

**Chiron: A Ray Based
Laser-Plasma Instability
Model for
Radiation-Hydrodynamics
Codes**

Arun Vijay Nair Nutter

Doctor of Philosophy

University of York

Physics

August 2024

Abstract

A new approach to modelling laser-plasma instabilities (LPI) has been developed that is designed to run in-line with ray tracing routines alongside radiation-hydrodynamics simulations of inertial confinement fusion (ICF) implosions. Accurate modelling of ICF experiments is a crucial aspect to advancing the field. Many approaches to LPI modelling exist [1–4] because, although predicting LPI levels is essential to ICF simulations, developing models that encompass the physics of the highly complex, nonlinear LPI processes is very difficult.

The model developed in this work, called Chiron, is based on analytic theory for the steady-state solutions to LPI levels. Chiron is unique in that it combines the simultaneous growth of stimulated Raman scattering, stimulated Brillouin scattering, and two-plasmon decay along with the inverse bremsstrahlung absorption mechanism, while featuring saturation by convection, Langmuir decay and pump depletion for the LPI models.

The famous Rosenbluth gain [5] has been adapted into a new expression for convective gain called the “bounded” gain. A new expression for saturation of instabilities via Langmuir decay has also been derived. The LPI models are saturated both convectively and absolutely depending on which mechanism the model determines is dominant. The regimes of applicability of both Rosenbluth and bounded gain are defined by comparing to a more physically complete numerically (rather than analytically) integrated LPI model. An iterative method has been developed to deal with how pump depletion nonlinearly affects LPI saturation levels, which has allowed the pump amplitude to be treated as a constant in the analytic theory, keeping it tractable.

Chiron was tested by simulating a planar target experiment carried out on the OMEGA laser facility. The spectra generated by stimulated Raman scattering shows strong agreement with the experimental scattered light data. The model also indicates two-plasmon decay was the dominant instability in the experiment, which can not be verified from the currently available experimental data.

Contents

Abstract	ii
List of Figures	vi
Acknowledgements	xviii
Declaration	xix
Role Of The Author	xx
1 Introduction	1
1.1 Inertial Confinement Fusion	2
1.2 Laser Plasma Instabilities	6
1.2.1 Mitigation	7
1.2.2 Modelling	9
2 Laser-Plasma Instabilities	12
2.1 Plasma Physics	12
2.1.1 Plasma parameters	12
2.1.2 Waves	15
2.1.3 Inverse Bremsstrahlung	23
2.2 Laser-Plasma Instabilities	25
2.2.1 Convective and Absolute instabilities	27
2.2.2 Temporal Amplification	28
2.2.3 Convective Amplification	29
2.2.4 Stimulated Raman Scattering	35
2.2.5 Stimulated Brillouin Scattering	42

2.2.6	Two-Plasmon Decay	46
2.2.7	Langmuir Decay Instability and absolute instability saturation	53
3	Computational Methodology	55
3.1	A Single Ray	55
3.1.1	Wave Parameters	56
3.1.2	Growth and Saturation of a Single Instability	59
3.1.3	Pump depletion	64
3.1.4	Defining Reflectivity	66
3.1.5	Comparing Analytically and Numerically Integrated Convective Amplification	70
3.1.6	Nonlinear Saturation by the Langmuir Decay Instability	75
3.1.7	Multiple Absorption mechanisms	81
3.1.8	Scattered Light Diagnostics	84
3.2	Moving to Three Dimensions	85
3.2.1	Ray Tracing and Radiation-Hydrodynamics Codes	85
3.2.2	3D Mesh, multiple rays	86
3.2.3	Convergence tests	92
3.3	3D Problems + Fixes	94
3.3.1	Mesh effects on parameters	94
3.3.2	Failure to converge	98
4	Results	106
4.1	Pump Depletion Benchmark	106
4.2	Conical Target Experiment	108
4.3	The Chiron Simulation	111
4.3.1	Stimulated Raman Scattering	114
4.3.2	Two-Plasmon Decay	120
4.3.3	Excessive TPD Depletion	126
4.3.4	SBS	129
5	Conclusions	132
5.1	Future Work	135

CONTENTS

v

A Ion-acoustic Waves in Multi-ion Plasmas

138

List of References

142

List of Figures

1.1	Schematic depicting the process of inertial confinement fusion. The cryogenic target is irradiated such that the CH layer ablates and the ice and gas layers are compressed. Stagnation is reached at the point of maximum compression when the velocity of the imploding shell reaches zero. Its kinetic energy is converted into internal energy of the hotspot triggering the fusion reactions and a burn wave to propagate through the shell before disassembly. Taken from [6].	3
1.2	Fusion reactivity for a variety of candidate fuels.	4
1.3	The indirect drive (top left) and direct drive (top right) configurations of ICF. The below panels show the evolution of the capsule when imploding. Taken from [7].	5
1.4	(a) Schematic depicting spatial and temporal smoothing of a laser beam. Below are images of a typical focal spot (b) with no smoothing, (c) with spatial smoothing using a RPP and (d) with spatial smoothing and temporal smoothing from 1-D SSD. Reproduced from [8].	9
2.1	The electron, ion and total Debye length for a hydrogen plasma with $T_e = 2T_i = 3\text{keV}$ (standard ICF ablation plasma temperatures) as a function of normalised electron density, where n_c is the critical density — defined in Sec. 2.1.2.1. . . .	15

2.2 (a) Phase velocity (v_ϕ) for the two distinct IAW modes compared to using single ion dispersion relation. In red is the 99%H 1%Xe (partially ionised) plasma. The dotted and dashed lines are the fast (F subscript) and slow (S subscript) mode respectively taken from Williams' paper. The solid line is the phase velocity using the single ion theory (ave subscript) in Eq.(2.26) with ion parameters found by averaging over the two species. Exactly the same applies to the blue lines except these were for the 50%C 50%H plasma. (b) Phase velocities for both the CH and XeH plasmas as a function of Hydrogen fraction for a constant $T_i = T_e/2$. Blue lines are for XeH plasma and red for CH, with the solid lines representing both fast modes and dashed for the slow modes. These results use Williams' kinetic theory. The three dotted lines are the phase velocity of single ion plasmas using Eq.(2.26) for H, C and Xe. 20

2.3 Evolution of the Langdon correction factor, f , with (a) α and (b) laser intensity at differing electron temperatures. 24

2.4 Schematic for the mechanism behind stimulated Raman backscattering that can be applied to any three wave parametric instability by just exchanging waves and directions. Adapted from Montgomery [8]. 25

2.5 Plots demonstrating the change in SRS wavevectors through a density gradient. (a) and (b) show the scattered EMW and EPW respectively, while (c) shows the wavevector mismatch described in Eq.(2.58). The black line on all three plots shows the values of that parameter that satisfy the wavematching conditions. The coloured lines show the evolution of individual waves as an instability moves in space. A wave starts where the coloured line meets the black line as the wavematching conditions are met at the point an instability is seeded, but then the wave's wavevector moves further away from wavematching in space. 30

2.6 (a) The amplification of three Raman scattered rays (colour coded) seeded at different densities. The amplification of each ray is modelled based on Eq.(2.71) (solid line), Eq.(2.72) (dotted line), Eq.(2.69) (dashed line) and Eq.(2.70) (dashdotted line). The pump ray had intensity 10^{15} W/cm² and the plasma had a linear density gradient with a scale length at $n_c/4$ of 500 μm . Pump depletion was ignored for this simulation. (b) Profiles of the wavevector mismatch, κ , along each back propagating scattered wave. 35

2.7	Spectrum of stimulated Raman scattered light from a NIF experiment. Taken from [9].	36
2.8	(a) The measured wavelength of a scattered ray based on the density at which it is seeded for plasmas with 10 different, but constant profiles of electron temperature. (b) The wavelength of scattered rays that are resonant with the Raman instability at each density for a plasma with constant $T_e = 5$ keV. In blue is the wavelength of the ray at the point of resonance in the plasma that satisfies the wavematching conditions. In orange is the wavelength each of the resonant rays has become by the time they have propagated out of the plasma, and thus is the wavelength at which they would be measured.	37
2.9	Comparison of the WKB and non-WKB growth rates given in eqs. (2.74) and (2.75). The WKB approximation depends on gradients in the wave parameters, so is plotted as a function of plasma density scale length at a fixed density of $0.235n_c$. The WKB cut-off is defined by the expression in Eq. (2.38).	39
2.10	The density each Raman scattered ray is generated at relative to its own critical density (n_{cs}). The three x-axes are all parameters of the scattered ray. The bottom one is the frequency of the ray normalised to the laser frequency. The lower x-axis on top of the plot is the wavelength of the ray after it has left the plasma i.e. the wavelength that is measured. And the upper x-axis is the density the ray has been generated at relative to the laser's critical density (n_{c0}).	41
2.11	Four different thresholds for the definition of the strong coupling regime of SBS taken from [10–12]. (a) uses a 5×10^{14} W/cm ² intensity pump, and (b) a 2×10^{15} W/cm ² pump, both with 400 μ m scale lengths. The black dashed line is at 1, the value that must be crossed for each threshold.	43
2.12	Wave vector diagram for the two-plasmon decay instability. The solid lines are the three wave vectors, with the EPW subscript being dropped from the daughter waves. The dotted lines are the components of the daughter waves parallel and perpendicular to the pump.	46
2.13	(a) TPD wavenumbers normalised to the vacuum wavenumber of the pump in a 2.5 keV plasma. (b) The angles (as defined in Fig. 2.12) between the wavevectors in (a).	48

2.14	Wavelength of the two TPD EPWs as a function the density they are seeded at. The different colours are for two different pairs of EPWs formed in 1 and 4 keV plasmas. The solid lines are for the forward EPW and dashed for the backwards EPWs. The second y axis gives the measured wavelength of light that scatters off the EPWs. Experimentally, it is like that only EPWs at $n_e \geq 0.2n_c$ will be amplified due to the Landau cutoff.	49
2.15	(a) The linear TPD growth rate (Eq. (2.89)) and the threshold for the instability to overcome wave damping (Eq. (2.55)) in blue. In orange is $k_{epw1}\lambda_{De}$. This example has $T_e = 2.5$ keV and $I_0 = 10^{15}$ W/cm ² . (b) The Landau, collisional and total damping on the forward EPW from TPD using the same parameters as in (a).	50
2.16	(a) The linear TPD growth rate and the modified convective growth rate in a 2.5 keV plasma. (b) The linear, convective and non-WKB growth rates as a function of density scale length at $0.235 n_c$ and 2.5 keV. The black dashed line indicates the scale length below which the WKB approximation (Eq. (2.38)) fails and it becomes appropriate to use the non-WKB growth rate.	52
3.1	The difference between the two calculated values of k_{epw2} for each starting guess of $k_{epw1,x}$ at 4 different pairings of plasma density and temperature. Normalising $k_{epw1,x}$ to k_0 means the x axis is equivalent to α in Eq. (3.6). The black dashed line is at zero, where the values are equal.	58
3.2	Growth and saturation of Raman backscattered rays. Each ray is coloured by its wavelength when it leaves the plasma. The pink line is the spectrally integrated Raman backscattered light, found by summing all of the individual rays (including those not pictured here). This example was for $I_0 = 10^{15}$ W/cm ² , $T_e = 2.5$ keV and an exponential gradient with $L_n = 500\mu\text{m}$. Rays are saturated convectively and via LDI.	60
3.3	(a) Variation of reflectivity with number of scattered rays launched for 9 simulations with different combinations of pump intensity and density scale length. (b) The same data but with the reflectivity of each pair of parameters normalised to the pair's reflectivity when 10,000 rays are launched. These simulations used the bounded gain to convectively amplify rays and ignored the effects of pump depletion.	62

- 3.4 The convective amplification for rays seeded at each density for simulations with pump intensity $2 \times 10^{15} \text{ W/cm}^2$, scale length $500 \mu\text{m}$ and varying numbers of Raman scattered rays. 62
- 3.5 Simulation results considering only convective SRS along a single pump ray with intensity $1 \times 10^{15} \text{ W/cm}^2$ in a plasma with scale length $700 \mu\text{m}$. SRS growth rates are calculated from an initial constant intensity pump (top), allowing us to calculate saturated amplitudes of Raman scattered light rays (middle) and then deplete the initial pump by conserving action (bottom). The middle panel shows 50 selected Raman rays with their measured wavelengths colour coded as well as the total intensity of scattered light in pink. 65
- 3.6 Intensity profile of the pump ray before doing LPI calculations (iteration 0) and after each iteration of the LPI calculation. The profile converges to within 1% of the previous iteration after the fourth iteration. 66
- 3.7 The variation of the coefficient from Eq. (3.9) with plasma density. After a scattered EMW at some density has been Raman amplified the pump is depleted by the product of the scattered wave's intensity and the coefficient. 67
- 3.8 Scattered light data from the same simulation as Fig. 3.5. (a) The colour scaled lines represent the intensity of individual rays, with the colour matching their wavelength upon exiting the plasma. The pink is the total intensity of scattered light leaving the plasma at each point in space. The solid lines show the intensity if we account for parametric amplification due to SRS. The dashed lines represent the increase in the intensity of waves due to the changing group velocity as they travel through new densities. (b) The Raman spectrum when considering the initial intensity I_s^i and the final intensity I_s^f , i.e. the intensities directly after Raman amplification and after Raman amplification and exiting the plasma respectively. 69
- 3.9 Reflectivity defined as $\frac{I_s c}{I_0 v_{g,s}^i}$ for a single pump ray with differing initial intensity and plasma scale lengths. Each simulation had $T_e = 2.5 \text{ keV}$, a density upper bound of $0.23n_c$ and was run for 20 iterations or until convergence. They only modelled convective SRS amplified by (a) the Rosenbluth gain, (b) the bounded gain or (c) the integrated gain. 71

3.10 (a) Spectra from the simulation in Fig.3.9 with $I_0 = 5 \times 10^{15} \text{ W/cm}^2$ and $L_n = 600 \mu\text{m}$. (b) Pump profiles from the same simulation. 73

3.11 Reflectivities for the 3 gain methods along a single pump ray of differing intensity in plasmas with differing scale lengths and constant temperature of 2.5 keV. The only difference between these simulations and Fig. 3.9 is that those simulations had the LPI calculation iterated over up to 20 times to include the effects of pump depletion, whereas these completely ignored pump depletion by using a constant intensity pump throughout the plasma and not iterating. 74

