, particularly Stimulated Raman Scattering (SRS) [69]. Studies have shown that SRS may be less sensitive to laser bandwidth [265] and therefore accurate modelling of this effect is crucial for the design of next generation direct-drive facilities. A significant challenge to modelling all LPIs is an accurate field reconstruction, which now exists in SOLAS, and therefore it should be possible to include reduced models for these interactions [170].

6.1.3 Potential CHIMERA Enhancements

Further enhancements could be made to the CHIMERA code to improve the predictive capability of direct-drive simulations. For instance, the CHIMERA Monte-Carlo alpha transport algorithm could be adapted to model hot-electron generation and pre-heat of the fuel [96]. Monte-Carlo models are the gold standard of Radiative-Hydrodynamics (Rad-Hydro) hot-electrons models [80, 266, 267], and much of the requisite code infrastructure already exists in CHIMERA. Hot-electrons are generated via the SRS instability, and therefore may well be

increasingly significant at larger scale facilities, due to the increasing plasma scale length and the potential inefficacy of broadband light to mitigate SRS [265]. Therefore, a hot-electron modelling capability could be highly useful for designing these next generation facilities.

Finally, laser direct-drive targets exhibit steep temperature and density gradients in the conduction zone. Therefore, the heatflow, which transports laser energy to the ablation surface, is not well modelled by a local, Spitzer conductivity [109]. Treatments for non-local heatflow could be implemented into CHIMERA, for example SNB [110, 111, 112], FAST-VFP [114] or RKM [115]. These models would more accurately predict the thermal transport and therefore drive and preheat from the conduction zone.

6.2 Effect of Beam Radius on Beam-Mode Asymmetry

A summary of the work conducted in Chap. 4 is presented in this section, along with some suggestions for further work. In this chapter, a series of simulations was conducted to attempt to understand the effect of the beam radius initial condition on the performance of direct-drive targets on OMEGA. In particular, statistical modelling of OMEGA implosions has demonstrated that the scaling of performance with respect to beam radius is not well understood [72, 73]. It was hypothesised that the varying effect of CBET at different beam radii may explain the complicated scaling relation observed in the data.

6.2.1 Summary of Simulations

In order to study the effect of CBET as a function of beam-radius, an ensemble of ~ 10 simulations was required. Additionally, because the energetics of an implosion decrease with increasing CBET (less energy is coupled and therefore less mass can be imploded), each simulation in the ensemble required $\mathcal{O}(10)$ tuning simulations to obtain the initial conditions. While the true effect of CBET is 3-D, performing $\mathcal{O}(100)$ 3-D CBET simulations was well beyond the maximum computational resources available for this project. A 2-D, cylindrical direct-drive simulation platform was therefore designed, which was able to qualitatively reproduce the physics of a spherical implosion. Because the simulation geometry was planar, a 2-D ray-trace fully described the laser propagation through the system. This reduced the number of rays required by $\mathcal{O}(100)$, compared to a fully 3-D simulation. The platform thus enabled the impact of CBET on stagnation to be studied as a function of beam radius, via a simulation ensemble.

A 1-D tuning procedure was initially conducted to obtain the required laser power to drive hydrodynamically similar implosions at all beam radii. Integrated implosion metrics indicated that the 1-D simulation results at all beam radii were sufficiently similar to each other. Therefore, the performance degradation as a function of beam radius, due to *only* multi-dimensional effects could be studied with 2-D simulations. These 2-D cylindrical simulations indicated that in the absence of CBET, symmetry of the stagnation state improved with increasing beam radius. When CBET was included, ‘modal-flips’ of the deposition were observed in time, due to CBET acting non-uniformly around the target azimuth. Broadly,

the symmetry of the stagnation state was observed to decrease with increasing beam radius when CBET was accounted for. However, the trend obtained was non-monotonic, due to the modal flips instantaneously improving deposition symmetry, depending on the beam radius and target convergence. It was hypothesised that this non-monotonic trend (if observed in 3-D spherical simulations as well) could be the origin of complicated statistical scaling with respect to beam radius on OMEGA.

6.2.2 Suggestions for Further Work

Two main improvements were identified to make the results from this chapter more relevant to the OMEGA statistical model. The first would be to use a spherical geometry for the simulations. The cylindrical convergence of the target and divergence of the coronal plasma is not directly applicable to spherical targets, even if qualitatively similar. However, it was only computationally possible to create such a large ensemble of simulations for this chapter by conducting a 2-D ray-trace (*i.e.* the rays move in a plane, not in a full 3-D volume), as this reduced the number of rays required by orders of magnitude. CBET is typically by far the most expensive effect to model in direct-drive calculations, so this reduction in cost enabled a large number of simulations to be conducted. Simulating CBET accurately in a 2-D r - z spherical Rad-Hydro simulation requires the full azimuthal extent to be included in the laser model. Therefore, an expensive 3-D ray-trace is required to compute CBET. Additionally, the power deposition in these simulations must be azimuthally averaged around the target, which artificially reduces the amplitude of the beam-mode perturbations. This means that the true performance scaling with beam radius could not be accurately captured with 2-D r - z simulations. However, a 2-D ‘planar’ simulation could provide the correct spherical convergence properties, while also maintaining the performance benefit of a 2-D ray-trace. This platform would be in r - ϕ space and have a polar angle extent $\theta \in [\lesssim \pi/2, \gtrsim \pi/2]$. A planar, 2-D raytrace could therefore be performed,¹ which would maintain the computational cost advantages over 3-D simulations.

A further enhancement would be to alter the 1-D tuning procedure to ensure that the results are more directly applicable to the OMEGA statistical model. In Chap. 4, the targets were kept hydrodynamically similar, in the sense that the energy coupled to the target was the same. This was achieved by increasing the incident laser power, such that the same energy was absorbed for each simulation, despite the larger beam radius simulations experiencing more CBET due to increased beam overlap. This is contrary to what happens on OMEGA, where the laser energy is kept constant, and the target is altered (decreasing radius and mass) to create larger beam radius to target radius ratios. The trends observed from this simulation campaign are therefore not directly comparable to the statistical model and only qualitative comparisons can be drawn. By performing a more in-depth target optimisation procedure that does not vary the laser pulse but rather the target itself,² then the trends obtained could

¹Note that the deposition would have to be altered as a function of r , to account for varying cell volume with convergence.

²Note that this procedure was not followed for the work in Chap. 4, because it is a higher dimensional optimisation procedure and therefore requires many more simulations.

be compared directly to the statistical model.

6.3 The Effect of Coronal Magnetisation on CBET Scattering

A summary of the work conducted in Chap. 5 is presented in this section, along with some suggestions for further work. Magnetisation of direct-drive targets are known to enhance the $\ell = 2$ stagnation asymmetry of pole-heavy-drive magnetised exploding pusher targets [243]. This asymmetry arises from anisotropisation of the heatflow, and therefore also leads to asymmetry in the magnetised, coronal plasma, where the lasers propagate. CBET is known to be sensitive to global, $\ell = 1$ coronal asymmetries [67]. Therefore, it was hypothesised that the asymmetry of the corona due to magnetisation could anisotropically affect CBET scattering.

6.3.1 Summary of Simulations

A series of simulations was performed with varying strengths of seed magnetic field and treatments of CBET, to assess both the role of magnetisation on exploding pusher implosions, and how the platform is affected by CBET. Initially, unmagnetised 1-D calculations were performed, which demonstrated that CBET is dynamically significant for exploding pushers, reducing absorption by $\sim 15\%$ and bangtime by $\sim 10\%$. It is worth noting that this is less significant than for conventional hotspot ignition OMEGA designs ($\gtrsim 20\%$ reduction). This is mainly due to the SiO_2 ablator material resulting in higher coronal plasma temperatures, which results in reduced gains. Unmagnetised 2-D simulations demonstrated that the polar-drive configuration led to significant asymmetry of the hydrodynamic profiles at bangtime.

Magnetised 2-D simulations without CBET were conducted to assess the role of various extended-Magnetohydrodynamics (MHD) terms on the implosion dynamics. Anisotropic thermal conduction was observed to be the most dynamically significant. The heatflow was highly anisotropic in the large Hall parameter corona, which enhanced the $\ell = 2$ asymmetry from the drive geometry by preventing equilibration of temperature around the target. High plasma β and magnetic Reynolds numbers meant that the Lorentz force and resistive diffusion, respectively, were relatively unimportant. While it did not significantly impact the integrated implosion metrics, the Nernst effect significantly altered the field structure at the low Hall parameter equator of the target, which led to an altered density profile in this region. Due to the global $\ell = 2$ asymmetry however, this alteration was substantially separated from the hottest plasma regions which produced the yield, and therefore limited effect was observed in the integrated neutron metrics.

Simulations were conducted in 2-D with varying initial field strengths and CBET treatments to understand how CBET was affected by magnetised coronal profiles. The main comparison was between full-CBET simulations and partial-CBET simulations, where only the reduction of deposited power due to CBET (not the redistribution of deposition) was fed into the hydrodynamics. The computational diagnostics from SOLAS, which can be used to un-

derstand how CBET affects the deposition asymmetry, demonstrated some minor changes to the scattering of light due to CBET. Particularly, higher magnetisation appeared to marginally reduce the deposition asymmetry from CBET, especially early in the implosion. However, comparison of the bangtime profiles did not demonstrate significant difference between the full- and partial-CBET simulations for any level of seed magnetisation. Therefore, it was concluded that CBET was not significantly affected by magnetisation for these direct-drive exploding pusher targets.

6.3.2 Suggestions for Further Work

The first avenue for further work should be to properly compare the computational results to the experiments, by using synthetic x-ray diagnostics to obtain comparable images of the bangtime asymmetry. This was not conducted in this thesis, as the focus of the work was to computationally investigate the action of CBET in these experiments.

Several suggestions were provided to design an experiment, which may observe a more significant impact on CBET scattering due to magnetisation. Firstly, a symmetric laser drive could be employed. The pole-heavy-drive which was used in these simulations dominated the symmetry evolution of the targets. It therefore obscured the much smaller changes to implosion morphology, which were a result of magnetised coronas altering CBET scattering. By utilising a symmetric drive, the zeroth order polar-drive asymmetry would be eliminated, such that the CBET changes to deposition may appear more obvious in both simulations and experiments. Additionally, a more conventional hotspot ignition style implosion could be utilised with a CH ablator. A lower Z ablator implosion would exhibit lower coronal plasma temperatures, resulting in more CBET scattering. This would make the changes to deposition symmetry from magnetised CBET more evident.