3.12 (a) The absolute difference in bounded and integrated reflectivity from the simulations that ignore pump depletion in Fig.3.11. The red line forms the boundary between simulations that had a difference of more or less than $0.1I_0$. (b) The spectra from the simulation with pump intensity $5 \times 10^{15} \text{ W/cm}^2$ and density scale length $600 \mu\text{m}$. The dotted blue line changes the integration limits of the Rosenbluth gain and is explained in the text. 74

3.13 Spectra from two simulations with (a) $I_0 = 2 \times 10^{15} \text{ W/cm}^2$ and $L_n = 500 \mu\text{m}$, and (b) $I_0 = 4 \times 10^{15} \text{ W/cm}^2$ and $L_n = 400 \mu\text{m}$. Both have switched off the effects of pump depletion for clarity. Each simulation shows the reflectivity for rays based on whether the saturation is convective or LDI, with the solid lines depicting the dominant saturation mechanism at each wavelength. The reflectivity is defined here as I_s^i/I_0 using a scattered ray's initial intensity after Raman amplification rather than final intensity after exiting the plasma because these are the values the model actually compares when selecting the dominant saturation mechanism. 76

3.14 Raman scattered light data from simulations to investigate the effect of our LDI model on the SRS model utilising the bounded gain for convective saturation. Each simulation was for a single pump ray in a 2.5 keV plasma with a maximum density of $0.25 n_c$ and includes the effects of pump depletion. (a) Allowed only convective saturation using the bounded gain while (b) used both this and saturation via the Langmuir decay instability. (c) Shows results from the same simulations as in (b) but this time plotting the total intensity of scattered light exiting the plasma without normalising to pump intensity. 79

- 3.15 Raman reflectivities from simulations using the (a) bounded and (b) integrated gain methods for convective gain combined with Langmuir decay saturation as well as (c) the difference between the two sets of data. The simulations ran for 1 iteration only, ignoring the effect of pump depletion. 79
- 3.16 (a) Reflectivity variation with number of scattered rays launched for 9 different combinations of pump intensity and density scale length. Each pair of parameters' reflectivity is normalised to the pair's reflectivity when 10,000 rays are launched. The simulations presented here are identical to those in Fig. 3.3 except for the inclusion of the Langmuir decay instability. (b) The convective amplification for rays seeded at each density for simulations with pump intensity $2 \times 10^{15} \text{ W/cm}^2$, scale length $500 \mu\text{m}$ and varying numbers of Raman scattered rays. These simulations are identical to those presented in Fig. 3.4 except for the inclusion of LDI. 81
- 3.17 (a) Fraction of the pump absorbed by each instability in simulations with inverse bremsstrahlung turned off. The solid lines are from a simulation that ignored pump depletion (lasted one iteration) while the dashed line is a simulation that included pump depletion (iterated until pump profile convergence). (b) The fraction of the pump absorbed by each absorption mechanism for three different implementations of the IB model after 1 iteration of LPI calculations and after converging iterations. For example, the blue markers show the fraction of the pump absorbed by the pump when the pre-LPI method was implemented. All simulations had $I_0 = 2 \times 10^{15} \text{ W/cm}^2$, $L_n = 500 \mu\text{m}$ and $T_e = 2.5 \text{ keV}$ and modelled a ray propagating in one direction through a plasma extending up to its critical density. 83
- 3.18 The fraction of laser energy absorbed by each mechanism for simulations with increasing number of rays making up the beam. This simulation was of a plasma with a cartesian mesh with a maximum density of $0.75 n_c$ and scale length of 400 nm . The beam propagates parallel to the density gradient in the x direction while the mesh only has one cell in the y and z directions. 92

- 3.19 The fraction of laser energy absorbed by each mechanism for simulations with a single beam made up of increasing numbers of rays incident upon a spherical plasma with a scale length $\sim 500 \mu\text{m}$ at $n_c/4$. (a) Shows the results after only running the code for 1 iteration, whilst (b) is the results after iterating until pump profile convergence. 94
- 3.20 (a) Intensity profile of an individual ray up to its turning point and the intensity that ray sees in each cell it passes through as described by Eq.(3.17). This simulation had ~ 150 rays in the cells at the outermost layer of the plasma. (b) Schematic of a ray passing through three cells with significantly different path lengths in each cell as it propagates into a spherical mesh. The scale here is exaggerated for illustrative purposes, most simulations would have higher angular and radial resolution. 95
- 3.21 (a) Density profile of a ray propagating into and out of a spherical plasma mesh which has a smooth density profile that is constant on each radial layer of the mesh. The inset zooms in on a few points of the profile where the ray steps angularly through cells. (b) Density gradient of the same ray with an inset again zoomed in on a similar region to examine behaviour when the ray steps angularly. 97
- 3.22 (a) The density profile (top) and its first (middle) and second (bottom) order derivatives of the ray from Fig.3.21 zoomed in to examine one instance where the cell steps angularly and the gradient fluctuates sharply. The orange line labelled fixed shows the gradient after interpolating when an angular step is detected. (b) The two derivatives of the density profile over the entire course of the ray, with red crosses highlighting every time the ray steps angularly and the orange profile again showing the fixed gradient once these fluctuations have been removed. 97
- 3.23 Pump profiles from simulations using (a) $4.3 \times 10^{15} \text{ W/cm}^2$ and (b) $5 \times 10^{15} \text{ W/cm}^2$ intensities, in a homogeneous plasma with $n_e = 0.1 n_c$, $T_e = 3.5 \text{ keV}$ and a seed amplitude $I_{seed} = 10^{-4} I_0$. Neither simulation converges so the pump profiles after the first 20 iterations over the LPI calculation are presented. . . . 99

3.24	Results from a simulation using the same set-up as in Fig. 3.23a. Now the pump profiles after iteration 19 and 20 are averaged to create a new guess at the answer in between them, and the simulation is continued from here for 60 iterations in total.	100
3.25	The pump profiles from the two respective simulations in Fig. 3.23 but now with an automated method that continually guesses in between the diverging pumps to search for the saddle point-like solution.	101
3.26	Various absorption profiles of a single ray from a simulation after (a) 9 iterations and (b) 10 iterations. This ray was one out of 900 rays making up a laser beam with intensity $\sim 3 \times 10^{15}$ W/cm ² . Each line represents the amount the pump is depleted by by an individual absorption mechanism, all of which are then combined to create the "total" depleted pump profile. TPD was not included in this simulation for simplicity.	103
3.27	Various absorption profiles of a single ray from a simulation after (a) 9 iterations and (b) 10 iterations. This ray was one out of 900 rays making up a laser beam with intensity $\sim 3 \times 10^{15}$ W/cm ² . Each line represents the amount the pump is depleted by by an individual absorption mechanism, all of which are then combined to create the "total" depleted pump profile. TPD was not included in this simulation for simplicity.	104
4.1	Various (a) SRS and (b) SBS reflectivities in a homogeneous hydrogen plasma with $n_e = 0.1n_c$, $T_e = 3.5$ keV and $T_i = 1.6$ keV. FLAME is the EM code developed by Hao <i>et al.</i> [13] and the data series labelled analytic Hao and analytic Arun are their solutions and mine to Tang's analytic reflectivity (eq. (4.1)). . .	107
4.2	The open cone target design with one beam shown for each cone of 5 beams at polar angles of 23°, 48° and 62° out of the low intensity beams shown. The 48° and 62° were turned on initially at lower intensity to generate the long scale length plasma and then the high intensity 23° beams were used to investigate laser-plasma interactions. Two of the 23° beams had FABS detectors behind them, one to collect light generated by SRS and TPD and the other for SBS and CBET generated light.	108

4.3	Backscattered light captured by the FABS diagnostic, which was positioned behind one of the 23° beams. White lines show laser pulse. The detector shown in (a) was sensitive to scattered light generated by SRS and TPD, while (b) was sensitive to SBS scattered light.	110
4.4	Fraction of laser energy absorbed by IB, each LPI or not absorbed after applying Chiron to the hydrodynamic outputs from h2d simulations of experiment 88980. Two different simulation set ups were tested: (a) used 3 laser beams with one beam per cone of beams at constant polar angle, (b) simulated the propagation of all 15 beams. The total power used in each simulation was the same.	112
4.5	Fraction of laser power absorbed by IB and the summed total of all three LPI for the three beam simulation (blue) and the all beam simulation (orange).	112
4.6	The same data presented in fig. 4.4 but with a reduced limit on the y axis to allow a more detailed view of the LPI interplay. (a) is from the three beam simulation, (b) is from the all beam simulation.	113
4.7	Density profiles of all 1600 rays launched by the high intensity, 23° beam in the 3 beam simulation at (a) 1.6 ns, (b) 1.8 ns and (c) 2.0 ns.	114
4.8	Synthetic FABS image created by collecting all of the Raman scattered light generated by the high intensity beam from the three beam simulation. The correction to power due to varying group velocity described in sec.3.1.4 is applied to the scattered rays. (a) collisionally absorbs the scattered rays as they propagate out of the plasma, whilst (b) does not.	116
4.9	Simulated Raman scattered light diagnostics created from the high intensity, 23° beams in the all beams simulation including collisional absorption and the group velocity correction on the scattered rays. (a) For the beam that had the FABS in its port in the experiment (b) The simulated Raman light summed from all of the five high intensity beams and using the same colour scale as the three beam simulation FABS images in fig. 4.8.	117

4.10	SRS threshold parameter as formulated by Scott <i>et al.</i> for the two sets of beams and the scale lengths they considered using eq. (4.4). The alternative threshold I provide for the low intensity beams uses Liu's expression which is a factor of two smaller than eq. (4.4) and a wider range of scale lengths based on the distribution of values seen in the ray tracer. The intensities used are the maximum overlapped intensities of 0.9 and 2.5×10^{15} W/cm ² for the low and high beams respectively.	119
4.11	Power of scattered light generated by TPD from (a) 62°, (b) 48° and (c) the high intensity 23° degree beams in the three beam simulation. Collisional absorption was applied to the light rays. The initial powers of the scattered rays were assumed to have powers equal to the EPWs they scatter off, as is explained in Sec. 3.1.8.	121
4.12	Wavelength integrated spectra for TPD scattered light from each laser beam from the spectra presented in fig. 4.11. The solid and dashed lines are from the shorter and longer wavelength scattered light lobes respectively.	123
4.13	Power of scattered light generated by TPD from (a) 62°, (b) 48° and (c) the high intensity 23° degree beams in the three beam simulation. This is from the same simulation as in fig. 4.11 except it ignores collisional absorption on the scattered light.	124
4.14	Total TPD scattered light from the all beam simulation created by summing each of the five beams in (a) the 62° cone of beams, (b) the 48° cone and (c) the 23° cone.	125
4.15	Wavelength integrated spectra for TPD scattered light from each laser beam from the spectra presented in fig. 4.14. This light is only from the lower wavelength lobe of scattered light (<700 nm).	126
4.16	TPD scattered light from one of the five (a) 62° beams and (b) 48° beams from the all beam simulation.	126

4.17 (a) Daughter wave amplitudes for TPD occurring along the single pump ray that contributes most to the largest spike at 0.8 ns in the beam presented in fig.4.16b. LDI indicates these amplitudes are only considering saturation via LDI, whereas those without it have also considered convective saturation. (b) The conversion factor and the two contributing variables as defined in the text to move between intensity and power for the pump ray considered in (a). . . . 128

4.18 Fraction of laser energy absorbed by SBS in the two simulations 130

4.19 (a) Spectrally integrated signal from the experimental FABS data in fig.4.3b. (b) Spectrally integrated scattered light from SBS in the three beam simulation.131

A.1 Phase velocity (v_ϕ) of the fast (red) and slow (blue) ion-acoustic mode from a 99%H 1%Xe plasma. The solid lines are the results presented by Williams *et al.* [14], and the pluses and minuses are for our results using their expression with either the plus or minus sign. Velocities are normalised to $v_r = \sqrt{k_B T_e / m_p}$.139

A.2 (a) Phase velocity (v_ϕ) of the fast (red) and slow (blue) ion acoustic modes as well as the thermal velocities (v_{th}) of the light (magenta) and heavy (cyan) ion species. The solid lines are results taken from Williams' paper, while the crosses and pluses represent my results using the kinetic expressions of Williams and Vu respectively. Velocities are normalised to $v_r = \sqrt{k_B T_e / m_p}$. (b) Landau damping rate (ν) of the two ion modes again comparing results presented by Williams with my model using kinetic theory from Williams and Vu using the same legend as in (a). Damping rate is normalised to the wave's frequency. . . . 140

A.3 (a) The phase velocities (v_ϕ) of the two IAWs and the thermal velocities in a 50% C 50%H plasma. The phase velocity data was taken directly from Williams *et al.* [14]. (b) The Landau damping rates (ν) for the fast IAW (red) and slow IAW (blue) in a CH plasma. The data from Williams paper are the solid lines and our solutions using Williams' and Vu's theory are the crosses and pluses respectively. 140

A.4 Phase velocities (v_ϕ) for the (a) fast and (b) slow ion mode in a CH plasma. The solid lines are the results presented in Williams paper from his new kinetic (blue) and fluid (orange) theory. Our results using Williams' and Vu's knietic theory based on the unfulfilled condition in eq.(A.1) are the blue crosses and pluses respectively. The orange pluses are our results using Williams' fluid theory.141

Acknowledgments

My first thanks must go to my Mum. I am the person that I am because of her. I am in a position fortunate enough to carry out a PhD because of the opportunity she has provided for me. To Kiran and Maya, I find it surprising when I see other siblings with comparatively distant relationships, because my own experience has thankfully been so different to that. You guys are my day ones and will continue to be my best friends for the rest of my life. Keating n the meme kidz 'til I die.

Thanks must go to colleagues at York and the CLF, Adam, Calum and Kevin, who have provided camaraderie, academic support and irreverent conversations about physics. Special thanks go to Matt, who first welcomed me to the YPI and has offered regular conversation and advice about LPI and PhD life in general, and Alessandro, who offered computational support, theoretical advice and ran many simulations for me.

To my friends from various parts of my life: Katya, Matt, Kieran, Gary, Michael, James, Hannah, Amy, Marianne, Charlie, Noah, Joey, Sam, Isla, Yorick, Laura, Cyd and Rob (and others). You may not have contributed directly to this work, but you have shaped me into the person that has undertaken this work. I am fortunate I have such a wonderful group of people whose friendship I can revel in and always count upon.

To my housemates of the last year: Kasper, Reuben, Jago, Juan and Callum. Twelve months ago we didn't know each other, but since then you have provided company, conversation and recreation to keep me sane in this most difficult of years. Of course, the many members of Footie Mob Mk. II also deserve a mention as regularly playing good vibes football again has been a great outlet for me during this last slog of a year (and has reminded me how much fun football is when people aren't shouting at each other).

My supervisors undoubtedly have played the biggest role in getting this project over the line and have developed me as a researcher along the way. Thank you Nigel and Robbie for your unwavering support, guidance and enduring patience during my PhD.

Declaration

I declare that this thesis is a presentation of original work and I am the sole author. This work has not previously been presented for a degree or other qualification at this University or elsewhere. All sources are acknowledged as references.

Role Of The Author

The bulk of the work presented in chapters 2, 3 and 4 is original work carried out by the author. The derivations presented for equations 2.53 and 2.69 are well known and in this work the formulation provided by Michel [12] is followed. The ray tracer described in chapter 3 provided the backbone for the LPI simulations carried out in chapter 4 was developed by R. H. H. Scott. The author was not involved in the experiment that is described and simulated in chapter 4. The scattered light data from the experiment presented in fig. 4.3 was analysed by K. Glize.

Chapter 1

Introduction

Climate change is the most urgent existential threat to humanity and the entire natural world. It has already begun to cause and will continue causing: loss of food sources through decrease in arability of land and losses in pollinators and fisheries; increases in extreme weather events including floods and droughts; forced mass migrations as regions of the world become uninhabitably warm, bereft of food sources, or underwater as ocean levels rise; mass extinctions as animals and plants are unable to adapt to the rate of such drastic changes; and reduction in quality of human health as a result of air pollution, rising temperatures, worse nutrition and diseases.

Global electricity consumption increases almost every year and is forecast to continue doing so with the development of emerging economies and electrification replacing fossil fuel technologies [15]. Electricity production has always been dominated by the fossil fuels coal and gas, which, with oil, contributed over 60% of the global electricity production in 2023 [16].

Sustainable and renewable energy sources, particularly solar and wind energy, have an increasing role in electricity generation and they offer much lower greenhouse gas emissions per unit of energy generated. Solar and wind energy require energy storage to account for their intermittent supply, which may not align with the energy demand. Current energy storage capabilities are not at the required capacity needed to power the world solely by renewables, although heavy investment is being made, particularly in lithium-ion batteries [17]. Scaling renewable energy supply to a global supply using the current technology also may face further challenges due to sustainability concerns over the sourcing of certain metals and rare-earth elements [18].

Nuclear fission can provide a regular supply of energy and is similarly sustainable in that it produces minimal greenhouse gas emissions. However, the limited amount of uranium that can be mined on Earth and the radioactive fission waste products with half lives that can be over a thousand years mean it is not long term solution to the humanity's energy demands.

Nuclear fusion promises an alternate green energy source that can surpass renewables and fission by solving their deficiencies. A working fusion reactor would run on abundant fuel, would be able to meet variable energy demands and should produce shorter lived radioactive waste. Unfortunately, the conditional tense in that last sentence is essential as a working fusion reactor that can put electricity onto an electrical grid is still many years from existing. Clearly, fusion alone is not the immediate answer that is needed for mitigating the effects of climate change. Significant investment is needed in renewables and fission to reduce dependence on fossil fuels in the short term, but the scalability of fusion to a potential global renewable energy source make it a necessary target for continued investment, research and development.

Despite a pervasive feeling of futility when one considers the scale of the challenge presented by climate change, there is cause for optimism with regards to nuclear fusion at least. The UK is investing heavily in a new tokamak, STEP, which it sees as an alternative to the ITER project, using modern physics and technology to push towards achieving fusion via magnetic confinement. Many private companies have also appeared with unique approaches to achieving fusion, hoping to accelerate advancement towards this goal. Most significantly, in December 2022 the National Ignition Facility (NIF) achieved net energy gain for the first time using laser driven inertial confinement fusion — a monumental milestone for the field [19].

1.1 Inertial Confinement Fusion

To fuse positively charged nuclei together, they must have sufficient energy to overcome electrostatic repulsion and be close enough together for attraction by the strong nuclear force bind them into a heavier nucleus and other products, while releasing energy. The energetic particles created can in turn heat the fuel, generating more fusion reactions in a self-sustaining process.

The challenge for nuclear fusion as an energy source is confining the fuel for long enough





Figure 1.1: Schematic depicting the process of inertial confinement fusion. The cryogenic target is irradiated such that the CH layer ablates and the ice and gas layers are compressed. Stagnation is reached at the point of maximum compression when the velocity of the imploding shell reaches zero. Its kinetic energy is converted into internal energy of the hotspot triggering the fusion reactions and a burn wave to propagate through the shell before disassembly. Taken from [6].

to achieve net energy gain. Magnetic confinement fusion uses large, superconducting magnets wrapped around a vessel containing the plasma to generate fields and currents in the plasma which work to confine it. Inertial confinement fusion (ICF), the focus of this project, uses the fuel's own inertia to confine it over nanosecond timescales.

ICF uses a millimetre scale spherical capsule that has three main layers: a thin outer layer typically made of a plastic like CH; a thick layer of cryogenically cooled fusion fuel, normally deuterium and tritium; and a gas core of the same fuel. Energy is deposited in the outer layer of a spherical capsule causing it to rapidly ablate outwards. Ablation pressure drives the compression of the ice and gas layers by creating shock waves which propagate into the capsule. The gas in the centre of the capsule can reach temperatures up to 10^8 K or 10 keV, hotter than the core of the Sun. This much energy is sufficient to trigger nuclear fusion reactions in the core that, if successful, will trigger more reactions, burning through the ice layer in the capsule.

The reactivity of a fuel is the product of the fusion cross section and the ion velocity thermally averaged over the velocity distribution function of the plasma. Figure 1.2 shows the reactivity of possible fuels averaged over a Maxwell-Boltzmann distribution of velocities. Given that current ICF implosion designs achieve temperatures up to 5 keV, deuterium and tritium (DT) is the obvious candidate for fusion. The DT reaction,



creates a Helium nucleus, or alpha particle, and a highly energetic neutron. Deuterium is a




Figure 1.2: Fusion reactivity for a variety of candidate fuels.

stable isotope of hydrogen that is abundant in the form of heavy water in the world's oceans. Tritium is an unstable isotope of hydrogen with a half life of ~ 12 years and will need to be created within reactors by using the fusion neutron to activate lithium-6, which will decay into tritium.

The uncharged fusion neutrons carry 80% of the energy from the reaction and will also be used to generate electricity via boiling water and turning a turbine. For fusion to be self-sustaining the alpha particles must deposit their energy in the fuel to drive more fusion reactions. The recent experiments on the NIF that have achieved net energy gain because of significant alpha heating have had hotspot areal densities $\rho R > 0.4 \text{ g/cm}^2$ [19].

If the hotspot has successfully ignited it will launch a nuclear burn wave through the dense fuel layer of the capsule, where the bulk of the fusion reactions will take place. This fuel must have been compressed to high areal densities to maximise the fusion reactions that can occur before the fuel disassembles. For maximal compression of the dense fuel layer, it is kept at $\sim 18 \text{ K}$ [20], keeping the pressure in this layer as close as possible to the Fermi degeneracy pressure by minimising thermal pressure. Compression is achieved by a series of shocks deliberately launched through the capsule by specific shaping of the laser pulse. The shock waves are designed to keep thermal pressures low, but do increase the entropy of the system, thus increasing the Fermi pressure. A compromise is made by launching several weak shocks rather than one strong shock or pressure waves to achieve compression while limiting the increase in entropy and keeping the fuel cold.

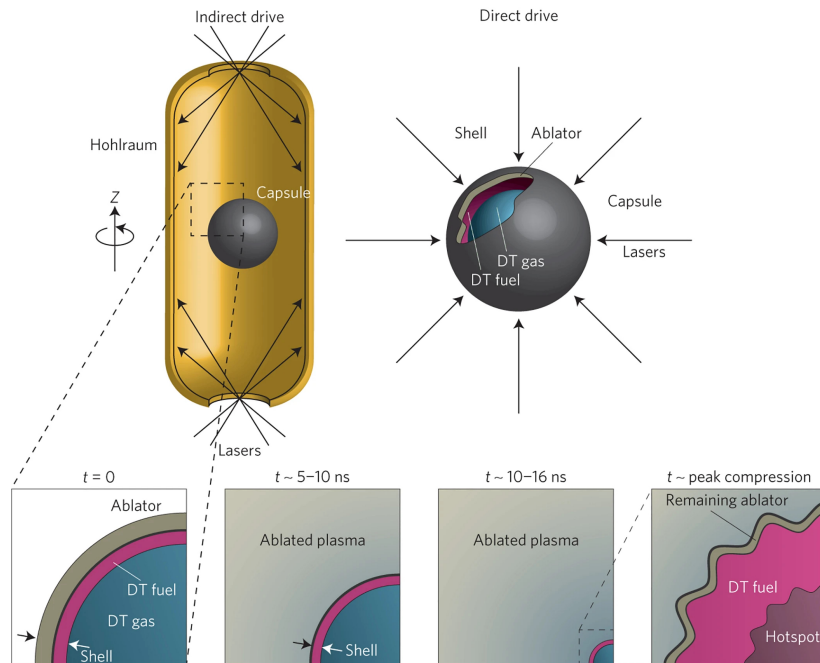


Figure 1.3: The indirect drive (top left) and direct drive (top right) configurations of ICF. The below panels show the evolution of the capsule when imploding. Taken from [7].

Inertial confinement fusion has taken a long time to reach the experiments which produce energy gain. The NIF has taken about twelve years of constant operation for our understanding of the field and its many obstacles to reach this point. ICF is hindered by asphericities in the of the imploding capsule that can limit maximal compression. Hydrodynamic instabilities, like the Raleigh-Taylor instability, can cause mixing between layers of the shell, cooling the hotspot, reducing the fuel purity and even breaking up the shell entirely [21]. The efficiency of compression is also reduced by highly energetic particles from the ablator depositing energy in the fuel layer. Laser energy coupling into the plasma is reduced by the presence of of laser-plasma instabilities (LPI) that scatter laser light and drive large amplitude plasma waves.

The two leading approaches to laser driven ICF, direct [22] and indirect drive [23], are shown schematically in Fig. 1.3. At the NIF they pursue indirect drive; the lasers illuminate the inside of a high Z (normally gold plating on the inside of depleted uranium [24]) cavity called a hohlraum, which emits x-rays that couple energy into the capsule, driving the implosion. The OMEGA laser facility is configured for direct drive ICF; the lasers are incident upon the fuel capsule, coupling their energy directly into the ablator to drive the implosion. There is currently no ignition scale direct drive facility in the world. Other

advanced ignition schemes exist like fast ignition [25] and shock ignition [26], which decouple the compression phase and the heating of the fuel. They initially promised better heating efficiencies and larger yields but interest has recently cooled as problems with fast electron coupling in fast ignition and laser-plasma instabilities in shock ignition have reduced their viability [22].

Direct drive is prone to non-uniformities in the capsule surface and laser intensity distribution seeding hydrodynamic and laser-plasma instabilities in the capsule. The more homogeneous x-ray driver of indirect drive reduces the impact of these non-uniformities and allows better absorption of x-rays that are incident upon the capsule. Indirect drive suffers from inefficiency in the indirect coupling of laser energy into the capsule. There is a small loss of efficiency in converting the laser to x-rays and a much larger loss caused by x-rays in the hohlraum not being incident upon the capsule. It also has higher costs because of the expensive materials in the hohlraum, which sets higher requirements on gain for an energy plant. The hohlraum materials are normally high-Z so are subject to neutron activation, so would present a radiological hazard in a fusion energy plant

1.2 Laser Plasma Instabilities

The five main LPI that have affected ICF are stimulated Raman scattering (SRS), two-plasmon decay (TPD), stimulated Brillouin scattering (SBS), cross-beam energy transfer (CBET) and filamentation. These are three wave processes where an initial pump wave, the laser in the case of ICF, decays into two daughter waves. The daughter waves deplete energy from the laser as they grow unstably via a feedback loop created by the beating of the laser wave with the daughter waves. In addition to reducing laser coupling, one of the daughter waves of SRS and TPD is an electron plasma wave (EPW), which, when Landau damped, can accelerate electrons to suprathermal velocities, which in turn can deposit their kinetic energy in the cold fuel causing preheat. The mechanism behind LPI and the theory that dictates their growth is discussed in Sec. 2.2.

Experimentally, the presence of LPI is detected by scattered light amplified by SRS and SBS or light generated by secondary processes involving the TPD EPWs. The presence of hot electrons generated by SRS and TPD can also be inferred by measuring the

bremsstrahlung emission from the electrons decelerating in the plasma with hard x-ray detectors. Additionally, embedded layers containing a high-Z material can be used where hot electron bombardment induces characteristic K_α line emission.

Laser-plasma instabilities are present in both direct and indirect drive ICF configurations. In direct drive, they occur in the ablation plasma, preventing laser energy reaching higher densities where it is more likely to be absorbed. In indirect drive, LPI occur in the low density plasma in the hohlraum and the windows at its entrance hole, preventing the laser energy from reaching the hohlraum walls and being converted to x-rays.

Early experiments on the NIF were marred by the presence of LPI. SRS was found to be absorbing 15% of the drive energy before it could reach the hohlraum wall, contributing to flux asymmetries that were causing factor eight decreases in yield [27]. CBET is a form of SBS which redirects a laser beam's energy into the path of a different overlapping beam, is very prevalent in the hohlraum laser entrance hole, but is now used to deliberately transfer power between the outer and inner cones to get more energy to the hohlraum waist [28].

CBET has also been a prevalent problem in direct drive experiments at the OMEGA laser facility. At the overlapping beam edges, energy can be redirected away from the target and potentially cause optics damage as it propagates back up the path of a beam. Two-plasmon decay is often seen as the dominant LPI on OMEGA [22], in comparison to SRS and SBS on the NIF. This is because TPD occurs more readily as a multi-beam instability, and the lower single beam intensity on the sub-ignition scale OMEGA facility requires overlapping beams to reach the necessary intensities to drive LPI. The lower energy of the OMEGA laser also creates shorter density scale-length plasmas, which suppresses the spatially amplified SRS and SBS more than the locally amplified TPD. It is likely they will all be present in an ignition scale direct drive facility.