6.4 Final Remarks on CBET Modelling

The work in this manuscript has outlined the development, validation and use of the SOLAS ray-trace and CBET model. IFRIIT is the only other direct-drive suitable, 3-D CBET model, which has been integrated into a 3-D Rad-Hydro code (ASTER). The lack of computational models capable of studying the inherently 3-D physics of CBET mean that there is ample possibility for further computational study with this code. It is worth noting however, that CBET modelling is very expensive and dominates the computational runtimes of 1-, 2- and 3-D CHIMERA-SOLAS simulations. This is also true of ASTER-IFRIIT simulations [67]. Memory overheads are also very large, due to the necessity of storing field information for many beams over an entire 3-D computational grid. The large memory requirements make effective load balancing across processors difficult. It is also hard to envisage how current multidimensional CBET models could be effectively refactored for Graphics Processing Unit (GPU) acceleration, due to the limited Random-Access Memory (RAM) available on current units. It is likely that a future code, which is capable of harnessing the vast performance improvements available from a GPU, would need to utilise a much more memory efficient method

to store the field.

It is hoped that CBET will become a problem of the past on next generation laser facilities, due to the development of broad bandwidth lasers. There are still, however, many experiments which have and shall in future be performed on current, low bandwidth facilities. Understanding the performance enhancement which can be expected from broadband lasers, and hence how to design these future facilities, requires a complete understanding of how LPIs degrade current implosions. One interesting example is that CBET acts to mitigate $\ell = 1$ asymmetries, which arise from target offsets and mispointing of the lasers [247]. Mitigation of CBET is therefore likely to enhance sensitivity to these highly damaging perturbations, therefore enhanced engineering capabilities will be required to reduce these random errors. CBET models may also be adapted to include additional LPIs, which are less of a concern for current direct-drive implosions, but may increase in significance as the laser energy and plasma scale lengths increase. In short, many problems remain to be solved and there is ample area for investigation in the inline modelling of CBET and other LPIs for ICF.

Appendices

A List of Acronyms

ADI Alternating-Direction Implicit

CBET Cross-Beam Energy Transfer

CEA Commissariat à l'Énergie Atomique et aux Energies Alternatives

CIFS Centre for Inertial Fusion Studies

CGT Caustic Gain Truncation

DFT Discrete Fourier Transform

DPRs Distributed Polarization Rotators

DPSSL Diode-Pumped Solid State Laser

EI Etalon Integral

EMW Electro-Magnetic Wave

EoS Equation of State

EPW Electron Plasma Wave

FL Field Limiter

FRT Forward Ray Tracing

GO Geometric Optics

GPU Graphics Processing Unit

HDC High Density Carbon

HEDP High Energy Density Physics

IAW Ion Acoustic Wave

ICF Inertial Confinement Fusion

IFAR In-Flight Aspect Ratio

IFE Inertial Fusion Energy

Inv-Brem Inverse-Bremsstrahlung

IRT Inverse Ray Tracing

LCOE Levelised Cost of Electricity

LEH Laser Entrance Hole

LLE Laboratory for Laser Energetics

LLNL Lawrence Livermore National Laboratory

LMJ Laser Mégajoule

LPIs Laser-Plasma Instabilities

MCF Magnetic Confinement Fusion

MHD Magnetohydrodynamics

MIT Massachusetts Institute of Technology

MPI Message Passing Interface

NIF National Ignition Facility

OpenMP Open Multi-Processing

PCGO Paraxial Complex Geometric Optics

PiC Particle in Cell

Rad-Hydro Radiative-Hydrodynamics

Rad-MHD Radiative-Magnetohydrodynamics

RAM Random-Access Memory

RTI Rayleigh–Taylor Instability

SBS Stimulated Brillouin Scattering

SRS Stimulated Raman Scattering

SSD Smoothing by Spectral Dispersion

TPD Two Plasmon Decay

VFP Vlasov-Fokker-Planck

B Copyright Permissions

Several figures under copyright have been reproduced with permission in this thesis. Tab. B.1 lists the relevant copyright information and associated Licence numbers. Fig. 1.6 was reproduced under the Creative Commons (CC BY 4.0) Licence and is also listed in Tab. B.1.

Table B.1: Copyright permissions for material reproduced in this thesis.

Fig. #	Ref. #	©	Licence
1.6	[47]	-	CC BY 4.0
2.4	[116]	Springer Nature	5874790453624
2.5	[116]	Springer Nature	5874790453624
2.9	[131]	AIP Publishing	5875480620949
4.2	[142]	American Physical Society	RNP/24/AUG/081941
4.3	[73]	AIP Publishing	5840820879747
5.2	[237]	AIP Publishing	5847050626966
5.3	[241]	American Physical Society	RNP/24/AUG/081941
5.4	[243]	American Physical Society	RNP/24/AUG/082231

Bibliography

- [1] J P Chittenden et al. “X-Ray Generation Mechanisms in Three-Dimensional Simulations of Wire Array Z-pinches”. In: *Plasma Physics and Controlled Fusion* 46.12B (Dec. 2004), B457–B476. ISSN: 0741-3335, 1361-6587. DOI: 10.1088/0741-3335/46/12B/039.
- [2] R. K. Follett et al. “Validation of Ray-Based Cross-Beam Energy Transfer Models”. In: *Physics of Plasmas* 29.11 (Nov. 2022), p. 113902. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/5.0123462.
- [3] Ember Institute. *Statistical Review of World Energy*. May 2024. URL: <https://ember-climate.org/app/uploads/2024/05/Report-Global-Electricity-Review-2024.pdf>.
- [4] Thomas Griffiths et al. “The Commercialisation of Fusion for the Energy Market: A Review of Socio-Economic Studies”. In: *Progress in Energy* 4.4 (Oct. 2022), p. 042008. ISSN: 2516-1083. DOI: 10.1088/2516-1083/ac84bf.
- [5] Stefano Atzeni and Jürgen Meyer-ter-Vehn. *The Physics of Inertial Fusion: Beam Plasma Interaction, Hydrodynamics, Hot Dense Matter*. Oxford Science Publications 125. Oxford : New York: Clarendon Press ; Oxford University Press, 2004. ISBN: 978-0-19-856264-1.
- [6] Todd H. Rider. “A General Critique of Inertial-Electrostatic Confinement Fusion Systems”. In: *Physics of Plasmas* 2.6 (June 1995), pp. 1853–1872. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/1.871273.
- [7] M.B. Chadwick et al. “ENDF/B-VII.1 Nuclear Data for Science and Technology: Cross Sections, Covariances, Fission Product Yields and Decay Data”. In: *Nuclear Data Sheets* 112.12 (Dec. 2011), pp. 2887–2996. ISSN: 00903752. DOI: 10.1016/j.nds.2011.11.002.
- [8] Lowell S. Brown and Gerald M. Hale. “Field Theory of the $d + t \rightarrow n + \alpha$ Reaction Dominated by a ${}^5\text{He}^*$ Unstable Particle”. In: *Physical Review C* 89.1 (Jan. 2014), p. 014622. ISSN: 0556-2813, 1089-490X. DOI: 10.1103/PhysRevC.89.014622.
- [9] W. Daughton et al. “Influence of Mass Ablation on Ignition and Burn Propagation in Layered Fusion Capsules”. In: *Physics of Plasmas* 30.1 (Jan. 2023), p. 012704. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/5.0129561.
- [10] 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.

- [11] 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.
 - [12] A. L. Schawlow and C. H. Townes. "Infrared and Optical Masers". In: *Physical Review* 112.6 (Dec. 1958), pp. 1940–1949. ISSN: 0031-899X. DOI: 10.1103/PhysRev.112.1940.
 - [13] T. H. Maiman. "Stimulated Optical Radiation in Ruby". In: *Nature* 187.4736 (Aug. 1960), pp. 493–494. ISSN: 0028-0836, 1476-4687. DOI: 10.1038/187493a0.
 - [14] J. H. Nuckolls Llnl. *Early Steps toward Inertial Fusion Energy (IFE) (1952 to 1962)*. Tech. rep. UCRL-ID–131075, 658936. June 1998, UCRL-ID–131075, 658936. DOI: 10.2172/658936.
 - [15] G. S. Fraley et al. "Thermonuclear Burn Characteristics of Compressed Deuterium-Tritium Microspheres". In: *The Physics of Fluids* 17.2 (Feb. 1974), pp. 474–489. ISSN: 0031-9171. DOI: 10.1063/1.1694739.
 - [16] Aidan Crilly. "Simulation of Nuclear Observables in Inertial Confinement Fusion Experiments". PhD thesis. London: Imperial College London, Apr. 2020. URL: <https://inspirehep.net/literature/1850910> (visited on 09/05/2024).
 - [17] M. Lafon et al. "Direct-Drive-Ignition Designs with Mid- Z Ablators". In: *Physics of Plasmas* 22.3 (Mar. 2015), p. 032703. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/1.4914835.
 - [18] J. L. Kline et al. "First Beryllium Capsule Implosions on the National Ignition Facility". In: *Physics of Plasmas* 23.5 (May 2016), p. 056310. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/1.4948277.
 - [19] S. X. Hu et al. "Laser-Direct-Drive Fusion Target Design with a High- Z Gradient-Density Pusher Shell". In: *Physical Review E* 108.3 (Sept. 2023), p. 035209. ISSN: 2470-0045, 2470-0053. DOI: 10.1103/PhysRevE.108.035209.
 - [20] D. T. Casey et al. "Performance and Mix Measurements of Indirect Drive Cu-Doped Be Implosions". In: *Physical Review Letters* 114.20 (May 2015), p. 205002. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.114.205002.
 - [21] V. N. Goncharov et al. "Improved Performance of Direct-Drive Inertial Confinement Fusion Target Designs with Adiabatic Shaping Using an Intensity Picket". In: *Physics of Plasmas* 10.5 (May 2003), pp. 1906–1918. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/1.1562166.
 - [22] O. A. Hurricane et al. "The High-Foot Implosion Campaign on the National Ignition Facility". In: *Physics of Plasmas* 21.5 (May 2014), p. 056314. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/1.4874330.
 - [23] M. Gatu Johnson et al. "Impact of Asymmetries on Fuel Performance in Inertial Confinement Fusion". In: *Physical Review E* 98.5 (Nov. 2018), p. 051201. ISSN: 2470-0045, 2470-0053. DOI: 10.1103/PhysRevE.98.051201.
-