1.2.1 Mitigation

The product of laser intensity and wavelength squared, $I\lambda^2$, is a parameter used to determine the threat posed by LPI. Instability growth rate is proportional to $I\lambda^2$, so by reducing this parameter LPI can be mitigated. Most ICF facilities use frequency tripled laser light because short wavelength light is collisionally absorbed more efficiently [29]. These shorter wavelengths are also beneficial in reducing LPI.

The interference between the monochromatic waves at varying phases within a laser

beam cause its focal spot to have a structured intensity distribution with large regions of increased intensity as depicted in Fig. 1.4b. The local intensity maxima can have much larger intensities than the beam average and its in these regions where LPI can be grow much more rapidly than may have been expected. Laser speckle also creates an imprint on the capsule where the local maxima are driven harder, which is susceptible to seeding Rayleigh-Taylor growth as well.

A number of techniques to smooth the intensity distribution of the laser have been developed, which have been shown to significantly reduce the impact of LPI [30, 31]. A phase plate, as shown in Fig. 1.4a, is an optic with varying thickness which shifts the phase of light an amount corresponding to wave's path length in the phase plate. Applying differing phase shifts across the beam breaks up the large scale non-uniformities seen in Fig. 1.4b and instead creates a speckled intensity profile pictured in Fig. 1.4c. Speckles are still prone to problems with imprint for hydrodynamic instabilities and laser-plasma instabilities due to the many intensity maxima across the focal spot.

Smoothing by spectral dispersion SSD is a temporal smoothing method, which varies the speckle intensity pattern on a timescale faster than the plasma's hydrodynamic response timescale [32]. A small bandwidth is applied to the laser beam which is then spectrally dispersed by a grating so that each element of the phase plate is irradiated by a different frequency. The relative phase between "beamlets" from different phase-plate elements will then vary in time according to their frequency differences. The focal spot in 1.4d has had one dimensional SSD applied to it which is apparent by the horizontal structures across it. Adding a second grating to spectrally disperse beam orthogonally to the first grating allows further smoothing of the beam profile and is an option on lasers at OMEGA, but not the NIF [33, 34]. By moving the speckle faster than the hydrodynamic response of the plasma, SSD was intended to prevent laser imprint from seeding hydrodynamic instabilities rather than LPI. The growth rates of the laser-plasma instabilities are faster than the characteristic time of SSD, so it can not directly prevent them. However, it was found that by preventing the ion density modulation that results from filamentation, SSD results in reduced self-focusing of laser light to high intensities within the plasma, thus indirectly reducing LPI growth [35].

Polarisation smoothing has also been shown to significantly reduce levels of SRS and SBS when used in tandem with SSD and phase plates [36]. By placing a wedge shaped optic

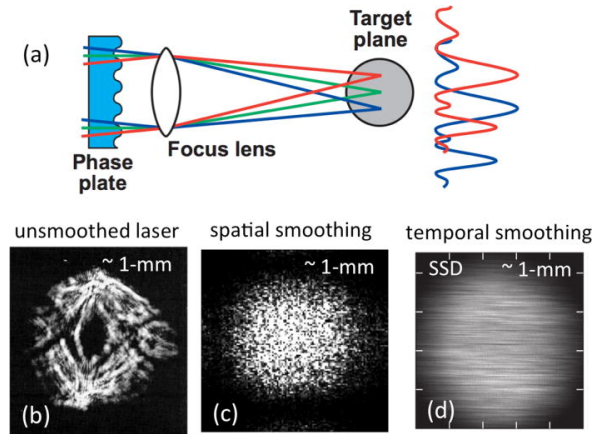


Figure 1.4: (a) Schematic depicting spatial and temporal smoothing of a laser beam. Below are images of a typical focal spot (b) with no smoothing, (c) with spatial smoothing using a RPP and (d) with spatial smoothing and temporal smoothing from 1-D SSD. Reproduced from [8].

before the phase plate, the differently polarised light is deflected by different amounts, thus creating spatially separated speckle patterns on the target [37]. To maximise the effect of polarisation smoothing the spatial displacement of speckles should be larger than the spectral dispersion to smooth out non-uniformities that the SSD failed to address [38].

1.2.2 Modelling

Being able to accurately model ICF is essential for designing experiments given the large costs associated with taking shots and the limited number of shots that can be conducted on the few reactors in the world. LPI modelling is also important in interpreting experimental results, especially given the limitations on LPI diagnoses. Different types of simulation codes allow different ways to examine LPI behaviour. Particle-in-cell (PIC) codes [39] offer a chance to examine details over very short timescales. Spencer *et al.* [40], for example, investigated inflation of SRS growth caused by kinetic effects over picosecond scales.

At the Laboratory for Laser Energetics (LLE), which houses the OMEGA laser facility, they have developed the code laser plasma simulation environment (LPSE) [1, 41]. It can either be used to simulate TPD or it can look at how SRS and SBS evolve together, including CBET, which is a significant problem in direct drive ICF on OMEGA [42]. LPSE is a hybrid code including a time-enveloped wave solver and particle tracking algorithm, but does not model the hydrodynamic response of the plasma to the LPI growth and laser energy deposition. It instead take plasma parameters from a radiation-hydrodynamics simulation

and use those to run a simulation.

Radiation-hydrodynamics codes (see Sec. 3.2.1) model the hydrodynamic evolution of plasma in response to the deposition of energy by electromagnetic radiation. This makes them ideal type of simulation code for modelling full ICF implosions. The plasma dynamics are modelled using fluid equations, but there is no inclusion of the wave or particle physics needed for LPI to emerge from. Incorporating methods to model LPI with hydrocodes is a key challenge in the development of more accurate ICF modelling that has had many attempts of varying approach. LPSE is too computationally expensive to run in tandem with a hydrocode and model how LPI evolve over a multiple nanosecond implosion [43].

CBET models developed at LLE to run in-line with their radiation-hydrodynamics code DRACO rely on ad-hoc multipliers to bring the simulation results close to experiments [44, 45]. There is a desire among many in the community to move away from needing multipliers and towards fully predictive, physics-based models for LPI.

At the Lawrence Livermore National Laboratory (LLNL), where the NIF is housed, they have the wave based laser-plasma interaction model pf3D [46] that has SRS and SBS models including both convective and absolute saturation. Strozzi [2] developed the 1D code DEplete which solves linear equations for steady-state levels of convective SRS and SBS, including nonlinear effects of pump depletion and physics based models for seed amplitudes. This is also used at LLNL, but is a ray based code and can offer faster results than pf3D. Assumptions included in DEplete mean it does not include modelling absolute SRS. Between pf3D, DEplete and the other in-house LPI codes at LLNL there does not seem to be any integrated TPD model [47].

At the University of Warwick, Angus [3] has developed a model for SRS with ray tracing and theory based on a work by Debayle *et al* [48], and further enhancing its predictive capabilities by using Gaussian processes [49]. At France's Commissariat à l'énergie atomique et aux énergies alternatives (CEA), Bénisti *et al.* [4, 50] have recently developed a predictive, physics based model for SRS that runs in-line with their in-house hydrocode, TROLL. This model is particularly notable for its ability to reproduce nonlinear kinetic effects seen in an experiment that are often not included as part of in-line LPI models.

This thesis details my contribution to this field with the LPI model called Chiron. This model is designed to run in-line with ray tracing packages as part of radiation-hydrodynamics codes. The long term plan for the project is for the model to be coupled with a radiation-

hydrodynamics so that LPI levels can be calculated during an ICF implosion simulation and these results fed back into the simulation. Currently, the model is not coupled with a radiation-hydrodynamics code so takes plasma parameters from a certain time during a hydrodynamic simulation to calculate LPI levels at that time.

To run in tandem with an implosion simulation, Chiron needs to be computationally efficient enough to not significantly increase the simulation's runtime. The LPI models are designed in the steady-state regime using analytically integrated expressions for the saturated amplitude of waves. Models for SRS, TPD and SBS (but not yet CBET) are combined with collisional absorption in Chiron with the aim to capture their interplay and how absorption by one mechanism can reduce other absorption mechanisms further along a ray.

Chapter 2

Laser-Plasma Instabilities

In this chapter the theory that underpins the laser-plasma instability models is presented. It begins with general plasma physics in Sec. 2.1 that is necessary for the laser-plasma instability theory. The bulk of the chapter is then in Sec. 2.2, which opens with general instability theory before going through the theory and behaviour of the three instabilities being modelled in detail.

2.1 Plasma Physics

Plasma, often called the fourth state of matter, can be thought of approximately as an ionised gas. It is a collection of charged particles whose dynamics are influenced by the electric and magnetic fields that can result from the motion of these particles allowing interaction and thus collective effects over much larger length scales and number of particles than would be possible in a gas.

2.1.1 Plasma parameters

The coronal plasmas of ICF capsules can be described as “quasi-neutral” or having no net charge. Displacing a slab of electrons from its quasi-neutral position by a small distance will create an electrostatic field that drives the electrons back to its initial position to re-establish neutrality. The electrons return to their initial position with a non-zero velocity and overshoot, creating another restoring electrostatic field as they oscillate at the (electron) plasma frequency:

$$\omega_{pe} = \left(\frac{n_e e^2}{m_e \epsilon_0} \right)^{1/2} \quad (2.1)$$

where n_e is the electron number density and m_e the electron mass. There is similarly an ion plasma frequency for the rate at which ions oscillate:

$$\omega_{pi} = \left(\frac{n_i Z e^2}{m_i \epsilon_0} \right)^{1/2}, \quad (2.2)$$

here Z is the charge of the ions, and subscript i refers to the ion number density and mass. Both of these descriptions are for cold particles, i.e. with temperature neglected. Section 2.1.2.2 details the effects of temperature on plasma oscillations. In a fully ionised plasma $\frac{\omega_{pe}}{\omega_{pi}} = \sqrt{\frac{Am_p}{m_e}} \sim 43A$, where A is the ion's atomic mass. The electron plasma frequency is much larger than the ion plasma frequency. When dealing with electron dynamics over short timescales the ions in the plasma can often be treated as stationary and their motion neglected.

Another key feature of a plasma is its ability to screen charge such that the electrostatic field of any particle is limited to a characteristic length — the Debye length, λ_D . Physically, this can be thought of as a cloud of negatively (or positively) charged particles surrounding an oppositely charged particle creating a sphere with radius λ_D beyond which the potential created by the charge in the centre has decayed significantly.

To derive the Debye length, use Poisson's equation to examine the potential in a small region of a quasi-neutral plasma:

$$\nabla^2 \Phi = \frac{-\rho}{\epsilon_0} = \frac{-e}{\epsilon_0} \left(\sum_i Z_i n_i - n_e \right), \quad (2.3)$$

where the sum is over all of the ion species in the plasma. Assuming the particles in the plasma obey a Maxwell-Boltzmann velocity distribution, the densities of electrons and ions in our region of plasma in response to the potential are defined as

$$\begin{aligned} n_e &= n_{e0} \exp\left(\frac{e\Phi}{k_B T_e}\right) \\ n_i &= n_{i0} \exp\left(\frac{-Ze\Phi}{k_B T_i}\right), \end{aligned} \quad (2.4)$$

where the zero subscript indicates the average number density across the region of plasma being considered. T_e and T_i are the electron and ion temperatures respectively. Often in a laser produced plasma they are not equal as the electrons will absorb most of the laser energy. The temperatures can equilibrate later in time as the electron energy is transferred

to the ions via collisions. Plasmas can often be treated as a fluid, using the Navier-Stokes equations to describe plasma dynamics. Typically, the electrons and ions will be treated as two separate fluids owing to differences in mass and temperature.

In the weak coupling limit, the kinetic energy of plasma particles ($\frac{3}{2}k_B T$) dominates their electrostatic potential energy ($q\Phi$), so Taylor expanding using $k_B T_{e,i} \gg q_{e,i}\Phi$ gives

$$n_{e,i} = n_{e0,i0} (1 \pm q_{e,i}\Phi/k_B T_{e,i}). \quad (2.5)$$

Using $\sum_i Z_i n_{i0} = n_{e0}$, Eq. (2.3) becomes

$$\begin{aligned} \nabla^2 \Phi &= \left(\frac{n_{e0} e^2}{\epsilon_0 k_B T_e} + \sum_i \frac{n_{i0} Z_i^2 e^2}{\epsilon_0 k_B T_i} \right) \Phi \\ &= \left(\frac{1}{\lambda_{De}^2} + \frac{1}{\lambda_{Di}^2} \right) \Phi \\ &= \frac{1}{\lambda_D^2} \Phi \end{aligned} \quad (2.6)$$

where the total Debye length is

$$\lambda_D = \left(\frac{\epsilon_0 k_B / e^2}{\frac{n_e}{T_e} + \sum_i \frac{Z_i^2 n_i}{T_i}} \right)^{1/2}. \quad (2.7)$$

Often, the electron Debye length is used, which is defined as

$$\lambda_{De} = \left(\frac{\epsilon_0 k_B T_e}{n_e e^2} \right)^{1/2}, \quad (2.8)$$

while the ion Debye length is:

$$\lambda_{Di} = \left(\frac{\epsilon_0 k_B / e^2}{\sum_i \frac{Z_i^2 n_i}{T_i}} \right)^{1/2}. \quad (2.9)$$

In a single ion plasma the difference between the electron and ion Debye lengths is just a factor of $\sqrt{\frac{Z_i T_e}{T_i}}$. This is demonstrated by Fig. 2.1, which shows that not only are the electron and ion Debye lengths of similar magnitude, but that the total Debye length is much closer to the ion Debye length. The reason to omit the ion term would be due to the differing




Figure 2.1: The electron, ion and total Debye length for a hydrogen plasma with $T_e = 2T_i = 3$ keV (standard ICF ablation plasma temperatures) as a function of normalised electron density, where n_c is the critical density — defined in Sec. 2.1.2.1.

time scales of ion and electron response in shielding a potential. Ions are thousands of times heavier than electrons, so respond over much slower timescales.

Electrons and ions in a plasma can often be characterised by a Maxwell-Boltzmann distribution with a characteristic temperature $T_{e,i}$ that describes the bulk of the distribution. The thermal velocity,

$$v_{e,i} = \left(\frac{k_B T_{e,i}}{m_{e,i}} \right)^{1/2}, \quad (2.10)$$

thus describes the characteristic velocity of the particles in the bulk of said distribution.

2.1.2 Waves

Unmagnetised plasmas can support three types of waves: transverse electromagnetic waves (EMW), and two longitudinal waves: electron plasma waves (EPW) and ion acoustic waves (IAW).

2.1.2.1 Electromagnetic waves

Electromagnetic waves in a plasma obey the dispersion relation

$$\omega_{emw}^2 = \omega_{pe}^2 + k_{emw}^2 c^2, \quad (2.11)$$

where ω_{emw} and k_{emw} are the frequency and wavenumber of the light wave and c is the

speed of light in a vacuum. From this equation, and remembering that plasma frequency increases with electron density, it is clear that above the density where $\omega_{emw} = \omega_{pe}$ the wavenumber becomes imaginary and the light wave can no longer propagate, instead becoming evanescent. This density is called the critical density and is found by setting the wavenumber in Eq.(2.11) to zero:

$$n_c = \frac{m_e \epsilon_0}{e^2} \omega_{emw}^2. \quad (2.12)$$

The rate at which a wave transports energy is called its group velocity, which is defined as

$$v_g = \frac{\partial \omega}{\partial k}. \quad (2.13)$$

Thus for an EMW,

$$v_{g,emw} = \frac{2k_{emw}c^2}{\omega_{emw}}. \quad (2.14)$$

The phase velocity of a wave is the rate at which any phase of a wave propagates, it is most easily thought of as the speed at which a wave's peak moves. A wave's phase velocity is defined as

$$v_\phi = \frac{\omega}{k}. \quad (2.15)$$

An electron in the presence of an EMW will oscillate in its electric field, perpendicular to the direction the wave is propagating. These oscillating electrons will in turn scatter light and drive the growth of parametric instabilities. This so called "quiver" velocity is

$$v_{osc}[m/s] = \frac{eE}{m_e \omega}, \quad (2.16)$$

or in more practical units for laser-plasmas,

$$v_{osc} = 0.256 \lambda_0 [\mu m] \sqrt{I_0 [W/cm^2]}, \quad (2.17)$$

where the 0 subscript signifies the wavelength and intensity for the pump EMW and the units of each are in the square brackets.

2.1.2.2 Electron plasma waves

Electron plasma waves are oscillations in electron density described by the Bohm-Gross relation (cf. Eq. (2.1) in a warm plasma):

$$\omega_{epw}^2 = \omega_{pe}^2 + 3k_{epw}^2 v_e^2. \quad (2.18)$$

This can be found by a fluid treatment of the plasma, but is also recovered from a kinetic approach when $k\lambda_D$ is small (a fluid approximation). As well as this real part to the frequency, the kinetic treatment yields an imaginary component of the frequency - the Landau damping rate. This is a collisionless wave damping that occurs between particles that have a similar thermal velocity to a wave's phase velocity; as these electrons remain in phase with the wave they can exchange energy with it. Electrons with $v_e < v_{\phi,epw}$ gain energy from the wave and are accelerated, whereas those with $v_e > v_{\phi,epw}$ lose energy to the wave. If the electron velocity distribution has a negative gradient with respect to velocity at $v_{\phi,epw}$ there will be more electrons with velocities below the wave's phase velocity than above it, so the net effect is that the wave is damped and hence electrons are accelerated. For a Maxwellian distribution of electrons EPWs will always be Landau damped, however with other electron distributions that have a positive gradient of the distribution function at a wave's phase velocity the wave instead will undergo Landau growth leading to unstable wave growth as is seen in the bump-on-tail instability [51]. The linear Landau damping rate for the amplitude of an EPW is

$$\nu_L = \sqrt{\frac{\pi}{8}} \frac{\omega_{pe}}{(k\lambda_D)^3} \exp\left(-\frac{1}{2(k\lambda_D)^2} - \frac{3}{2}\right). \quad (2.19)$$

Electron plasma waves and electromagnetic waves are both damped by electron-ion collisions. This process is explained in more detail in Sec. 2.1.3. The collisional damping rate for a wave with frequency ω is [52]

$$\nu_{coll} = \nu_{ei} \frac{\omega_{pe}^2}{\omega^2}. \quad (2.20)$$

The electron-ion collision frequency is

$$\nu_{ei} = \frac{1}{3(2\pi)^{3/2}} \frac{e^4}{\epsilon_0^2 m_e^{1/2}} \frac{Z_{eff} n_e}{(k_B T_e)^{3/2}} \ln \Lambda_{ei} \quad (2.21)$$

$$= 3.63 \times 10^{-6} \frac{Z^* n_e}{T_e^{3/2}} \ln \Lambda_{ei}, \quad (2.22)$$

where the Coulomb logarithm used is for electron-ion collisions defined on page 34 of the NRL plasma formulary [53]. $Z_{eff} = \Sigma_i n_i Z_i^2 / \Sigma_i n_i Z_i$ is the effective plasma ion charge [54].

2.1.2.3 Ion acoustic waves

Similarly to EPWs, ion acoustic waves are oscillations, but in ion density and, due to ion mass, of much lower frequency. The fluid dispersion relation for an ion acoustic wave in a cold, quasineutral plasma is

$$\omega_{iaw} = c_s k_{iaw}, \quad (2.23)$$

where $c_s = \sqrt{\frac{Z k_B T_e}{m_i}}$ is the ion sound speed. In a warm, quasineutral plasma, the ion temperature has an impact on the sound speed such that

$$c_s = \left(\frac{Z k_B T_e}{m_i} + 3v_i^2 \right)^{1/2}. \quad (2.24)$$

The factor 3 before the ion thermal velocity is its adiabatic index from treating the ions as an adiabatic fluid, similar to the electrons in an EPW as seen in Eq. (2.18). The electrons in an IAW are not adiabatic due to their larger thermal velocity allowing them to equilibrate during IAW oscillations. They are treated as isothermal with an adiabatic index of 1 as a coefficient on the first term in the sound speed equation.

To account for charge separation in the IAW, an additional factor is sometimes included on the electron term in the sound speed:

$$\omega_{iaw} = \left(\frac{Z k_B T_e}{m_i} \frac{1}{1 + k_{iaw}^2 \lambda_{De}^2} + 3v_i^2 \right)^{1/2} k_{iaw}. \quad (2.25)$$

Finally, in ICF plasmas there is a significant bulk plasma flow, \mathbf{u} , as the coronal plasma expands, which Doppler shifts the IAW by a frequency $\Delta\omega = \mathbf{u} \cdot \mathbf{k}$. Including all of these effects, the general dispersion relation for an IAW is:

$$\omega_{iaw} = \mathbf{u} \cdot \mathbf{k}_{iaw} + \left(\frac{Zk_B T_e}{m_i} \frac{1}{1 + k_{iaw}^2 \lambda_{De}^2} + 3v_i^2 \right)^{1/2} k_{iaw}. \quad (2.26)$$

The group velocity of this wave is

$$v_{g,iaw} = u_{\parallel} + C_1^{1/2} + \frac{k_{iaw}}{2C_1^{1/2}} \left(\frac{Zk_B T_e}{m_i} \frac{2k_{iaw} \lambda_{De}^2}{(1 + k_{iaw}^2 \lambda_{De}^2)^2} \right) \quad (2.27)$$

where

$$C_1 = \frac{Zk_B T_e}{m_i} \frac{1}{1 + k_{iaw}^2 \lambda_{De}^2} + 3v_i^2 \quad (2.28)$$

and u_{\parallel} is the component of plasma flow velocity parallel to the IAW wave vector. Ion acoustic waves are Landau damped by both electrons and ions, the damping rate is:

$$\nu_L = \sqrt{\frac{\pi}{8}} \frac{\omega_{iaw}}{(1 + k_{iaw}^2 \lambda_{De}^2)^{3/2}} \left(\left(\frac{Zm_e}{m_i} \right)^{1/2} + \left(\frac{ZT_e}{T_i} \right)^{3/2} \exp \left(-\frac{ZT_e}{2T_i} \frac{1}{1 + k_{iaw}^2 \lambda_{De}^2} - \frac{3}{2} \right) \right) \quad (2.29)$$

2.1.2.4 Ion Acoustic Waves in two ion plasmas

The coronas of directly driven ICF targets are typically a CH plasma. In a multi-ion plasma, multiple ion acoustic modes can be supported for each species in the plasma as each ion with a different mass has a unique oscillation frequency while also impacted by the presence of other ions. Williams *et al.* [14] and Vu *et al.* [55] both derived analytic expressions for IAW phase velocities and damping rates in two ion plasmas. Appendix A provides a detailed explanation of the attempts to use Williams' theory and the difficulties experienced in recreating the results in that paper.

For a two ion plasma, both authors describe the two ion modes as being fast and slow based on which has the larger phase velocity. Both of them focus on the situation where the two ion modes have phase velocities such that

$$v_{\phi,F} > v_{th,l} > v_{\phi,S} > v_{th,h}, \quad (2.30)$$

where the F and S subscripts refer to the fast and slow ion acoustic modes. The two thermal velocities are for each ion species where one species is termed h for heavy and one is subscripted l for the lighter ions. The kinetic theory first derived by Williams and that




Figure 2.2: (a) Phase velocity (v_ϕ) for the two distinct IAW modes compared to using single ion dispersion relation. In red is the 99%H 1%Xe (partially ionised) plasma. The dotted and dashed lines are the fast (F subscript) and slow (S subscript) mode respectively taken from Williams' paper. The solid line is the phase velocity using the single ion theory (ave subscript) in Eq. (2.26) with ion parameters found by averaging over the two species. Exactly the same applies to the blue lines except these were for the 50%C 50%H plasma. (b) Phase velocities for both the CH and XeH plasmas as a function of Hydrogen fraction for a constant $T_i = T_e/2$. Blue lines are for XeH plasma and red for CH, with the solid lines representing both fast modes and dashed for the slow modes. These results use Williams' kinetic theory. The three dotted lines are the phase velocity of single ion plasmas using Eq. (2.26) for H, C and Xe.

derived by Vu both require this inequality to remain true.

In cases with two ions that have much closer masses, like carbon and hydrogen, the ions and their modes no longer obey the condition in Eq. (2.30). In a CH plasma, the slow mode often has a phase velocity above the light ion's thermal velocity. Williams presents a new expression for the frequency and damping of ion modes based on kinetic theory without using the approximations they had previously used. This has not yet been solved numerically as part of this project so the original, approximated kinetic theory by Williams is currently used throughout this work. The work presented in the appendix shows that the phase velocity (from which the frequency and wavenumber can be found) predicted by this theory provide reasonable approximations to his numerical solutions.

It is not uncommon in literature to see single ion theory used with an average ion for a two ion plasma. Figure 2.2a shows that for the XeH and CH plasmas considered in Williams' paper, the phase velocity of the average ion sits in between the two modes Williams' theory actually predicts. In the case of the CH plasma the fast and slow mode are much closer together, because of the similar mass and charge of the ions. For the XeH plasma it seems

likely that the average ion phase velocity is much closer to the fast mode because the average ion is a lot closer to H than Xe given that the plasma has 99% H.

It was also anticipated that the two ion theory would converge with the single ion theory as the percentage of each ion tended towards 100%. Given the obvious parallels with Williams' original, approximated kinetic dispersion relation for a two ion plasma and the dispersion for a one ion plasma this seemed plausible. Figure 2.2b outlines this test for the same XeH and CH plasmas. As the percentage of hydrogen present approaches 0, all four ion modes approach the single ion phase velocity of C or Xe. At 100% H, both fast modes also reach the single ion phase velocity for hydrogen. The slow modes do not approach this same phase velocity at 100% H, instead decreasing rapidly away from it. It is unclear if this is correct behaviour or not as one might intuitively expect the two modes to converge upon each other as they do at the 0% hydrogen limit.

2.1.2.5 Wave Amplitudes

The growth of laser-plasma instabilities is dictated by the intensity of the laser's EMW, a higher laser intensity induces a larger electron quiver velocity, which creates larger instability growth rates. Intensity (energy per area per second) is a fairly common way to describe the amplitude of EMWs, but other quantities are also important in this work, especially for electrostatic waves where intensity is rarely used to describe them. The intensity of a wave is related to the amplitude of its electric field by

$$I = \frac{v_g \epsilon_0}{2} |E|^2. \quad (2.31)$$

Electrostatic waves are often described by their electric field amplitude or the magnitude of the density perturbation associated with them. For an electron plasma wave, the electrons are bunched in regions of high density with δn_e describing the difference in electron density between the extrema of the EPW and the background electron density. The relation between that density perturbation and the magnitude of the associated electric field is

$$E_{epw} = \frac{e}{\epsilon_0 k_{epw}} \delta n_e. \quad (2.32)$$

For an IAW

$$E_{iaw} = \frac{ek_{iaw}\lambda_{De}^2}{\varepsilon_0}\delta n_e = \frac{Z^*ek_{iaw}\lambda_{De}^2}{\varepsilon_0(1+k_{iaw}^2\lambda_{De}^2)}\delta n_i \quad (2.33)$$

Waves can also be described by their action density, A , in the literature this is often written as $|a|^2$, where a is a complex action amplitude (rather than normalised vector potential).

$$A = \frac{u}{\omega}, \quad (2.34)$$

where u is the energy density of the wave. A monochromatic wave has energy density $u = N\hbar\omega$, where N is the number density of quanta (quantised wave particles — photons, plasmons, phonons) and $\hbar\omega$ is the energy of a single quantum. From this, $A = \hbar N$ and action density can be seen to be a proxy for the density of quanta in a wave [56]. A wave's energy density can also be related to its intensity. For a wave with some energy per unit volume, its intensity is the rate at which this energy density passes through a surface, so $I = uv_g = v_g\omega A$. Subsequently, the action densities for our three waves of interest are:

$$A_{emw} = \frac{\varepsilon_0}{2\omega_{emw}} |E_{emw}|^2 = \frac{I_{emw}}{v_{g,emw}\omega_{emw}}, \quad (2.35)$$

$$A_{epw} = \frac{\varepsilon_0\omega_{epw}}{2\omega_{pe}^2} |E_{epw}|^2, \quad (2.36)$$

$$A_{iaw} = \frac{\varepsilon_0}{2\omega_{iaw}k_{iaw}^2\lambda_{De}^2} |E_{iaw}|^2 \quad (2.37)$$

2.1.2.6 WKB Approximation

For slowly varying plasma inhomogeneities, the Wentzel-Kramers-Brillouin (WKB) approximation allows us to simplify our wave equations [12]. This is used in the treatment of EMWs, EPWs and IAWs to find growth rates of the various instabilities. The WKB approximation is appropriate for a wave when the gradients in the plasma cause variation over larger lengths than its wavelength. The approximation fails in regions with steep gradients such that this no longer holds or near the turning point of a wave, where its wavelength increases dramatically. In this work, the WKB condition for a wave propagating parallel to x is defined as

$$\left| \frac{dk}{dx} \right| \ll k^2, \quad (2.38)$$

where k is its wavevector and x is its position in space.