- [24] V A Smalyuk et al. “Review of Hydrodynamic Instability Experiments in Inertially Confined Fusion Implosions on National Ignition Facility”. In: *Plasma Physics and Controlled Fusion* 62.1 (Jan. 2020), p. 014007. ISSN: 0741-3335, 1361-6587. DOI: 10 . 1088/1361-6587/ab49f4.
 - [25] 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.
 - [26] L. J. Perkins et al. “Shock Ignition: A New Approach to High Gain Inertial Confinement Fusion on the National Ignition Facility”. In: *Physical Review Letters* 103.4 (July 2009), p. 045004. ISSN: 0031-9007, 1079-7114. DOI: 10 . 1103/PhysRevLett . 103 . 045004.
 - [27] 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.
 - [28] Max Tabak et al. “Ignition and High Gain with Ultrapowerful Lasers*”. In: *Physics of Plasmas* 1.5 (May 1994), pp. 1626–1634. ISSN: 1070-664X, 1089-7674. DOI: 10 . 1063/ 1 . 870664.
 - [29] L. C. Jarrott et al. “Visualizing Fast Electron Energy Transport into Laser-Compressed High-Density Fast-Ignition Targets”. In: *Nature Physics* 12.5 (May 2016), pp. 499–504. ISSN: 1745-2473, 1745-2481. DOI: 10 . 1038/nphys3614.
 - [30] T. Gong et al. “Direct Observation of Imploded Core Heating via Fast Electrons with Super-Penetration Scheme”. In: *Nature Communications* 10.1 (Dec. 2019), p. 5614. ISSN: 2041-1723. DOI: 10 . 1038/s41467-019-13574-8.
 - [31] Rasol Khoda-Bakhsh et al. “Advanced Fusion Fuel for Inertial Confinement Fusion”. In: *Fusion Technology* 22.1 (Aug. 1992), pp. 50–55. ISSN: 0748-1896. DOI: 10 . 13182/ FST92-A30053.
 - [32] Kim Molvig et al. “Low Fuel Convergence Path to Direct-Drive Fusion Ignition”. In: *Physical Review Letters* 116.25 (June 2016), p. 255003. ISSN: 0031-9007, 1079-7114. DOI: 10 . 1103/PhysRevLett . 116 . 255003.
 - [33] Saumyabrata Banerjee et al. “DiPOLE: A 10 J, 10 Hz Cryogenic Gas Cooled Multi-Slab Nanosecond Yb:YAG Laser”. In: *Optics Express* 23.15 (July 2015), p. 19542. ISSN: 1094-4087. DOI: 10 . 1364/OE . 23 . 019542.
 - [34] Dimitri Batani et al. “Future for Inertial-Fusion Energy in Europe: A Roadmap”. In: *High Power Laser Science and Engineering* 11 (2023), e83. ISSN: 2095-4719, 2052-3289. DOI: 10 . 1017/hpl . 2023 . 80.
 - [35] V. N. Goncharov et al. “Novel Hot-Spot Ignition Designs for Inertial Confinement Fusion with Liquid-Deuterium-Tritium Spheres”. In: *Physical Review Letters* 125.6 (Aug. 2020), p. 065001. ISSN: 0031-9007, 1079-7114. DOI: 10 . 1103 / PhysRevLett . 125 . 065001.
-

- [36] I. V. Igumenshchev et al. “Proof-of-Principle Experiment on the Dynamic Shell Formation for Inertial Confinement Fusion”. In: *Physical Review Letters* 131.1 (July 2023), p. 015102. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.131.015102.
 - [37] N. Metzler and J. Meyer-Ter-Vehn. “Target Study for Heavy Ion Beam Fusion”. In: *Laser and Particle Beams* 2.1 (Feb. 1984), pp. 27–48. ISSN: 0263-0346, 1469-803X. DOI: 10.1017/S0263034600000604.
 - [38] John Lindl. “Development of the Indirect-Drive Approach to Inertial Confinement Fusion and the Target Physics Basis for Ignition and Gain”. In: *Physics of Plasmas* 2.11 (Nov. 1995), pp. 3933–4024. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/1.871025.
 - [39] George H. Miller. “The National Ignition Facility”. In: *Optical Engineering* 43.12 (Dec. 2004), p. 2841. ISSN: 0091-3286. DOI: 10.1117/1.1814767.
 - [40] M. L. Spaeth et al. “Description of the NIF Laser”. In: *Fusion Science and Technology* 69.1 (Feb. 2016), pp. 25–145. ISSN: 1536-1055, 1943-7641. DOI: 10.13182/FST15-144.
 - [41] Noël Fleurot et al. “The Laser Mégajoule (LMJ) Project Dedicated to Inertial Confinement Fusion: Development and Construction Status”. In: *Fusion Engineering and Design* 74.1-4 (Nov. 2005), pp. 147–154. ISSN: 09203796. DOI: 10.1016/j.fusengdes.2005.06.251.
 - [42] D. S. Clark et al. “Three-Dimensional Simulations of Low Foot and High Foot Implosion Experiments on the National Ignition Facility”. In: *Physics of Plasmas* 23.5 (May 2016), p. 056302. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/1.4943527.
 - [43] B. J. MacGowan et al. “Laser–Plasma Interactions in Ignition-Scale Hohlraum Plasmas”. In: *Physics of Plasmas* 3.5 (May 1996), pp. 2029–2040. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/1.872000.
 - [44] P. Michel et al. “Tuning the Implosion Symmetry of ICF Targets via Controlled Crossed-Beam Energy Transfer”. In: *Physical Review Letters* 102.2 (Jan. 2009), p. 025004. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.102.025004.
 - [45] J. D. Moody et al. “Multistep Redirection by Cross-Beam Power Transfer of Ultrahigh-Power Lasers in a Plasma”. In: *Nature Physics* 8.4 (Apr. 2012), pp. 344–349. ISSN: 1745-2473, 1745-2481. DOI: 10.1038/nphys2239.
 - [46] A. L. Kritcher et al. “Energy Transfer between Lasers in Low-Gas-Fill-Density Hohlraums”. In: *Physical Review E* 98.5 (Nov. 2018), p. 053206. ISSN: 2470-0045, 2470-0053. DOI: 10.1103/PhysRevE.98.053206.
 - [47] 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.
 - [48] O. A. Hurricane et al. “Physics Principles of Inertial Confinement Fusion and U.S. Program Overview”. In: *Reviews of Modern Physics* 95.2 (June 2023), p. 025005. ISSN: 0034-6861, 1539-0756. DOI: 10.1103/RevModPhys.95.025005.
-

- [49] John Lindl et al. “Review of the National Ignition Campaign 2009-2012”. In: *Physics of Plasmas* 21.2 (Feb. 2014), p. 020501. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/1.4865400.
 - [50] A. J. MacKinnon et al. “High-Density Carbon Ablator Experiments on the National Ignition Facility”. In: *Physics of Plasmas* 21.5 (May 2014), p. 056318. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/1.4876611.
 - [51] A. B. Zylstra et al. “Record Energetics for an Inertial Fusion Implosion at NIF”. In: *Physical Review Letters* 126.2 (Jan. 2021), p. 025001. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.126.025001.
 - [52] A. L. Kritcher et al. “Design of the First Fusion Experiment to Achieve Target Energy Gain $G > 1$ ”. In: *Physical Review E* 109.2 (Feb. 2024), p. 025204. ISSN: 2470-0045, 2470-0053. DOI: 10.1103/PhysRevE.109.025204.
 - [53] A. B. Zylstra et al. “Burning Plasma Achieved in Inertial Fusion”. In: *Nature* 601.7894 (Jan. 2022), pp. 542–548. ISSN: 0028-0836, 1476-4687. DOI: 10.1038/s41586-021-04281-w.
 - [54] A. L. Kritcher et al. “Design of Inertial Fusion Implosions Reaching the Burning Plasma Regime”. In: *Nature Physics* 18.3 (Mar. 2022), pp. 251–258. ISSN: 1745-2473, 1745-2481. DOI: 10.1038/s41567-021-01485-9.
 - [55] 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.
 - [56] *NIF Sets Power and Energy Records | National Ignition Facility & Photon Science*. URL: <https://lasers.llnl.gov/about/keys-to-success/nif-sets-power-energy-records> (visited on 09/06/2024).
 - [57] A. Simon. *LLE Review Quarterly Report (April-June 1989). Volume 39*. Tech. rep. DOE/DP/40200–95, 1191446. June 1989, DOE/DP/40200–95, 1191446. DOI: 10.2172/1191446.
 - [58] T. R. Boehly et al. “The Upgrade to the OMEGA Laser System”. In: *Review of Scientific Instruments* 66.1 (Jan. 1995), pp. 508–510. ISSN: 0034-6748, 1089-7623. DOI: 10.1063/1.1146333.
 - [59] M. Gatu Johnson et al. “Impact of Stalk on Directly Driven Inertial Confinement Fusion Implosions”. In: *Physics of Plasmas* 27.3 (Mar. 2020), p. 032704. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/1.5141607.
 - [60] Stephen E. Bodner. “Critical Elements of High Gain Laser Fusion”. In: *Journal of Fusion Energy* 1.3 (July 1981), pp. 221–240. ISSN: 0164-0313, 1572-9591. DOI: 10.1007/BF01050355.
 - [61] Y. Kato et al. “Random Phasing of High-Power Lasers for Uniform Target Acceleration and Plasma-Instability Suppression”. In: *Physical Review Letters* 53.11 (Sept. 1984), pp. 1057–1060. ISSN: 0031-9007. DOI: 10.1103/PhysRevLett.53.1057.
-

- [62] S. Skupsky et al. "Improved Laser-Beam Uniformity Using the Angular Dispersion of Frequency-Modulated Light". In: *Journal of Applied Physics* 66.8 (Oct. 1989), pp. 3456–3462. ISSN: 0021-8979, 1089-7550. DOI: 10.1063/1.344101.
 - [63] M. Hohenberger et al. "Optical Smoothing of Laser Imprinting in Planar-Target Experiments on OMEGA EP Using Multi-FM 1-D Smoothing by Spectral Dispersion". In: *Physics of Plasmas* 23.9 (Sept. 2016), p. 092702. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/1.4962185.
 - [64] K. Tsubakimoto et al. "Suppression of Interference Speckles Produced by a Random Phase Plate, Using a Polarization Control Plate". In: *Optics Communications* 91.1-2 (July 1992), pp. 9–12. ISSN: 00304018. DOI: 10.1016/0030-4018(92)90091-5.
 - [65] T. R. Boehly et al. "Reduction of Laser Imprinting Using Polarization Smoothing on a Solid-State Fusion Laser". In: *Journal of Applied Physics* 85.7 (Apr. 1999), pp. 3444–3447. ISSN: 0021-8979, 1089-7550. DOI: 10.1063/1.369702.
 - [66] J. Huba D. *NRL Plasma Formulary*. Tech. rep. NRL/PU/6790–13-589. Washington, DC: Naval Research Laboratory, 2013.
 - [67] A. Colaïtis et al. "Inverse Ray Tracing on Icosahedral Tetrahedron Grids for Non-Linear Laser Plasma Interaction Coupled to 3D Radiation Hydrodynamics". In: *Journal of Computational Physics* 443 (Oct. 2021), p. 110537. ISSN: 0021-9991. DOI: 10.1016/j.jcp.2021.110537.
 - [68] T. Afshar-rad et al. "Evidence for Whole-Beam Self-Focusing of Induced Spatially Incoherent Laser Light in Large Underdense Plasma". In: *Physical Review Letters* 68.7 (Feb. 1992), pp. 942–945. ISSN: 0031-9007. DOI: 10.1103/PhysRevLett.68.942.
 - [69] 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.
 - [70] C. D. Zhou and R. Betti. "Hydrodynamic Relations for Direct-Drive Fast-Ignition and Conventional Inertial Confinement Fusion Implosions". In: *Physics of Plasmas* 14.7 (July 2007), p. 072703. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/1.2746812.
 - [71] R. Betti et al. "Deceleration Phase of Inertial Confinement Fusion Implosions". In: *Physics of Plasmas* 9.5 (May 2002), pp. 2277–2286. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/1.1459458.
 - [72] A. Lees et al. "Experimentally Inferred Fusion Yield Dependencies of OMEGA Inertial Confinement Fusion Implosions". In: *Physical Review Letters* 127.10 (Aug. 2021), p. 105001. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.127.105001.
 - [73] A. Lees et al. "Understanding the Fusion Yield Dependencies in OMEGA DT-layered Implosion Experiments Using a Physics-Based Statistical Mapping Model". In: *Physics of Plasmas* 30.1 (Jan. 2023), p. 012709. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/5.0106515.
-

- [74] V. Gopalaswamy et al. “Tripled Yield in Direct-Drive Laser Fusion through Statistical Modelling”. In: *Nature* 565.7741 (Jan. 2019), pp. 581–586. ISSN: 0028-0836, 1476-4687. DOI: 10.1038/s41586-019-0877-0.
 - [75] C. A. Williams et al. “High Yields in Direct-Drive Inertial Confinement Fusion Using Thin-Ice DT Liner Targets”. In: *Physics of Plasmas* 28.12 (Dec. 2021), p. 122708. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/5.0069372.
 - [76] C. A. Williams et al. “Demonstration of Hot-Spot Fuel Gain Exceeding Unity in Direct-Drive Inertial Confinement Fusion Implosions”. In: *Nature Physics* 20.5 (May 2024), pp. 758–764. ISSN: 1745-2481. DOI: 10.1038/s41567-023-02363-2.
 - [77] V. Gopalaswamy et al. “Demonstration of a Hydrodynamically Equivalent Burning Plasma in Direct-Drive Inertial Confinement Fusion”. In: *Nature Physics* 20.5 (May 2024), pp. 751–757. ISSN: 1745-2473, 1745-2481. DOI: 10.1038/s41567-023-02361-4.
 - [78] Francis F. Chen. *Introduction to Plasma Physics and Controlled Fusion*. Third edition, corrected publication. Cham: Springer, 2018. ISBN: 978-3-319-22308-7.
 - [79] P. Gibbon et al. “Tree-Code Simulations of Proton Acceleration from Laser-Irradiated Wire Targets”. In: *Physics of Plasmas* 11.8 (Aug. 2004), pp. 4032–4040. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/1.1767096.
 - [80] D. Barlow et al. “Role of Hot Electrons in Shock Ignition Constrained by Experiment at the National Ignition Facility”. In: *Physics of Plasmas* 29.8 (Aug. 2022), p. 082704. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/5.0097080.
 - [81] E. M. Epperlein and R. W. Short. “A Practical Nonlocal Model for Electron Heat Transport in Laser Plasmas”. In: *Physics of Fluids B: Plasma Physics* 3.11 (Nov. 1991), pp. 3092–3098. ISSN: 0899-8221. DOI: 10.1063/1.859789.
 - [82] J. A. Bittencourt. *Fundamentals of Plasma Physics*. 3rd ed. New York: Springer, 2004. ISBN: 978-0-387-20975-3.
 - [83] A.G.R. Thomas et al. “A Review of Vlasov–Fokker–Planck Numerical Modeling of Inertial Confinement Fusion Plasma”. In: *Journal of Computational Physics* 231.3 (Feb. 2012), pp. 1051–1079. ISSN: 00219991. DOI: 10.1016/j.jcp.2011.09.028.
 - [84] R.J. Kingham and A.R. Bell. “An Implicit Vlasov–Fokker–Planck Code to Model Non-Local Electron Transport in 2-D with Magnetic Fields”. In: *Journal of Computational Physics* 194.1 (Feb. 2004), pp. 1–34. ISSN: 00219991. DOI: 10.1016/j.jcp.2003.08.017.
 - [85] W.T. Taitano et al. “An Eulerian Vlasov-Fokker–Planck Algorithm for Spherical Implosion Simulations of Inertial Confinement Fusion Capsules”. In: *Computer Physics Communications* 263 (June 2021), p. 107861. ISSN: 00104655. DOI: 10.1016/j.cpc.2021.107861.
-

- [86] C. P. Ridgers et al. “The Inadequacy of a Magnetohydrodynamic Approach to the Biermann Battery”. In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 379.2189 (Jan. 2021), p. 20200017. ISSN: 1364-503X, 1471-2962. DOI: 10.1098/rsta.2020.0017.
- [87] 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.
- [88] John I. Castor. *Radiation Hydrodynamics*. 1st ed. Cambridge University Press, Sept. 2004. ISBN: 978-0-521-83309-7 978-0-521-54062-9 978-0-511-53618-2. DOI: 10.1017/CB09780511536182.
- [89] S. I. Braginskii. “Transport Processes in a Plasma”. In: *Reviews of Plasma Physics* 1 (Jan. 1965), p. 205. URL: <https://ui.adsabs.harvard.edu/abs/1965RvPP...1..205B> (visited on 09/16/2024).
- [90] E. M. Epperlein and M. G. Haines. “Plasma Transport Coefficients in a Magnetic Field by Direct Numerical Solution of the Fokker–Planck Equation”. In: *The Physics of Fluids* 29.4 (Apr. 1986), pp. 1029–1041. ISSN: 0031-9171. DOI: 10.1063/1.865901.
- [91] J. P. Chittenden et al. “Signatures of Asymmetry in Neutron Spectra and Images Predicted by Three-Dimensional Radiation Hydrodynamics Simulations of Indirect Drive Implosions”. In: *Physics of Plasmas* 23.5 (May 2016), p. 052708. ISSN: 1070-664X. DOI: 10.1063/1.4949523.
- [92] C. A. Walsh. “Extended Magneto-Hydrodynamic Effects in Indirect Drive Inertial Confinement Fusion Experiments”. PhD thesis. London: Imperial College London, 2018.
- [93] Griffin Farrow. “Extended Magnetohydrodynamic Effects in Z-Pinch Plasmas”. PhD thesis. London: Imperial College London, Aug. 2023.
- [94] Sam Thomas Jack O’Neill. “Modelling Ignition and Burn in Pre-magnetised Inertial Confinement Fusion Experiments”. PhD thesis. London: Imperial College London, Dec. 2023.
- [95] Nikita Chaturvedi. “Simulating Radiatively Cooled Pulsed-Power Driven Systems Using Static Mesh Refinement”. PhD thesis. London: Imperial College London, Dec. 2024.
- [96] J.K. Tong et al. “Burn Regimes in the Hydrodynamic Scaling of Perturbed Inertial Confinement Fusion Hotspots”. In: *Nuclear Fusion* 59.8 (Aug. 2019), p. 086015. ISSN: 0029-5515, 1741-4326. DOI: 10.1088/1741-4326/ab22d4.
- [97] J. E. Morel. “Diffusion-Limit Asymptotics of the Transport Equation, the P1/3 Equations, and Two Flux-Limited Diffusion Theories”. In: *Journal of Quantitative Spectroscopy and Radiative Transfer* 65 (Feb. 2000), pp. 769–778. ISSN: 0022-4073. DOI: 10.1016/S0022-4073(99)00148-X.

- [98] Chris Jennings. “Radiation Transport Effects in Wire Array Z-Pinches and Magneto-Hydrodynamic Modelling Techniques”. PhD thesis. London: Imperial College London, 2005.
 - [99] K. McGlinchey. “Radiation-Hydrodynamics Simulations of the Impact of Instabilities and Asymmetries of Inertial Confinement Fusion”. PhD thesis. London: Imperial College London, 2017.
 - [100] T. S. Ramazanov and S. K. Kodanova. “Coulomb Logarithm of a Nonideal Plasma”. In: *Physics of Plasmas* 8.11 (Nov. 2001), pp. 5049–5050. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/1.1407820.
 - [101] S. K. Kodanova et al. “Investigation of Coulomb Logarithm and Relaxation Processes in Dense Plasma on the Basis of Effective Potentials”. In: *Contributions to Plasma Physics* 55.2-3 (Feb. 2015), pp. 271–276. ISSN: 0863-1042, 1521-3986. DOI: 10.1002/ctpp.201400094.
 - [102] Chengliang Lin et al. “Temperature Relaxation and Generalized Coulomb Logarithm in Two-Temperature Dense Plasmas Relevant to Inertial Confinement Fusion Implosions”. In: *Nuclear Fusion* 63.10 (Oct. 2023), p. 106005. ISSN: 0029-5515, 1741-4326. DOI: 10.1088/1741-4326/acee96.
 - [103] George K. Parks. “Why Space Physics Needs to Go beyond the MHD Box”. In: *Space Science Reviews* 113.1/2 (2004), pp. 97–121. ISSN: 0038-6308. DOI: 10.1023/b:spac.0000042941.14372.9b.
 - [104] M. Mathioudakis et al. “Alfvén Waves in the Solar Atmosphere: From Theory to Observations”. In: *Space Science Reviews* 175.1-4 (June 2013), pp. 1–27. ISSN: 0038-6308, 1572-9672. DOI: 10.1007/s11214-012-9944-7.
 - [105] Allen H. Boozer. “Equations for Studies of Feedback Stabilization”. In: *Physics of Plasmas* 5.9 (Sept. 1998), pp. 3350–3357. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/1.873048.
 - [106] M. M. Kuznetsova et al. “Multiscale Modeling of Magnetospheric Reconnection”. In: *Journal of Geophysical Research: Space Physics* 112.A10 (Oct. 2007). ISSN: 0148-0227. DOI: 10.1029/2007ja012316.
 - [107] Judit Pérez-Coll Jiménez et al. “The Role of Resistivity on the Efficiency of Magnetic Reconnection in MHD”. In: *Journal of Geophysical Research: Space Physics* 127.6 (June 2022). ISSN: 2169-9380, 2169-9402. DOI: 10.1029/2021ja030134.
 - [108] S. A. Slutz et al. “Pulsed-Power-Driven Cylindrical Liner Implosions of Laser Preheated Fuel Magnetized with an Axial Field”. In: *Physics of Plasmas* 17.5 (May 2010). ISSN: 1070-664X, 1089-7674. DOI: 10.1063/1.3333505.
 - [109] W Q Yuan et al. “Space-Time Dependent Non-Local Thermal Transport Effects on Laser Ablation Dynamics in Inertial Confinement Fusion”. In: *Plasma Physics and Controlled Fusion* 66.4 (Apr. 2024), p. 045023. ISSN: 0741-3335, 1361-6587. DOI: 10.1088/1361-6587/ad27f1.
-