2.1.3 Inverse Bremsstrahlung

When an electron undergoes a collision with an ion in the presence of a driving external field (from an EMW or EPW) the electron will gain energy and that external field will lose energy. This process is called collisional absorption or inverse bremsstrahlung, named after it being the opposite process to bremsstrahlung emission. The final power of a laser after propagating between two points, x_i to x_f , in a plasma is

$$P_f = P_i \exp \left(\int_{x_i}^{x_f} \kappa_{coll} dx \right), \quad (2.39)$$

where P_i is the initial power at position x_i . κ_{coll} is the inverse bremsstrahlung spatial damping rate and is defined, in SI units, as [57]

$$\kappa_{coll} = \frac{64\pi^3 Z_{eff} n_e^2 \ln \Lambda}{3c(2\pi m_e k_B T_e)^{3/2}} \left(\frac{e^2}{4\pi\epsilon_0} \right)^3 \frac{1}{\omega^2} \left(1 - \left(\frac{\omega_{pe}}{\omega} \right)^2 \right)^{-1/2} \quad (2.40)$$

Substituting out most of the constants leaves

$$\kappa_{coll} = 1.98 \times 10^{-45} \frac{Z^* n_e^2 \ln \Lambda}{(k_B T_e)^{3/2}} \frac{1}{\omega^2} \left(1 - \left(\frac{\omega_{pe}}{\omega} \right)^2 \right)^{-1/2}. \quad (2.41)$$

Here ω is the frequency of the wave being absorbed and $\ln \Lambda$ is the Coulomb logarithm. While defined slightly differently in many different places, here we use the NRL plasma formulary definition for the Coulomb logarithm from page 58 as it is specifically for inverse bremsstrahlung [53].

Inverse bremsstrahlung is the desired absorption mechanism to drive an ICF implosion because its absorption rate is proportional to $T_e^{-3/2}$. Thus it preferentially heats the lower energy electrons in the thermal bulk of the distribution. Being of lower energy these electrons are more likely to heat the capsule's ablator via collisions with ions, aiding its expansion and the subsequent compression of the core. If the laser energy was absorbed by electrons in the high energy tail of the distribution, they would be more likely to gain enough energy to propagate into the ice layer of the capsule, depositing energy there, causing fuel preheat and reducing compression efficiency.

The absorption rate is also proportional to n_e^2 — a photon is more likely to be absorbed by




Figure 2.3: Evolution of the Langdon correction factor, f , with (a) α and (b) laser intensity at differing electron temperatures.

the process in regions with a larger density of electrons to absorb it. In an ICF implosion, as the laser propagates up the plasma density profile it is often significantly absorbed by other mechanisms in the underdense plasma before it can reach the regions near the laser's critical density where the collisional absorption rate becomes much larger. These other absorption mechanisms that risk removing laser energy from driving the implosion are laser-plasma instabilities.

The collisional absorption rate is found using the collision frequency of electrons and ions in a Maxwell-Boltzmann distribution. The Langdon effect describes the fact that high laser intensities can move the electron distribution function away from a Maxwellian and change the collision frequency [58]. It has been shown that the inverse bremsstrahlung absorption coefficient is reduced by a factor called the Langdon correction:

$$f = 1 - \frac{0.533}{1 + (0.27/\alpha)^{3/4}}, \quad (2.42)$$

where $\alpha = Z_e f v_{osc}^2 / v_e^2$ using the electron thermal and quiver velocities from eqs. (2.10) and (2.17).

Figure 2.3a shows the relationship between the Langdon correction and alpha. When $\alpha \rightarrow 0$ the electrons mostly follow a Maxwellian distribution so $f \rightarrow 1$. Larger values of α lead to the Langdon correction asymptoting at 0.467. Using laser intensity instead of α and three different, constant electron temperatures, figure 2.3b shows that the Langdon effect makes a significant difference in conditions relevant to this model. The upper bound of Chiron,




Figure 2.4: Schematic for the mechanism behind stimulated Raman backscattering that can be applied to any three wave parametric instability by just exchanging waves and directions. Adapted from Montgomery [8].

which is based on fluid theory, is estimated to be around an intensity of $3 \times 10^{15} \text{ W/cm}^2$. At this intensity, the Langdon effect reduces IB absorption by 17% for a 5 keV plasma. Recent work by Turnbull *et al.* [59, 60] showed that the Langdon correction along with a further ion screening correction were essential for their IB models to correctly compare to their experimental results. Chiron does not currently include a model for the ion screening correction, but its implementation could be considered in future developments.

2.2 Laser-Plasma Instabilities

Figure 2.4 outlines the feedback process that leads to the growth of light scattering laser-plasma instabilities, such as SRS, SBS and CBET. When a sufficiently large amplitude laser propagates through a plasma, electrons will oscillate in its field and in turn Thomson scatter light. This scattered light wave beats with the laser wave creating a beat-wave envelope whose ponderomotive force bunches electrons (or ions) into its minima creating an electrostatic

wave. If this wave matches a supported electrostatic mode of the plasma then the three waves are resonant and will continue beating together and growing unstably. The figure depicts the process for stimulated Raman backscattering, with the oppositely propagating light scattering off an electron background. The process is the same for stimulated Brillouin scattering and crossed-beam energy transfer but with ions in place of the electrons in the picture. For two-plasmon decay, both daughter waves are EPWs that are not collinear with the laser, however the general feedback loop that drives the instability is similar. Rather than scattering light off the electrostatic mode, the beating of the laser light with each EPW drives the unstable growth of the other EPW.

In addition to the need for waves to obey the linear dispersion relations, the waves are constrained to be in resonance with each other for the beating process to occur and drive growth. The three waves must have their frequencies matched so that

$$\omega_0 = \omega_1 + \omega_2 \quad (2.43)$$

and their wavevectors must also be matched,

$$\mathbf{k}_0 = \mathbf{k}_1 + \mathbf{k}_2. \quad (2.44)$$

The subscript 0 refers to the parent wave, typically the laser's EMW, that decays into the two daughter waves labelled 1 and 2. These two equations are comparable to conserving energy and momentum respectively between the three waves and must be satisfied for LPI to grow. There has been recent interest in using broadband lasers for ICF to mitigate LPI growth [61, 62]. For a monochromatic laser, all of the laser's energy can go into driving the growth of an instability for a single pair of daughter waves. However if the laser's energy can be distributed across a wider spectrum of frequencies, each frequency bandwidth of laser energy is resonant with a different set of daughter waves, thus reducing the energy available to each instance of unstable growth and increasing the chance of the pump being below the threshold intensity for the instability to grow.

An additional conservation law that must be maintained between the three waves is the conservation of action density, known as the Manley-Rowe relations [63],

$$\frac{dA_0}{dt} = -\frac{dA_1}{dt} = -\frac{dA_2}{dt}. \quad (2.45)$$

This describes the fact that every quantum lost from the pump wave results in one new quantum of each of the daughter waves.

In a homogeneous plasma the coupled equations between the three waves are derived by Michel [12] using the aid of the slowly varying envelope approximation

$$\begin{aligned} \left(\frac{\partial}{\partial t} + \nu_0 + v_{g0} \frac{\partial}{\partial x} \right) \tilde{a}_0 &= iK_0 \tilde{a}_1 \tilde{a}_2, \\ \left(\frac{\partial}{\partial t} + \nu_1 + v_{g1} \frac{\partial}{\partial x} \right) \tilde{a}_1 &= iK_1 \tilde{a}_0 \tilde{a}_2^*, \\ \left(\frac{\partial}{\partial t} + \nu_2 + v_{g2} \frac{\partial}{\partial x} \right) \tilde{a}_2 &= iK_2 \tilde{a}_0 \tilde{a}_1^*, \end{aligned} \quad (2.46)$$

where K_j are real coupling coefficients and the tilde indicate that these quantities are the slowly varying envelope of the wave vector potentials. The amplitude of the j th wave is

$$a_j = \frac{1}{2} \tilde{a}_j \exp(i\mathbf{k}_j \cdot \mathbf{r}_j - \omega_j t), \quad (2.47)$$

and the slowly varying approximation means that

$$\begin{aligned} \left| \frac{\partial^2}{\partial x^2} \tilde{a}_j \right| &\ll |k_j^2 \tilde{a}_j| \\ \left| \frac{\partial^2}{\partial t^2} \tilde{a}_j \right| &\ll |\omega_j^2 \tilde{a}_j|. \end{aligned} \quad (2.48)$$

These equations mean that the wave envelope evolves slowly in space and time compared to the wave's period and wavelength. Throughout this section pump depletion is ignored to simplify the problem, thus the first of the three coupled equations can be ignored as \tilde{a}_0 is constant. The changing amplitude of the pump is an analytically complex, nonlinear aspect to LPI. This problem is addressed computationally rather than analytically and is discussed further in Sec. 3.1.3.

2.2.1 Convective and Absolute instabilities

There are two distinct modes of amplification for laser-plasma instabilities which have important consequences for how they grow. When an instability is described as growing convectively the daughter waves amplify in a moving frame [12]. The coupling between the waves becomes weaker as they move further from the location of their resonance, reducing the rate of amplification in space. This decoupling process is explained in detail Sec. 2.2.3.

An instability grows absolutely when the daughter waves amplify exponentially in time at a fixed location in space [64]. The following section contains thresholds for both convective and absolute growth of an instability. This typically happens when the group velocities of the waves are very small so they are effectively trapped. These waves are not moving away from their resonance at a significant rate so do not undergo saturation by decoupling. Instead they grow exponentially in time until saturated by other mechanisms, such as pump depletion and nonlinear, secondary wave decays. This rapid amplification has made absolute instabilities a particular concern for laser fusion.

2.2.2 Temporal Amplification

The temporal growth rate of an instability in an infinite, homogeneous plasma is derived by setting the spatial derivatives to zero in the three wave equations. Ignoring pump depletion such that \tilde{a}_0 is constant and ignoring damping on both waves, the coupled equations become

$$\frac{\partial}{\partial t} \tilde{a}_1 = iK_1 \tilde{a}_0 \tilde{a}_2^*, \quad (2.49)$$

$$\frac{\partial}{\partial t} \tilde{a}_2 = iK_2 \tilde{a}_0 \tilde{a}_1^*. \quad (2.50)$$

Differentiating Eq. (2.49) and substituting the complex conjugate of Eq. (2.50) into it gives

$$\frac{\partial^2}{\partial t^2} \tilde{a}_1 = K_1 K_2 |\tilde{a}_0|^2 \tilde{a}_1. \quad (2.51)$$

The solution for this is

$$\tilde{a}_1(t) = \tilde{a}_1(t=0) \exp(\gamma_0 t), \quad (2.52)$$

where

$$\gamma_0^2 = K_1 K_2 |\tilde{a}_0|^2 \quad (2.53)$$

is the temporal growth rate, which is clearly dependent on the amplitude of the pump wave. Thus, the growth rate of the instability is dependent on the intensity of the pump wave that it is simultaneously depleting, as is evident from Eq. (2.46). This is the previously mentioned nonlinear relation between pump depletion and LPI growth rate that makes steady-state modelling difficult.

Equation (2.53) is an expression for the maximum possible growth rate of a parametric

instability in the absence of damping. Nishikawa [65] showed the effect that damping has on the growth rate of the two daughter waves. He arrives at

$$\gamma = \sqrt{\gamma_0^2 + \frac{1}{4}(\nu_2 - \nu_1)^2} - \frac{1}{2}(\nu_1 + \nu_2) \quad (2.54)$$

Rearranging this provides a threshold for the growth of the instabilities. For an instability to have $\gamma > 0$,

$$\gamma_0 > \sqrt{\nu_1 \nu_2}, \quad (2.55)$$

otherwise the damping of the daughter waves is too large for the instability to overcome. Although Eq. (2.54) provides a more complete description of instability growth rate, it is important to recognise the maximum growth rate in Eq. (2.53) as this appears in a number of later equations.

The threshold for an instability to grow absolutely was given by Briggs [66]:

$$\gamma_0 = \sqrt{v_{g1} v_{g2}} \left(\frac{\nu_1}{v_{g1}} + \frac{\nu_2}{v_{g2}} \right). \quad (2.56)$$

The terms in the brackets can be thought of as a spatial damping rate — the rate the wave is being damped per unit distance it travels. Even if wave velocities are small, high damping rates can prevent instabilities from growing absolutely. In the scenario where the growth rate exceeds the threshold in Eq. (2.55), but not the threshold in Eq. (2.56), an instability will grow convectively.

2.2.3 Convective Amplification

When considering parametric instabilities in an inhomogeneous plasma, the coupling of waves is reduced by an increasing mismatch in their wavevectors, limiting the growth. The wave vector mismatch, is defined as

$$\kappa(x) = k_0(x) - k_1(x) - k_2(x). \quad (2.57)$$

Explicitly for backscattering instabilities, it can be written as:

$$\kappa(x) = |k_0(x)| + |k_1(x)| - |k_2(x)|. \quad (2.58)$$




Figure 2.5: Plots demonstrating the change in SRS wavevectors through a density gradient. (a) and (b) show the scattered EMW and EPW respectively, while (c) shows the wavevector mismatch described in Eq. (2.58). The black line on all three plots shows the values of that parameter that satisfy the wavenumber matching conditions. The coloured lines show the evolution of individual waves as an instability moves in space. A wave starts where the coloured line meets the black line as the wavenumber matching conditions are met at the point an instability is seeded, but then the wave’s wavevector moves further away from wavenumber matching in space.

Using the magnitudes of the wavevectors provides clarity compared to some of the literature which often uses k_1 , which they have defined as $-|k_1(x)|$. At the position an instability is resonant, $\kappa(x = 0) = 0$, as required by Eq. (2.44). However, as the waves propagate away from this point, assuming their angular frequency remains constant, gradients in plasma parameters drive the wavevectors to change relative to each other. Thus, κ increases with distance from the seed location and coupled growth slows down.