- [110] G. P. Schurtz et al. "A Nonlocal Electron Conduction Model for Multidimensional Radiation Hydrodynamics Codes". In: *Physics of Plasmas* 7.10 (Oct. 2000), pp. 4238–4249. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/1.1289512.
- [111] Ph. D. Nicolaï et al. "A Practical Nonlocal Model for Heat Transport in Magnetized Laser Plasmas". In: *Physics of Plasmas* 13.3 (Mar. 2006), p. 032701. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/1.2179392.
- [112] Duc Cao et al. "Improved Non-Local Electron Thermal Transport Model for Two-Dimensional Radiation Hydrodynamics Simulations". In: *Physics of Plasmas* 22.8 (Aug. 2015), p. 082308. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/1.4928445.
- [113] M. Sherlock et al. "A Comparison of Non-Local Electron Transport Models for Laser-Plasmas Relevant to Inertial Confinement Fusion". In: *Physics of Plasmas* 24.8 (Aug. 2017), p. 082706. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/1.4986095.
- [114] Anthony R Bell and Mark Sherlock. "The *Fast* VFP Code for Solution of the Vlasov-Fokker-Planck Equation". In: *Plasma Physics and Controlled Fusion* (Jan. 2024). ISSN: 0741-3335, 1361-6587. DOI: 10.1088/1361-6587/ad2278.
- [115] N T Mitchell et al. "A Reduced Kinetic Method for Investigating Non-Local Ion Heat Transport in Ideal Multi-Species Plasmas". In: *Plasma Physics and Controlled Fusion* 66.7 (July 2024), p. 075005. ISSN: 0741-3335, 1361-6587. DOI: 10.1088/1361-6587/ad4740.
- [116] Pierre Michel. *Introduction to Laser-Plasma Interactions*. Graduate Texts in Physics. Cham: Springer, 2023. ISBN: 978-3-031-23423-1 978-3-031-23426-2.
- [117] T. Tajima and J. M. Dawson. "Laser Electron Accelerator". In: *Physical Review Letters* 43.4 (July 1979), pp. 267–270. ISSN: 0031-9007. DOI: 10.1103/PhysRevLett.43.267.
- [118] E. Esarey et al. "Physics of Laser-Driven Plasma-Based Electron Accelerators". In: *Reviews of Modern Physics* 81.3 (Aug. 2009), pp. 1229–1285. ISSN: 0034-6861, 1539-0756. DOI: 10.1103/RevModPhys.81.1229.
- [119] A. Colaïtis et al. "Towards Modeling of Nonlinear Laser-Plasma Interactions with Hydrocodes: The Thick-Ray Approach". In: *Physical Review E* 89.3 (Mar. 2014), p. 033101. DOI: 10.1103/PhysRevE.89.033101.
- [120] Thierry Dewandre et al. "Doppler Shift of Laser Light Reflected from Expanding Plasmas". In: *The Physics of Fluids* 24.3 (Mar. 1981), pp. 528–536. ISSN: 0031-9171. DOI: 10.1063/1.863401.
- [121] W. Seka et al. "Time-Resolved Absorption in Cryogenic and Room-Temperature Direct-Drive Implosions". In: *Physics of Plasmas* 15.5 (May 2008), p. 056312. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/1.2898405.
- [122] F B Bunkin et al. "INTERACTION OF INTENSE OPTICAL RADIATION WITH FREE ELECTRONS (NONRELATIVISTIC CASE)". In: *Soviet Physics Uspekhi* 15.4 (Apr. 1973), pp. 416–435. ISSN: 0038-5670. DOI: 10.1070/PU1973v015n04ABEH004990.

- [123] A. Colaïtis et al. “Real and Complex Valued Geometrical Optics Inverse Ray-Tracing for Inline Field Calculations”. In: *Physics of Plasmas* 26.3 (Mar. 2019), p. 032301. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/1.5082951.
 - [124] A. Bruce Langdon. “Nonlinear Inverse Bremsstrahlung and Heated-Electron Distributions”. In: *Physical Review Letters* 44.9 (Mar. 1980), pp. 575–579. ISSN: 0031-9007. DOI: 10.1103/PhysRevLett.44.575.
 - [125] C. P. Ridgers et al. “Transport in the Presence of Inverse Bremsstrahlung Heating and Magnetic Fields”. In: *Physics of Plasmas* 15.9 (Sept. 2008), p. 092311. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/1.2978092.
 - [126] D. Turnbull et al. “Inverse Bremsstrahlung Absorption”. In: *Physical Review Letters* 130.14 (Apr. 2023), p. 145103. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.130.145103.
 - [127] 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.
 - [128] E. A. Williams et al. “The Frequency and Damping of Ion Acoustic Waves in Hydrocarbon (CH) and Two-Ion-Species Plasmas”. In: *Physics of Plasmas* 2.1 (Jan. 1995), pp. 129–138. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/1.871101.
 - [129] P. Michel et al. “Saturation of Multi-Laser Beams Laser-Plasma Instabilities from Stochastic Ion Heating”. In: *Physics of Plasmas* 20.5 (May 2013), p. 056308. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/1.4802828.
 - [130] C. J. Randall et al. “Theory and Simulation of Stimulated Brillouin Scatter Excited by Nonabsorbed Light in Laser Fusion Systems”. In: *The Physics of Fluids* 24.8 (Aug. 1981), pp. 1474–1484. ISSN: 0031-9171. DOI: 10.1063/1.863551.
 - [131] J. A. Marozas et al. “Wavelength-Detuning Cross-Beam Energy Transfer Mitigation Scheme for Direct Drive: Modeling and Evidence from National Ignition Facility Implosions”. In: *Physics of Plasmas* 25.5 (May 2018), p. 056314. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/1.5022181.
 - [132] D. H. Edgell et al. “Nonuniform Absorption and Scattered Light in Direct-Drive Implosions Driven by Polarization Smoothing”. In: *Physical Review Letters* 127.7 (Aug. 2021), p. 075001. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.127.075001.
 - [133] A. Colaïtis et al. “3D Simulations Capture the Persistent Low-Mode Asymmetries Evident in Laser-Direct-Drive Implosions on OMEGA”. In: *Physical Review Letters* 129.9 (Aug. 2022), p. 095001. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.129.095001.
-

- [134] A Colaïtis et al. “3D Simulations of Inertial Confinement Fusion Implosions Part 2: Systematic Flow Anomalies and Impact of Low Modes on Performances in OMEGA Experiments”. In: *Plasma Physics and Controlled Fusion* 65.1 (Jan. 2023), p. 014005. ISSN: 0741-3335, 1361-6587. DOI: 10.1088/1361-6587/aca78d.
- [135] A Colaïtis et al. “3D Simulations of Inertial Confinement Fusion Implosions Part 1: In-line Modeling of Polarized Cross Beam Energy Transfer and Subsequent Drive Anomalies on OMEGA and NIF”. In: *Plasma Physics and Controlled Fusion* 65.1 (Jan. 2023), p. 014003. ISSN: 0741-3335, 1361-6587. DOI: 10.1088/1361-6587/aca78e.
- [136] O. M. Mannion et al. “Mitigation of Mode-One Asymmetry in Laser-Direct-Drive Inertial Confinement Fusion Implosions”. In: *Physics of Plasmas* 28.4 (Apr. 2021), p. 042701. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/5.0041554.
- [137] Bedros B. Afeyan et al. “Kinetic Theory of Electron-Plasma and Ion-Acoustic Waves in Nonuniformly Heated Laser Plasmas”. In: *Physical Review Letters* 80.11 (Mar. 1998), pp. 2322–2325. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.80.2322.
- [138] David Turnbull et al. “Impact of the Langdon Effect on Crossed-Beam Energy Transfer”. In: *Nature Physics* 16.2 (Feb. 2020), pp. 181–185. ISSN: 1745-2481. DOI: 10.1038/s41567-019-0725-z.
- [139] P. Michel et al. “Stochastic Ion Heating from Many Overlapping Laser Beams in Fusion Plasmas”. In: *Physical Review Letters* 109.19 (Nov. 2012), p. 195004. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.109.195004.
- [140] V N Goncharov et al. “National Direct-Drive Program on OMEGA and the National Ignition Facility”. In: *Plasma Physics and Controlled Fusion* 59.1 (Jan. 2017), p. 014008. ISSN: 0741-3335, 1361-6587. DOI: 10.1088/0741-3335/59/1/014008.
- [141] E.M. Campbell et al. “Laser-Direct-Drive Program: Promise, Challenge, and Path Forward”. In: *Matter and Radiation at Extremes* 2.2 (Mar. 2017), pp. 37–54. ISSN: 2468-2047, 2468-080X. DOI: 10.1016/j.mre.2017.03.001.
- [142] D. H. Froula et al. “Increasing Hydrodynamic Efficiency by Reducing Cross-Beam Energy Transfer in Direct-Drive-Implosion Experiments”. In: *Physical Review Letters* 108.12 (Mar. 2012), p. 125003. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.108.125003.
- [143] D. H. Froula et al. “Mitigation of Cross-Beam Energy Transfer: Implication of Two-State Focal Zooming on OMEGA”. In: *Physics of Plasmas* 20.8 (Aug. 2013), p. 082704. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/1.4818427.
- [144] I. V. Igumenshchev et al. “Laser-Beam Zooming to Mitigate Crossed-Beam Energy Losses in Direct-Drive Implosions”. In: *Physical Review Letters* 110.14 (Apr. 2013), p. 145001. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.110.145001.
- [145] D. M. Kehne et al. “Implementation of Focal Zooming on the Nike KrF Laser”. In: *Review of Scientific Instruments* 84.1 (Jan. 2013), p. 013509. ISSN: 0034-6748, 1089-7623. DOI: 10.1063/1.4789313.
-

- [146] Bedros Afeyan and Stefan Hüller. “Optimal Control of Laser Plasma Instabilities Using Spike Trains of Uneven Duration and Delay (STUD Pulses) for ICF and IFE”. In: *EPJ Web of Conferences* 59 (2013). Ed. by P. Mora et al., p. 05009. ISSN: 2100-014X. DOI: 10.1051/epjconf/20135905009.
 - [147] D. H. Edgell et al. “Mitigation of Cross-Beam Energy Transfer in Symmetric Implosions on OMEGA Using Wavelength Detuning”. In: *Physics of Plasmas* 24.6 (June 2017), p. 062706. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/1.4985315.
 - [148] J. A. Marozas et al. “First Observation of Cross-Beam Energy Transfer Mitigation for Direct-Drive Inertial Confinement Fusion Implosions Using Wavelength Detuning at the National Ignition Facility”. In: *Physical Review Letters* 120.8 (Feb. 2018), p. 085001. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.120.085001.
 - [149] J. W. Bates et al. “Mitigation of Cross-Beam Energy Transfer in Inertial-Confinement-Fusion Plasmas with Enhanced Laser Bandwidth”. In: *Physical Review E* 97.6 (June 2018), p. 061202. ISSN: 2470-0045, 2470-0053. DOI: 10.1103/PhysRevE.97.061202.
 - [150] J. W. Bates et al. “Suppressing Parametric Instabilities in Direct-Drive Inertial-Confinement-Fusion Plasmas Using Broadband Laser Light”. In: *Physics of Plasmas* 30.5 (May 2023), p. 052703. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/5.0150865.
 - [151] A. G. Seaton et al. “Cross-Beam Energy Transfer in Direct-Drive ICF. II. Theory and Simulation of Mitigation through Increased Laser Bandwidth”. In: *Physics of Plasmas* 29.4 (Apr. 2022), p. 042707. ISSN: 1070-664X. DOI: 10.1063/5.0078801.
 - [152] A. Colaïtis et al. “Exploration of Cross-Beam Energy Transfer Mitigation Constraints for Designing an Ignition-Scale Direct-Drive Inertial Confinement Fusion Driver”. In: *Physics of Plasmas* 30.8 (Aug. 2023), p. 082701. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/5.0150813.
 - [153] R. K. Follett et al. “Ray-Based Cross-Beam Energy Transfer Modeling for Broadband Lasers”. In: *Physics of Plasmas* 30.4 (Apr. 2023), p. 042102. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/5.0137420.
 - [154] C. Deeney. “Advances in Direct-Drive Fusion, High-Energy-Density Physics, and Laser Technologies at the Laboratory for Laser Energetics”. In: *2022 IEEE International Conference on Plasma Science (ICOPS)*. Seattle, WA, USA: IEEE, May 2022, pp. 1–2. ISBN: 978-1-66547-925-7. DOI: 10.1109/ICOPS45751.2022.9813166.
 - [155] P. Tzeferacos et al. “Laboratory Evidence of Dynamo Amplification of Magnetic Fields in a Turbulent Plasma”. In: *Nature Communications* 9.1 (Feb. 2018), p. 591. ISSN: 2041-1723. DOI: 10.1038/s41467-018-02953-2.
 - [156] F. Fiuza et al. “Electron Acceleration in Laboratory-Produced Turbulent Collisionless Shocks”. In: *Nature Physics* 16.9 (Sept. 2020), pp. 916–920. ISSN: 1745-2481. DOI: 10.1038/s41567-020-0919-4.
-

- [157] Jena Meinecke et al. “Strong Suppression of Heat Conduction in a Laboratory Replica of Galaxy-Cluster Turbulent Plasmas”. In: *Science Advances* 8.10 (Mar. 2022), eabj6799. ISSN: 2375-2548. DOI: 10.1126/sciadv.abj6799.
- [158] Andrea L. Kritcher et al. “A Measurement of the Equation of State of Carbon Envelopes of White Dwarfs”. In: *Nature* 584.7819 (Aug. 2020), pp. 51–54. ISSN: 0028-0836, 1476-4687. DOI: 10.1038/s41586-020-2535-y.
- [159] R. F. Smith et al. “Ramp Compression of Diamond to Five Terapascals”. In: *Nature* 511.7509 (July 2014), pp. 330–333. ISSN: 1476-4687. DOI: 10.1038/nature13526.
- [160] Stephen A. Slutz and Roger A. Vesey. “High-Gain Magnetized Inertial Fusion”. In: *Physical Review Letters* 108.2 (Jan. 2012), p. 025003. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.108.025003.
- [161] K. L. Nguyen et al. “Cross-Beam Energy Transfer Saturation by Ion Trapping-Induced Detuning”. In: *Physics of Plasmas* 28.8 (Aug. 2021), p. 082705. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/5.0054008.
- [162] J. F. Myatt et al. “A Wave-Based Model for Cross-Beam Energy Transfer in Direct-Drive Inertial Confinement Fusion”. In: *Physics of Plasmas* 24.5 (Apr. 2017), p. 056308. ISSN: 1070-664X. DOI: 10.1063/1.4982059.
- [163] Andrew J. Schmitt and Stephen P. Obenschain. “The Importance of Laser Wavelength for Driving Inertial Confinement Fusion Targets. II. Target Design”. In: *Physics of Plasmas* 30.1 (Jan. 2023), p. 012702. ISSN: 1070-664X. DOI: 10.1063/5.0118093.
- [164] Andrew J. Schmitt and Stephen P. Obenschain. “The Importance of Laser Wavelength for Driving Inertial Confinement Fusion Targets. I. Basic Physics”. In: *Physics of Plasmas* 30.1 (Jan. 2023), p. 012701. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/5.0118080.
- [165] Tudor Wyatt Johnston and John M. Dawson. “Correct Values for High-Frequency Power Absorption by Inverse Bremsstrahlung in Plasmas”. In: *Physics of Fluids* 16 (May 1973), pp. 722–722. ISSN: 0899-8213/1070-6631. DOI: 10.1063/1.1694419.
- [166] R. W. Paddock et al. “One-Dimensional Hydrodynamic Simulations of Low Convergence Ratio Direct-Drive Inertial Confinement Fusion Implosions”. In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 379.2189 (Jan. 2021), p. 20200224. ISSN: 1364-503X, 1471-2962. DOI: 10.1098/rsta.2020.0224.
- [167] A. Colaitis et al. “Adaptive Inverse Ray-Tracing for Accurate and Efficient Modeling of Cross Beam Energy Transfer in Hydrodynamics Simulations”. In: *Physics of Plasmas* 26.7 (July 2019), p. 072706. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/1.5108777.
- [168] D. J. Strozzi et al. “Interplay of Laser-Plasma Interactions and Inertial Fusion Hydrodynamics”. In: *Physical Review Letters* 118.2 (Jan. 2017), p. 025002. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.118.025002.

- [169] S. Liberatore et al. "First Indirect Drive Inertial Confinement Fusion Campaign at Laser Megajoule". In: *Physics of Plasmas* 30.12 (Dec. 2023), p. 122707. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/5.0176446.
- [170] A. Debayle et al. "A Unified Modeling of Wave Mixing Processes with the Ray Tracing Method". In: *Physics of Plasmas* 26.9 (Sept. 2019), p. 092705. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/1.5110247.
- [171] A. Colaïtis et al. "Crossed Beam Energy Transfer: Assessment of the Paraxial Complex Geometrical Optics Approach versus a Time-Dependent Paraxial Method to Describe Experimental Results". In: *Physics of Plasmas* 23.3 (Mar. 2016), p. 032118. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/1.4944496.
- [172] A. Colaïtis et al. "Modeling of the Cross-Beam Energy Transfer with Realistic Inertial-Confinement-Fusion Beams in a Large-Scale Hydrocode". In: *Physical Review E* 91.1 (Jan. 2015), p. 013102. ISSN: 1539-3755, 1550-2376. DOI: 10.1103/PhysRevE.91.013102.
- [173] R. K. Follett et al. "Ray-Based Modeling of Cross-Beam Energy Transfer at Caustics". In: *Physical Review E* 98.4 (Oct. 2018), p. 043202. ISSN: 2470-0045, 2470-0053. DOI: 10.1103/PhysRevE.98.043202.
- [174] I. V. Igumenshchev et al. "Crossed-Beam Energy Transfer in Implosion Experiments on OMEGA". In: *Physics of Plasmas* 17.12 (Dec. 2010), p. 122708. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/1.3532817.
- [175] I. V. Igumenshchev et al. "Crossed-Beam Energy Transfer in Direct-Drive Implosions". In: *Physics of Plasmas* 19.5 (May 2012), p. 056314. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/1.4718594.
- [176] Jason F. Myatt et al. "LPSE: A 3-D Wave-Based Model of Cross-Beam Energy Transfer in Laser-Irradiated Plasmas". In: *Journal of Computational Physics* 399 (Dec. 2019), p. 108916. ISSN: 00219991. DOI: 10.1016/j.jcp.2019.108916.
- [177] R. K. Follett et al. "Thresholds of Absolute Two-Plasmon-Decay and Stimulated Raman Scattering Instabilities Driven by Multiple Broadband Lasers". In: *Physics of Plasmas* 28.3 (Mar. 2021), p. 032103. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/5.0037869.
- [178] R. L. Berger et al. "On the Dominant and Subdominant Behavior of Stimulated Raman and Brillouin Scattering Driven by Nonuniform Laser Beams". In: *Physics of Plasmas* 5.12 (Dec. 1998), pp. 4337–4356. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/1.873171.
- [179] A. G. Seaton et al. "Cross-Beam Energy Transfer in Direct-Drive ICF. I. Nonlinear and Kinetic Effects". In: *Physics of Plasmas* 29.4 (Apr. 2022), p. 042706. ISSN: 1070-664X. DOI: 10.1063/5.0078800.
-

- [180] Hyeong-Bin Cheong and Hyun-Gyu Kang. “Eigensolutions of the Spherical Laplacian for the Cubed-sphere and Icosahedral-hexagonal Grids”. In: *Quarterly Journal of the Royal Meteorological Society* 141.693 (Oct. 2015), pp. 3383–3398. ISSN: 0035-9009, 1477-870X. DOI: 10.1002/qj.2620.
- [181] Zinovy Malkin. “A New Equal-area Isolatitude Grid on a Spherical Surface”. In: *The Astronomical Journal* 158.4 (Oct. 2019), p. 158. ISSN: 0004-6256, 1538-3881. DOI: 10.3847/1538-3881/ab3a44.
- [182] Thomas B. Kaiser. “Laser Ray Tracing and Power Deposition on an Unstructured Three-Dimensional Grid”. In: *Physical Review E* 61.1 (Jan. 2000), pp. 895–905. DOI: 10.1103/PhysRevE.61.895.
- [183] William H. Press, ed. *Numerical Recipes: The Art of Scientific Computing*. 3. ed. Cambridge: Cambridge University Press, 2007. ISBN: 978-0-521-88068-8.
- [184] Jacob Williams. *Bspline-Fortran*. Jan. 2024. URL: <https://zenodo.org/records/8423936>.
- [185] Charles K. Birdsall and A. Bruce Langdon. *Plasma Physics via Computer Simulation*. New York: McGraw-Hill, 1985. ISBN: 978-0-07-005371-7.
- [186] Christophe Cornet and Dixon T.K. Kwok. “A New Algorithm for Charge Deposition for Multiple-Grid Method for PIC Simulations in r-z Cylindrical Coordinates”. In: *Journal of Computational Physics* 225.1 (July 2007), pp. 808–828. ISSN: 00219991. DOI: 10.1016/j.jcp.2007.01.004.
- [187] Brian M. Haines et al. “Coupling Laser Physics to Radiation-Hydrodynamics”. In: *Computers & Fluids* 201 (Apr. 2020), p. 104478. ISSN: 00457930. DOI: 10.1016/j.compfluid.2020.104478.
- [188] J. Denavit and D. W. Phillion. “Laser Ionization and Heating of Gas Targets for Long-Scale-Length Instability Experiments”. In: *Physics of Plasmas* 1.6 (June 1994), pp. 1971–1984. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/1.870653.
- [189] I. V. Igumenshchev et al. “The Effects of Target Mounts in Direct-Drive Implosions on OMEGA”. In: *Physics of Plasmas* 16.8 (Aug. 2009), p. 082701. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/1.3195065.
- [190] I. V. Igumenshchev et al. “Three-Dimensional Modeling of Direct-Drive Cryogenic Implosions on OMEGA”. In: *Physics of Plasmas* 23.5 (May 2016), p. 052702. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/1.4948418.
- [191] William L. Kruer. *The Physics of Laser Plasma Interactions*. Frontiers in Physics. Boulder, Colo. [u.a.] Westview: CRC Press, 2003. ISBN: 978-0-8133-4083-8.
- [192] Jurij A. Kravcov and Jurij I. Orlov. *Caustics, Catastrophes and Wave Fields*. 2. ed. Springer Series on Wave Phenomena 15. Berlin Heidelberg New York: Springer, 1999. ISBN: 978-3-540-64227-5.
- [193] Yury Alexandrovich Kratsov and Ning Yan Zhu. *Theory of Diffraction: Heuristic Approaches*. Oxford: Alpha Science International, 2010. ISBN: 978-1-84265-372-2.