This is illustrated by examining SRS, which has an EMW and EPW as the daughter waves. It is clear from the two dispersion relations in eqs. (2.11) and (2.18), that as the waves propagate through new densities, the plasma frequency shifts, and so the wavevectors must also change to compensate as the wave frequency remains constant. For the EPW, the wavevector is also affected by gradients in temperature which will change the electron thermal velocity. Figure 2.5 illustrates how the SRS wavevectors diverge from the values that correspond to wavenumber matching (the black lines).

For TPD, which is an inherently 2D instability the mismatch is

$$\kappa(x) = |k_0(x)| - |k_{1x}(x)| + |k_{2x}(x)|, \quad (2.59)$$

only the component of the daughter wavevectors parallel to the pump (along x) is considered [67]. Here the signs in the equation have also been switched because in TPD convention,

wave 1 is the forward wave and 2 is travelling backwards, this is depicted in Fig. 2.12 and explained in depth in Sec. 2.2.6.

2.2.3.1 Convective Gain

Considering now the case of a pump wave decaying into two daughter waves, which amplify convectively in an inhomogeneous plasma. Daughter wave number 2 is the electrostatic mode that is being amplified by the beating of the pump wave and the daughter wave numbered 1, which can be either electromagnetic or electrostatic here. Ignoring pump depletion (treating \tilde{a}_0 as a constant and ignoring its coupled equation) and assuming steady-state ($\frac{\partial}{\partial t} \rightarrow 0$), the slowly varying coupled equations now become:

$$\left(v_{g1} \frac{\partial}{\partial x} + \nu_1\right) \tilde{a}_1 = iK_1 \tilde{a}_0 \tilde{a}_2^* \exp\left(i \int \kappa(x) dx\right), \quad (2.60)$$

$$\left(v_{g2} \frac{\partial}{\partial x} + \nu_2\right) \tilde{a}_2 = iK_2 \tilde{a}_0 \tilde{a}_1^* \exp\left(i \int \kappa(x) dx\right), \quad (2.61)$$

with the loss of coupling due to inhomogeneity encompassed within the exponential term. Using Michel's derivation [12] as a rough guide, a new variable is introduced,

$$\alpha_2 = \tilde{a}_2 e^{-i \int \kappa dx}, \quad (2.62)$$

and then substituted into the LHS of Eq. (2.61):

$$\left(\nu_2 + v_{g2} \frac{\partial}{\partial x}\right) \alpha_2 e^{i \int \kappa dx} = \left(\nu_2 + i\kappa v_{g2} \alpha_2 + v_{g2} \frac{\partial \alpha_2}{\partial x}\right) e^{i \int \kappa dx}. \quad (2.63)$$

Assuming the electrostatic wave 2 is strongly damped ($\nu_2 \alpha_2 \gg v_{g2} \partial_x \alpha_2$) the coupled equations are now

$$\left(v_{g1} \frac{\partial}{\partial x} + \nu_1\right) \tilde{a}_1 = iK_1 \tilde{a}_0 \alpha_2^*, \quad (2.64)$$

$$(\nu_2 + i\kappa v_{g2}) \alpha_2 = iK_2 \tilde{a}_0 \tilde{a}_1^*. \quad (2.65)$$

Taking the complex conjugate of Eq. (2.65) and substituting into Eq. (2.64) for α_2^* gives

$$\left(v_{g1} \frac{\partial}{\partial x} + \nu_1\right) \tilde{a}_1 = \frac{\gamma_0^2}{\nu_2 - i v_{g2} \kappa} \tilde{a}_1, \quad (2.66)$$

where γ_0^2 is as defined in Eq. (2.53). Finally, multiplying by \tilde{a}_1^* and adding the complex conjugate gives an expression in terms of \tilde{a}_1^2 , which is proportional to the wave's intensity:

$$\frac{\partial}{\partial x} |\tilde{a}_1|^2 = \frac{2\gamma_0^2}{v_{g1}v_2} \left(\frac{1}{1 + \left(\frac{v_{g2}\kappa}{v_2}\right)^2} - \frac{v_1v_2}{\gamma_0^2} \right) |\tilde{a}_1|^2. \quad (2.67)$$

First, the second term in the brackets, which describes the decrease in wave 1's amplitude as a result of damping dissipating energy, is ignored. The first term includes the information about the parametric growth of the wave amplitude. In his seminal paper on parametric instabilities [5], Rosenbluth worked with a plasma that had linear gradients in κ , so it could be approximated as $\kappa = \kappa'x$, where the prime refers to the spatial derivative. Applying this approximation to Eq. (2.67) and integrating gives

$$|\tilde{a}_1(x)|^2 = |\tilde{a}_1(x_i)|^2 \exp \left[\frac{2\gamma_0^2}{v_{g1}v_{g2}\kappa'} (\arctan(x/L_d) - \arctan(x_i/L_d)) \right], \quad (2.68)$$

where $L_d = \frac{v_2}{\kappa'v_{g2}}$ is the ‘‘amplification length’’ (a scale length for the rate of spatial amplification). When integrating over all space, i.e. between $\pm\infty$, using the limits of the arctangent function, the wave amplitude is

$$|\tilde{a}_1(x_f)|^2 = |\tilde{a}_1(x_i)|^2 \exp \left(\frac{2\pi\gamma_0^2}{v_{g1}v_{g2}\kappa'} \right). \quad (2.69)$$

The well known Rosenbluth gain [5] has been recovered. This is notable as included damping was included in the derivation, yet the same result was arrived at despite Rosenbluth ignoring damping on the waves. This is because of the fact that this result is derived by integrating between $\pm\infty$ — the first daughter wave will still amplify to the same saturated level regardless of the amount of damping on the second wave, but it will require amplification over more space to reach this value as damping increases. The amplitude of the plasma wave is significantly affected by damping, however. It will grow within the amplification region, but beyond this its energy is quickly damped away into the heating the thermal plasma particles (collisional damping) or further heating the suprathermal tail of the electron population (Landau damping).

Rather than integrating between $\pm\infty$, Eq. (2.67) can be integrated along only the path of the scattered ray, starting at zero where it is seeded and finishing at δx , the length of the

scattered ray once it has left the plasma. This gives the new result

$$|\tilde{a}_1(x_f)|^2 = |\tilde{a}_1(x_i)|^2 \exp\left(\frac{2\gamma_0^2}{v_{g1}v_{g2}\kappa'} \arctan\left(\frac{v_{g2}\kappa'}{\nu_2} \delta x\right)\right). \quad (2.70)$$

In this work, Eq. (2.70) is referred to as the “bounded” gain due to the finite bounds on the analytic integration in contrast to the Rosenbluth gain. Notably, this now has some dependence on the damping of the second daughter wave because the integrating is no longer bounded at infinity. The fact that damping decreases the rate the wave amplifies in space is important because the wave has limited space over which it grows.

The linear approximation made regarding the wavevector mismatch for the analytic integration is potentially quite weak in ICF plasmas where gradients are often exponentially steepening. Instead, Eq. (2.67) can be integrated keeping κ in place of $\kappa'x$, while still ignoring the second term (damping of wave 1) on the right:

$$|\tilde{a}_1(x)|^2 = |\tilde{a}_1(x_i)|^2 \exp\left(\int \frac{2\Gamma}{1 + \left(\frac{v_{g2}\kappa}{\nu_2}\right)^2} dx\right), \quad (2.71)$$

where $\Gamma = \frac{\gamma_0^2}{v_{g1}\nu_2}$ is the spatial growth rate. This integral is not solvable analytically but can be done numerically, allowing the variation of all of the plasma parameters to be tracked along the path of the amplifying wave.

2.2.3.2 Including Damping

Returning again to Eq. (2.67), the second term on the right hand side that relates to wave damping was previously neglected. In reality, this is often not the case and the wave damping will play a role in instability evolution. Keeping the second term in Eq. (2.67) and integrating yields:

$$|\tilde{a}_1(x)|^2 = |\tilde{a}_1(x_i)|^2 \exp\left(\int 2\Gamma \left[\frac{1}{1 + \left(\frac{v_{g2}\kappa}{\nu_2}\right)^2} - \frac{\nu_1\nu_2}{\gamma_0^2} \right] dx\right), \quad (2.72)$$

which can again be integrated numerically along the path of a scattered ray.

Figure 2.6a, shows the convective amplification of three Raman scattered rays seeded at 0.15, 0.2 and 0.245 n_c as they propagate to the left out of the plasma. Each ray is amplified following 4 of the above equations. Those equations are ordered in the legend as

the numerical integration of eqs. (2.71), (2.72) and the analytically integrated Rosenbluth gain and bounded gain from eqs. (2.69), (2.70).

This simulation was carried out in a linear density gradient with the hope that this would keep the wavevector mismatch as close to linear as possible as this approximation is part of the Rosenbluth gain. Figure 2.6b shows that the wavevector mismatch for the three rays is close to linear along the path of the rays.

For all three rays, the Rosenbluth gain (dashed) predicts a significantly higher saturated level than the newly derived bounded gain (dashdotted). This behaviour is expected given the only difference is that the Rosenbluth gain is integrated over a longer distance. The model assumes both of these analytically integrated gains saturate over the very short distance of one amplification length. In reality, the scattered ray will not saturate so quickly in space, but this is a clear and easy way for to model and depict these two approximations. Importantly, the bounded gain is a lot closer to the solid line “integrated” gain, which is assumed to be the most accurate model for convective amplification, as it contains fewer approximations and considers the evolution of plasma parameters. The integrated method sees maximum spatial amplification in closest proximity to the seed, where the waves are all resonant. Further away from this, the rate of amplification in space decreases. A more detailed comparison of the results from different amplification methods can be found in Sec. 3.1.5.

The dotted line for each scattered ray shows amplification following the numerical integration of Eq. (2.72). This expression has two terms, the first being identical to that in Eq. (2.71), the solid line in the plot. Multiplying out the coefficient of the second term shows that it has the form $-\frac{\nu_1}{v_{g1}}$, thus this term is a decrease in the amplitude of the wave due to it being damped. In this example for SRS the scattered light wave is being collisionally damped as it propagates. The fact it is undergoing collisional absorption means that not all of the laser energy that is absorbed by SRS is lost, some of it still is absorbed collisionally, which will enhance ablation. An important distinction to make regarding the two numerically integrated descriptions of amplification (eqs. (2.71) and (2.72)) is that although they result in a different amplitude of scattered wave, they both result in the same loss of energy from the pump. Including the collisional damping term also accounts for further energy losses from the scattered wave. Williams [68] has provided an analytically integrated solution for the amplitude of the first daughter wave when it is undergoing both convective amplification and being damped. His solution is an additional factor that can be applied straight to the

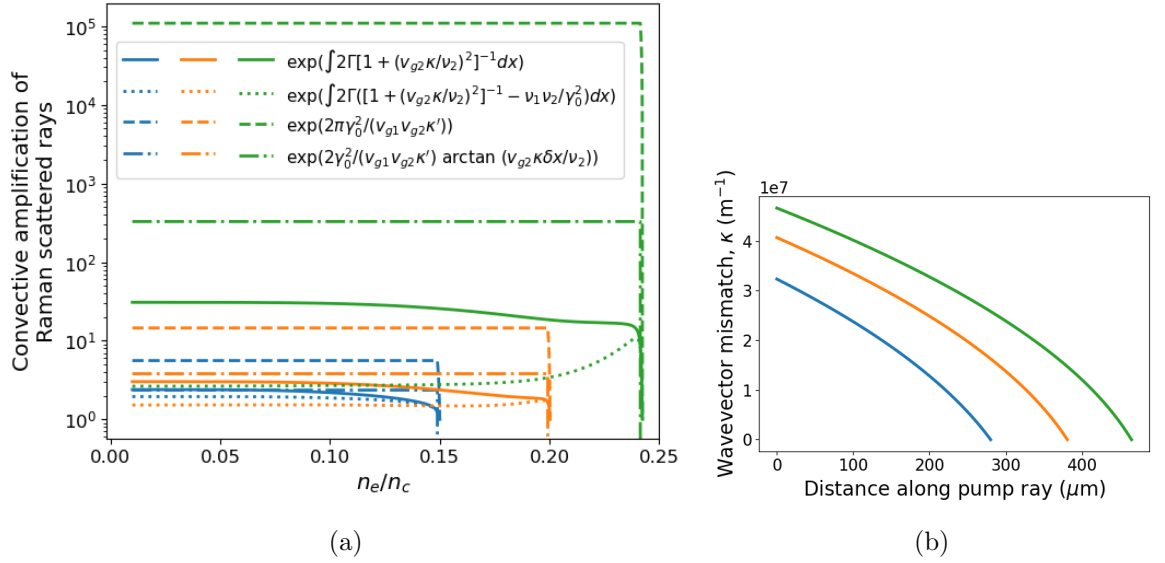


Figure 2.6: (a) The amplification of three Raman scattered rays (colour coded) seeded at different densities. The amplification of each ray is modelled based on Eq. (2.71) (solid line), Eq. (2.72) (dotted line), Eq. (2.69) (dashed line) and Eq.(2.70) (dashdotted line). The pump ray had intensity 10^{15} W/cm^2 and the plasma had a linear density gradient with a scale length at $n_c/4$ of $500 \mu\text{m}$. Pump depletion was ignored for this simulation. (b) Profiles of the wavevector mismatch, κ , along each back propagating scattered wave.

Rosenbluth gain.

Based on the results with these two numerically integrated expressions for amplification, I propose the collisional absorption of the scattered wave as a possible mechanism for the widely seen but little understood Raman gap, discussed in more detail and with clearer examples in Sec. 2.2.4.2.

2.2.4 Stimulated Raman Scattering

Stimulated Raman scattering describes the process where an EMW decays into a secondary, scattered EMW and an EPW. SRS is typically stated as not being able to occur at densities above $n_c/4$, but it is actually limited slightly below $n_c/4$. Looking at its dispersion relation, Eq. (2.11), the scattered EMW must have $\omega_s > \omega_{pe}$ to be able to propagate. The frequency matching condition, Eq. (2.43), becomes $\omega_0 > \omega_{pe} + \omega_{epw}$ and, by substituting in n_e and n_c :

$$n_e < \frac{n_c}{\left(1 + \sqrt{1 + 3k_{epw}^2 \lambda_{De}^2}\right)^2}. \quad (2.73)$$

In the limit of $T_e \rightarrow 0$ this reduces to $n_e < n_c/4$.