- [194] N. A. Lopez and I. Y. Dodin. “Metaplectic Geometrical Optics for Ray-Based Modeling of Caustics: Theory and Algorithms”. In: *Physics of Plasmas* 29.5 (May 2022), p. 052111. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/5.0082241.
 - [195] Shan-chieh Chang and Jian-Ming Jin. *Computation of Special Functions*. New York: Wiley, 1996. ISBN: 978-0-471-11963-0.
 - [196] A.A. Asatryan and Yu.A. Kravtsov. “Fresnel Zones of Hyperbolic Type from the Physical Point of View”. In: *Wave Motion* 10.1 (Jan. 1988), pp. 45–57. ISSN: 01652125. DOI: 10.1016/0165-2125(88)90005-4.
 - [197] Yu.A. Kravtsov. “IV Rays and Caustics as Physical Objects”. In: *Progress in Optics*. Vol. 26. Elsevier, 1988, pp. 227–348. ISBN: 978-0-444-87096-4. DOI: 10.1016/S0079-6638(08)70177-X.
 - [198] R. K. Follett. *LPSE Data for Ray-Based CBET Test Cases*. Aug. 2022. DOI: 10.5281/ZENODO.6962934.
 - [199] Jiachen Ding et al. “Identify the Limits of Geometric Optics Ray Tracing by Numerically Solving the Vector Kirchhoff Integral”. In: *Optics Express* 28.7 (Mar. 2020), p. 10670. ISSN: 1094-4087. DOI: 10.1364/OE.389097.
 - [200] G. P. M. Poppe and C. M. J. Wijers. “Algorithm 680: Evaluation of the Complex Error Function”. In: *ACM Transactions on Mathematical Software* 16.1 (Mar. 1990), p. 47. ISSN: 0098-3500, 1557-7295. DOI: 10.1145/77626.77630.
 - [201] Burton D. Fried and Samuel D. Conte. *The Plasma Dispersion Function: The Hilbert Transform of the Gaussian*. Academic Press, 1961. ISBN: 978-1-4832-2929-4.
 - [202] Dhrumir Patel et al. “Novel Double-Spike Pulse Shape for OMEGA Cryogenic Implosions”. In: 2018 (Jan. 2018), GO6.006. URL: <https://ui.adsabs.harvard.edu/abs/2018APS..DPPG06006P> (visited on 07/28/2024).
 - [203] D. Cao et al. “Interpreting the Electron Temperature Inferred from X-Ray Continuum Emission for Direct-Drive Inertial Confinement Fusion Implosions on OMEGA”. In: *Physics of Plasmas* 26.8 (Aug. 2019), p. 082709. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/1.5112759.
 - [204] A. J. Crilly et al. “Neutron Backscatter Edge: A Measure of the Hydrodynamic Properties of the Dense DT Fuel at Stagnation in ICF Experiments”. In: *Physics of Plasmas* 27.1 (Jan. 2020), p. 012701. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/1.5128830.
 - [205] James McHardy. *An Introduction to the Theory and Use of SESAME Equations of State*. Tech. rep. LA-14503, 1487368. Dec. 2018, LA-14503, 1487368. DOI: 10.2172/1487368.
 - [206] A.J. Crilly et al. “SpK: A Fast Atomic and Microphysics Code for the High-Energy-Density Regime”. In: *High Energy Density Physics* 48 (Sept. 2023), p. 101053. ISSN: 15741818. DOI: 10.1016/j.hedp.2023.101053.
 - [207] 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.
-

- [208] W. Trickey et al. "The Physics of Gain Relevant to Inertial Fusion Energy Target Designs". In: *Physics of Plasmas* 31.1 (Jan. 2024), p. 012702. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/5.0167405.
- [209] K. S. Anderson et al. "Enhanced Sensitivity to Target Offset When Using Cross-Beam Energy Transfer Mitigation Techniques in Direct-Drive Inertial Confinement Fusion Implosions". In: *Physics of Plasmas* 31.3 (Mar. 2024), p. 032704. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/5.0191277.
- [210] V. Gopalaswamy et al. "Using Statistical Modeling to Predict and Understand Fusion Experiments". In: *Physics of Plasmas* 28.12 (Dec. 2021), p. 122705. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/5.0056662.
- [211] R. Ejaz et al. "Deep Learning-Based Predictive Models for Laser Direct Drive at the Omega Laser Facility". In: *Physics of Plasmas* 31.5 (May 2024), p. 052703. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/5.0195675.
- [212] C. A. Thomas et al. "Quantifying the Effects of Scale and Illumination Geometry in Laser Direct Drive". In: 2020 (Jan. 2020), BO09.014. URL: <https://ui.adsabs.harvard.edu/abs/2020APS..DPPB09014T> (visited on 08/02/2024).
- [213] S. X. Hu et al. "Mitigating Laser Imprint in Direct-Drive Inertial Confinement Fusion Implosions with High- Z Dopants". In: *Physical Review Letters* 108.19 (May 2012), p. 195003. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.108.195003.
- [214] Max Karasik et al. "Suppression of Laser Nonuniformity Imprinting Using a Thin High- Z Coating". In: *Physical Review Letters* 114.8 (Feb. 2015), p. 085001. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.114.085001.
- [215] S. X. Hu et al. "Mitigating Laser-Imprint Effects in Direct-Drive Inertial Confinement Fusion Implosions with an above-Critical-Density Foam Layer". In: *Physics of Plasmas* 25.8 (Aug. 2018), p. 082710. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/1.5044609.
- [216] T. R. Dittrich et al. "Diagnosis of Pusher-Fuel Mix in Indirectly Driven Nova Implosions". In: *Physical Review Letters* 73.17 (Oct. 1994), pp. 2324–2327. ISSN: 0031-9007. DOI: 10.1103/PhysRevLett.73.2324.
- [217] S. C. Miller and V. N. Goncharov. "Instability Seeding Mechanisms Due to Internal Defects in Inertial Confinement Fusion Targets". In: *Physics of Plasmas* 29.8 (Aug. 2022), p. 082701. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/5.0091949.
- [218] A. G. MacPhee et al. "Stabilization of High-Compression, Indirect-Drive Inertial Confinement Fusion Implosions Using a 4-Shock Adiabatic-Shaped Drive". In: *Physics of Plasmas* 22.8 (Aug. 2015), p. 080702. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/1.4928909.
- [219] V.A. Smalyuk et al. "Mix and Hydrodynamic Instabilities on NIF". In: *Journal of Instrumentation* 12.06 (June 2017), pp. C06001–C06001. ISSN: 1748-0221. DOI: 10.1088/1748-0221/12/06/C06001.
-

- [220] Brian M. Haines et al. “Robustness to Hydrodynamic Instabilities in Indirectly Driven Layered Capsule Implosions”. In: *Physics of Plasmas* 26.1 (Jan. 2019), p. 012707. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/1.5080262.
- [221] S.P. Regan et al. “The National Direct-Drive Inertial Confinement Fusion Program”. In: *Nuclear Fusion* 59.3 (Mar. 2019), p. 032007. ISSN: 0029-5515, 1741-4326. DOI: 10.1088/1741-4326/aae9b5.
- [222] K. M. Woo et al. “Inferring Thermal Ion Temperature and Residual Kinetic Energy from Nuclear Measurements in Inertial Confinement Fusion Implosions”. In: *Physics of Plasmas* 27.6 (June 2020), p. 062702. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/1.5144460.
- [223] Rahman Ejaz et al. *Can Kans (Re)Discover Predictive Models for Direct-Drive Laser Fusion?* Sept. 2024. arXiv: 2409.08832 [cs]. URL: <http://arxiv.org/abs/2409.08832> (visited on 11/14/2024).
- [224] G. Pérez-Callejo et al. “Cylindrical Implosion Platform for the Study of Highly Magnetized Plasmas at Laser MegaJoule”. In: *Physical Review E* 106.3 (Sept. 2022), p. 035206. ISSN: 2470-0045, 2470-0053. DOI: 10.1103/PhysRevE.106.035206.
- [225] D.L. Tubbs et al. “Direct-Drive Cylindrical Implosion Experiments: Simulations and Data”. In: *Laser and Particle Beams* 17.3 (July 1999), pp. 437–449. ISSN: 0263-0346, 1469-803X. DOI: 10.1017/S0263034699173117.
- [226] D. W. Peaceman and H. H. Rachford Jr. “The Numerical Solution of Parabolic and Elliptic Differential Equations”. In: *Journal of the Society for Industrial and Applied Mathematics* 3.1 (Mar. 1955), pp. 28–41. ISSN: 0368-4245, 2168-3484. DOI: 10.1137/0103003.
- [227] A. Bose et al. “The Physics of Long- and Intermediate-Wavelength Asymmetries of the Hot Spot: Compression Hydrodynamics and Energetics”. In: *Physics of Plasmas* 24.10 (Oct. 2017), p. 102704. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/1.4995250.
- [228] A. Bose et al. “Analysis of Trends in Experimental Observables: Reconstruction of the Implosion Dynamics and Implications for Fusion Yield Extrapolation for Direct-Drive Cryogenic Targets on OMEGA”. In: *Physics of Plasmas* 25.6 (June 2018), p. 062701. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/1.5026780.
- [229] D. Patel et al. “Effects of Laser Bandwidth in Direct-Drive High-Performance DT-Layered Implosions on the OMEGA Laser”. In: *Physical Review Letters* 131.10 (Sept. 2023), p. 105101. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.131.105101.
- [230] T. R. Dittrich et al. “Design of a High-Foot High-Adiabatic ICF Capsule for the National Ignition Facility”. In: *Physical Review Letters* 112.5 (Feb. 2014), p. 055002. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.112.055002.
-

- [231] H. F. Robey et al. "Performance of Indirectly Driven Capsule Implosions on the National Ignition Facility Using Adiabatic-Shaping". In: *Physics of Plasmas* 23.5 (May 2016), p. 056303. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/1.4944821.
- [232] S. X. Hu et al. "Direct-Drive Double-Shell Implosion: A Platform for Burning-Plasma Physics Studies". In: *Physical Review E* 100.6 (Dec. 2019), p. 063204. ISSN: 2470-0045, 2470-0053. DOI: 10.1103/PhysRevE.100.063204.
- [233] I. V. Igumenshchev et al. "Rarefaction Flows and Mitigation of Imprint in Direct-Drive Implosions". In: *Physical Review Letters* 123.6 (Aug. 2019), p. 065001. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.123.065001.
- [234] I.R. Lindemuth and R.C. Kirkpatrick. "Parameter Space for Magnetized Fuel Targets in Inertial Confinement Fusion". In: *Nuclear Fusion* 23.3 (Mar. 1983), pp. 263–284. ISSN: 0029-5515, 1741-4326. DOI: 10.1088/0029-5515/23/3/001.
- [235] R.D. Jones and W.C. Mead. "The Physics of Burn in Magnetized Deuterium-Tritium Plasmas: Spherical Geometry". In: *Nuclear Fusion* 26.2 (Feb. 1986), pp. 127–137. ISSN: 0029-5515, 1741-4326. DOI: 10.1088/0029-5515/26/2/001.
- [236] G. Fiksel et al. "Note: Experimental Platform for Magnetized High-Energy-Density Plasma Studies at the Omega Laser Facility". In: *Review of Scientific Instruments* 86.1 (Jan. 2015), p. 016105. ISSN: 0034-6748, 1089-7623. DOI: 10.1063/1.4905625.
- [237] L. J. Perkins et al. "The Potential of Imposed Magnetic Fields for Enhancing Ignition Probability and Fusion Energy Yield in Indirect-Drive Inertial Confinement Fusion". In: *Physics of Plasmas* 24.6 (June 2017), p. 062708. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/1.4985150.
- [238] J. D. Moody et al. "The Magnetized Indirect Drive Project on the National Ignition Facility". In: *Journal of Fusion Energy* 41.1 (June 2022), p. 7. ISSN: 0164-0313, 1572-9591. DOI: 10.1007/s10894-022-00319-7.
- [239] C. A. Walsh et al. "Perturbation Modifications by Pre-Magnetisation of Inertial Confinement Fusion Implosions". In: *Physics of Plasmas* 26.2 (Feb. 2019), p. 022701. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/1.5085498.
- [240] C.A. Walsh et al. "Magnetized Directly-Driven ICF Capsules: Increased Instability Growth from Non-Uniform Laser Drive". In: *Nuclear Fusion* 60.10 (Oct. 2020), p. 106006. ISSN: 0029-5515, 1741-4326. DOI: 10.1088/1741-4326/abab52.
- [241] J. D. Moody et al. "Increased Ion Temperature and Neutron Yield Observed in Magnetized Indirectly Driven D₂-Filled Capsule Implosions on the National Ignition Facility". In: *Physical Review Letters* 129.19 (Nov. 2022), p. 195002. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.129.195002.
- [242] Chris Walsh et al. "Application of Non-Uniform Magnetic Fields to Inertial Confinement Fusion Implosions". In: 2023 (Jan. 2023), JO07.014. URL: <https://ui.adsabs.harvard.edu/abs/2023APS..DPPJ07014W> (visited on 08/17/2024).

- [243] A. Bose et al. "Effect of Strongly Magnetized Electrons and Ions on Heat Flow and Symmetry of Inertial Fusion Implosions". In: *Physical Review Letters* 128.19 (May 2022), p. 195002. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.128.195002.
 - [244] O. V. Gotchev et al. "Laser-Driven Magnetic-Flux Compression in High-Energy-Density Plasmas". In: *Physical Review Letters* 103.21 (Nov. 2009), p. 215004. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.103.215004.
 - [245] P. Y. Chang et al. "Fusion Yield Enhancement in Magnetized Laser-Driven Implosions". In: *Physical Review Letters* 107.3 (July 2011), p. 035006. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.107.035006.
 - [246] M. Hohenberger et al. "Inertial Confinement Fusion Implosions with Imposed Magnetic Field Compression Using the OMEGA Laser". In: *Physics of Plasmas* 19.5 (May 2012), p. 056306. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/1.3696032.
 - [247] K. S. Anderson et al. "Effect of Cross-Beam Energy Transfer on Target-Offset Asymmetry in Direct-Drive Inertial Confinement Fusion Implosions". In: *Physics of Plasmas* 27.11 (Nov. 2020), p. 112713. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/5.0015781.
 - [248] Yuan Shi. "Benchmarking Magnetised Three-Wave Coupling for Laser Backscattering: Analytic Solutions and Kinetic Simulations". In: *Journal of Plasma Physics* 89.3 (June 2023), p. 905890305. ISSN: 0022-3778, 1469-7807. DOI: 10.1017/S0022377823000405.
 - [249] Yuan Shi et al. "Three-Wave Scattering in Magnetized Plasmas: From Cold Fluid to Quantized Lagrangian". In: *Physical Review E* 96.2 (Aug. 2017), p. 023204. ISSN: 2470-0045, 2470-0053. DOI: 10.1103/PhysRevE.96.023204.
 - [250] Yuan Shi et al. "Laser-Plasma Interactions in Magnetized Environment". In: *Physics of Plasmas* 25.5 (May 2018), p. 055706. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/1.5017980.
 - [251] C. Leland Ellison et al. "Development and Modeling of a Polar-Direct-Drive Exploding Pusher Platform at the National Ignition Facility". In: *Physics of Plasmas* 25.7 (July 2018), p. 072710. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/1.5025724.
 - [252] C.B. Yeamans et al. "High Yield Polar Direct Drive Fusion Neutron Sources at the National Ignition Facility". In: *Nuclear Fusion* 61.4 (Apr. 2021), p. 046031. ISSN: 0029-5515, 1741-4326. DOI: 10.1088/1741-4326/abe4e6.
 - [253] Lorin X. Benedict et al. "Molecular Dynamics Simulations and Generalized Lenard-Balescu Calculations of Electron-Ion Temperature Equilibration in Plasmas". In: *Physical Review E* 86.4 (Oct. 2012), p. 046406. ISSN: 1539-3755, 1550-2376. DOI: 10.1103/PhysRevE.86.046406.
 - [254] O. M. Mannion et al. "Evidence of Non-Maxwellian Ion Velocity Distributions in Spherical Shock-Driven Implosions". In: *Physical Review E* 108.3 (Sept. 2023), p. 035201. ISSN: 2470-0045, 2470-0053. DOI: 10.1103/PhysRevE.108.035201.
-

- [255] O. Larroche et al. “Ion-Kinetic Simulations of D-3He Gas-Filled Inertial Confinement Fusion Target Implosions with Moderate to Large Knudsen Number”. In: *Physics of Plasmas* 23.1 (Jan. 2016), p. 012701. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/1.4939025.
 - [256] Cody Wu Chang and J. A. Frenje. Apr. 2023.
 - [257] C. A. Walsh et al. “Extended-Magnetohydrodynamics in under-Dense Plasmas”. In: *Physics of Plasmas* 27.2 (Feb. 2020), p. 022103. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/1.5124144.
 - [258] Bhargav Vaidya et al. “Scalable Explicit Implementation of Anisotropic Diffusion with Runge–Kutta–Legendre Super-Time Stepping”. In: *Monthly Notices of the Royal Astronomical Society* 472.3 (Dec. 2017), pp. 3147–3160. ISSN: 0035-8711, 1365-2966. DOI: 10.1093/mnras/stx2176.
 - [259] J A Frenje. “Nuclear Diagnostics for Inertial Confinement Fusion (ICF) Plasmas”. In: *Plasma Physics and Controlled Fusion* 62.2 (Feb. 2020), p. 023001. ISSN: 0741-3335, 1361-6587. DOI: 10.1088/1361-6587/ab5137.
 - [260] H.-S Bosch and G.M Hale. “Improved Formulas for Fusion Cross-Sections and Thermal Reactivities”. In: *Nuclear Fusion* 32.4 (Apr. 1992), pp. 611–631. ISSN: 0029-5515. DOI: 10.1088/0029-5515/32/4/I07.
 - [261] K. M. Woo et al. “Impact of Three-Dimensional Hot-Spot Flow Asymmetry on Ion-Temperature Measurements in Inertial Confinement Fusion Experiments”. In: *Physics of Plasmas* 25.10 (Oct. 2018), p. 102710. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/1.5048429.
 - [262] D. T. Casey et al. “Three Dimensional Low-Mode Areal-Density Non-Uniformities in Indirect-Drive Implosions at the National Ignition Facility”. In: *Physics of Plasmas* 28.4 (Apr. 2021), p. 042708. ISSN: 1070-664X, 1089-7674. DOI: 10.1063/5.0043589.
 - [263] E. R. Tracy and Alain Jean Brizard. *Ray Tracing and beyond: Phase Space Methods in Plasma Wave Theory*. Cambridge, United Kingdom: Cambridge University Press, 2014. ISBN: 978-1-139-87060-3.
 - [264] J. P. Chittenden. *Cross-Beam Energy Transfer in Magnetised Laser Driven Hohlräume*. Denver, Sept. 2024.
 - [265] Yao Zhao et al. “Mitigation of Laser Plasma Parametric Instabilities with Broadband Lasers”. In: *Reviews of Modern Plasma Physics* 7.1 (Dec. 2022), p. 1. ISSN: 2367-3192. DOI: 10.1007/s41614-022-00105-0.
 - [266] A. Tentori et al. “3D Monte-Carlo Model to Study the Transport of Hot Electrons in the Context of Inertial Confinement Fusion. Part I”. In: *Matter and Radiation at Extremes* 7.6 (Nov. 2022), p. 065902. ISSN: 2468-2047, 2468-080X. DOI: 10.1063/5.0103631.
 - [267] A. Tentori et al. “3D Monte-Carlo Model to Study the Transport of Hot Electrons in the Context of Inertial Confinement Fusion. Part II”. In: *Matter and Radiation at Extremes* 7.6 (Nov. 2022), p. 065903. ISSN: 2468-2047, 2468-080X. DOI: 10.1063/5.0103632.
-