Experimentally, SRS is detected by the scattered light it generates. Figure 2.7 shows the scattered light spectrum as a function of time from an experiment conducted on the NIF to explore ignition relevant conditions for direct drive ICF [9]. The wavelengths of the scattered light correspond to where in the plasma a backscattered ray was first emitted. The broadband scattered light emitted between 600 and 650 nm is from convective SRS and the bright feature above 700 nm is from absolute SRS, which is discussed in Sec. 2.2.4.1. between 650 and 700 nm the lack of scattered light is known as the Raman gap, which is often present in SRS data and is not well understood. Section 2.2.4.2 discusses past and current theories including work done in this project to explain it.

This density and temperature variation of scattered light wavelength is shown in figure 2.8a where each coloured line represents the spectrum of Raman scattered light waves and the density at which they were seeded. The relationship between emitted wavelength and density is roughly linear in these calculations which used a constant temperature across the plasma, however the form of this relationship varies when the temperature profile is not constant. This plot also confirms what is described in Eq.(2.73): that increasing plasma temperature decreases the maximum density SRS can occur at. The higher temperature spectra have a lower maximum density than the lower temperature spectra.

For each position in density-temperature space there exist a set of SRS waves that are solutions to the dispersion relations and wave-matching conditions, each at different orien-




Figure 2.7: Spectrum of stimulated Raman scattered light from a NIF experiment. Taken from [9].

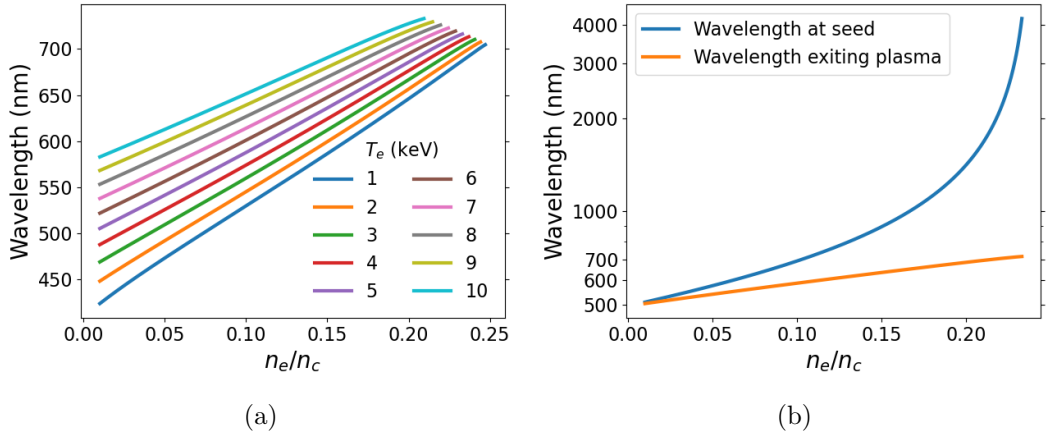


Figure 2.8: (a) The measured wavelength of a scattered ray based on the density at which it is seeded for plasmas with 10 different, but constant profiles of electron temperature. (b) The wavelength of scattered rays that are resonant with the Raman instability at each density for a plasma with constant $T_e = 5$ keV. In blue is the wavelength of the ray at the point of resonance in the plasma that satisfies the wavematching conditions. In orange is the wavelength each of the resonant rays has become by the time they have propagated out of the plasma, and thus is the wavelength at which they would be measured.

tations. This work focuses solely on SRS backscatter, i.e. when the pump and EPW are parallel and the scattered EMW is antiparallel. Backscatter has long been considered the dominant SRS mechanism. Each wave has constant frequency and a wavelength that varies as it propagates through varying plasma densities and temperature. Figure 2.8b shows the initial and final wavelength that each backscattered wave would have when seeded across the density profile.

Other SRS directions are documented as well, such as forward scatter, where the scattered EMW and EPW are parallel to the pump wave. This is generally considered a much smaller concern than backscatter for ICF as the orientation dictates that $k_{epw}^{forward} < k_{epw}^{back}$, and so with $\gamma_0 \propto k_{epw}$, backscatter will always grow faster as shown in simulation work by Estabrook *et al.* [69]. Further simulations also found that forward scatter is much more sensitive to density gradients, with those from ICF implosions expected to suppress it significantly more than backscatter [70].

Stimulated Raman sidescatter occurs when scattered light is resonant approximately perpendicular to the density gradient. Initially of great interest in the 1970s, interest has recently picked up after experiments on the NIF followed by new theory work that indicated that sidescatter could be as significant, if not more so, than backscatter [71, 72]. This has been complemented by experiments trying to isolate its spectrum from backscatter and

thus verify theory and understand how important it may be to ICF [73, 74]. Sidescattered rays are defined as those approximately perpendicular to the density gradient. These rays are launched approximately parallel to contours of constant density so remain in resonance with the instability for longer than a backscattered ray. Eventually, the scattered rays will refract and leave the regions of high resonance on their way out of the plasma, however it has been shown that the amplification length for sidescatter rays is typically longer than for backscatter rays. This, in combination with their transverse propagation across the laser beam focal spot, can lead to larger convective gains than for backscattered rays [72]. Raman sidescatter can also become absolute at any density as it occurs at the turning point of the scattered ray, whereas backscattered rays can only grow absolutely at $n_c/4$. These factors make sidescatter very plausibly the dominant SRS mechanism in certain scenarios. Experimental geometries with laser beams at higher incidence angles relative to the plasma density gradient are more prone to sidescatter for the reasons just listed. Raman sidescatter has not yet been included in this project, but the developed framework for modelling LPI will allow implementations of this and other models in the future.

Following the process outlined in sec 2.2.2, and using wave equations for driven EMWs and an EPW, Kruer shows the homogeneous temporal growth for SRS to be [52]

$$\gamma_0 = \frac{v_{osc} k_{epw} \omega_{pe}}{4\sqrt{\omega_s \omega_{epw}}}. \quad (2.74)$$

2.2.4.1 Raman growth at $n_c/4$

For backscattered SRS near $n_c/4$ (or for sidescatter anywhere) the daughter waves have small wavevector components along the density gradient and so the WKB approximation, which was used in the derivation of the above growth rate breaks down. Liu *et al.* [10] derived a growth rate for SRS in a linear density gradient, which does not rely on the WKB approximation:

$$\gamma_0 = \frac{v_{osc} \omega_{pe} k_{epw}}{2\omega_0} \left[1 - \frac{(n + \frac{1}{2}) \sqrt{\omega_{pe}/(ck_{epw})}}{\sqrt{2} (\frac{v_{osc}}{c})^{\frac{3}{2}} k_{epw} L_n} \right], \quad (2.75)$$

where n is a positive integer eigenvalue solutions and is set to $n = 0$ for the eigenmode with the largest growth rate. The first term of this rate is almost identical to the WKB growth rate (which is seen in Fig. 2.9), while the second term, which depends on the gradient, creates a difference between them. Figure 2.9 shows the two SRS growth rates derived with and

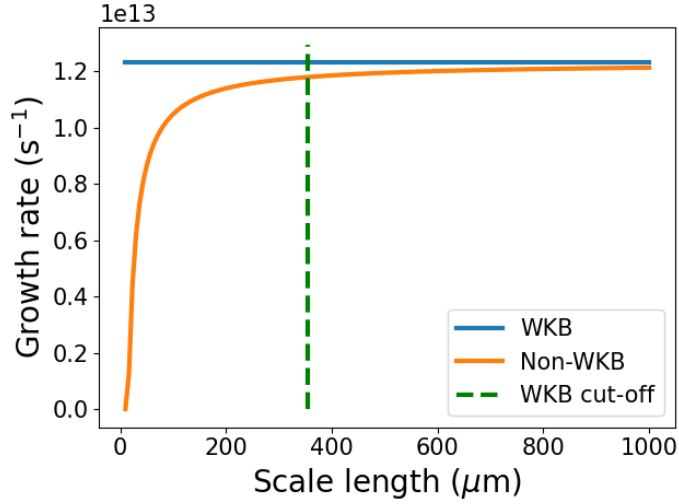


Figure 2.9: Comparison of the WKB and non-WKB growth rates given in eqs. (2.74) and (2.75). The WKB approximation depends on gradients in the wave parameters, so is plotted as a function of plasma density scale length at a fixed density of $0.235n_c$. The WKB cut-off is defined by the expression in Eq. (2.38).

without the WKB approximation at a single point in density space ($0.235n_c$) with varying density scale lengths. The simulation was for a 2.5 keV plasma with a linear density profile and laser intensity $10^{15}\text{W}/\text{cm}^2$. The linear density gradient was chosen in the hope that it would keep the non-WKB growth rate as accurate as possible given the constraints during its derivation, however when using exponential plasma gradients the plot was almost identical. The green dashed line marks the position in the plasma where the WKB approximation fails according to the condition in Eq. (2.38); at shorter scale lengths the density gradient is too steep and WKB theory is not applicable.

The WKB growth rate is constant in this plot because it has no dependence on density gradient. The Non-WKB growth rate diverges from the WKB growth rate as the density gradient steepens (scale length gets shorter). The WKB threshold condition functions exactly as anticipated — it occurs where the two growth rates diverge. At different points in density space the threshold also occurs around where the growth rates diverge, but for different densities this point occurs at different scale lengths: at lower densities the threshold is at lower scale lengths.

The exact position of the cut-off is somewhat arbitrary, as it is based on Eq.(2.38) having a \ll sign. In this work, the much less than sign has been equated to a factor of 100 between $|dk/dx|$ and k^2 , however that that could be changed to redefine where the

WKB approximation is valid.

Below the quarter critical surface, backscattered SRS is a convective instability that Chiron saturates using the previously described bounded gain. Near $n_c/4$ convective saturation is limited by the small group velocities of the waves. Here, the instability is primarily saturated by the parametric decay of its EPW into an IAW and a secondary, non-resonant EPW. This process is known as the Langmuir decay instability (LDI) and is described in more detail in Sec. 2.2.7.

Absolute SRS has been the subject of theoretical study, starting with Drake and Lee in 1973 finding a solution for its growth rate in a cold plasma with linear gradient [75]. Afeyan and Williams carried out particularly complicated analyses of the instability considering oblique pump incidence, non-zero temperature and power law density profiles [76, 77].

2.2.4.2 The Raman gap

The lack of scattered light between 650 nm and 700 nm in Fig. 2.7 is known as the “Raman gap” and has long been present without being understood, given that nothing in the linear theory of Raman indicate it would not grow in this region of the plasma. Baldis *et al.* suggested that the presence of SBS in this region of the plasma may be suppressing Raman by depleting the pump [78], which is possible given they do not see any higher wavelength Raman light. However, in the spectrum in Fig. 2.7, the presence of Raman light at 700 nm indicates that the pump can not have been depleted.

One suggestion by Drake and Batha [79], based on results from their experiment and a model they subsequently built, was that the Langmuir decay instability or lack thereof was the cause of the Raman gap. They point out the LDI threshold increases with damping on each of its daughter waves, and thus that low damping on the LDI EPW will allow for significant Langmuir decay and cause low saturation levels of SRS. In their experiment they modified the collisional damping of the EPW across two targets and found the Raman gap only present in the target with lower collisional damping.

My proposed explanation is that rather than SRS being suppressed by other mechanisms, the scattered light created by SRS is being collisionally absorbed as it propagates out of the plasma. Longer wavelength light is generated closer to the laser’s quarter critical surface as the wave’s frequency approaches $\omega_0/2$. Equation (2.12) shows that a wave with frequency $\omega_0/2$ has a critical surface at $n_{c0}/4$ if n_{c0} is the critical surface of a wave that has frequency

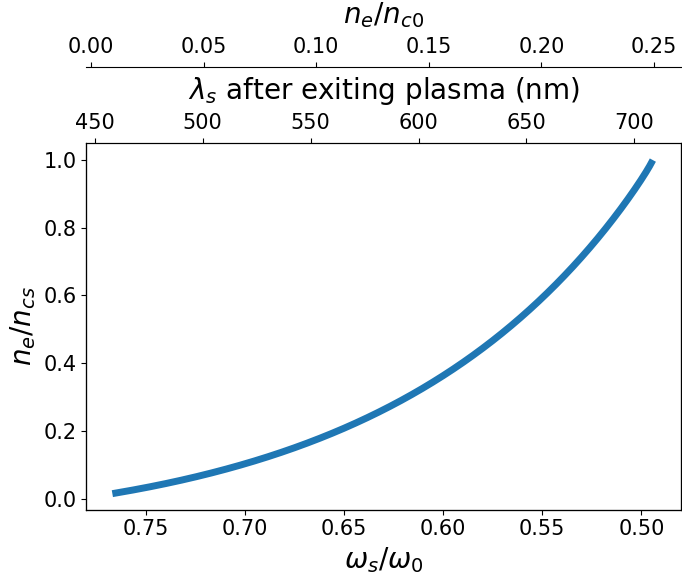


Figure 2.10: The density each Raman scattered ray is generated at relative to its own critical density (n_{cs}). The three x-axes are all parameters of the scattered ray. The bottom one is the frequency of the ray normalised to the laser frequency. The lower x-axis on top of the plot is the wavelength of the ray after it has left the plasma i.e. the wavelength that is measured. And the upper x-axis is the density the ray has been generated at relative to the laser's critical density (n_{c0}).

ω_0 . Figure 2.10 shows the plasma density that SRS rays are resonant at relative to their own critical density. Equation (2.40) shows that the rate that an EM wave is collisionally absorbed $\kappa_{coll} \propto \frac{\omega_{pe}}{\omega} \propto \left(\frac{n_e}{n_c}\right)^{1/2}$. Thus the higher wavelength SRS rays will be absorbed significantly more as they start at much higher relative densities before propagating out through lower densities.

This effect can be seen in the dotted lines in Fig. 2.6a that show the amplitude of three rays after amplification and absorption as they propagate out of the plasma. The higher the density the ray was seeded at, the larger the fraction of its energy was absorbed. The scattered ray seeded at $0.2n_c$ exits the plasma with lower intensity than the ray seeded at $0.15n_c$, despite undergoing larger amplification (the orange solid line is higher than the blue solid line).

Collisional absorption is also dependent on the distance a ray travels within the plasma, so at later times in experiments as the plasma has expanded further, there is a higher portion of these rays being collisionally absorbed. The ray path in the plasma being too short is likely why the rays in Fig. 2.6a do not display a full Raman gap. This effect is further discussed in Sec. 4.3.1.1. The simulation results using parameters output from a radiation-

hydrodynamics simulation show that at later times, when the plasma gets longer, the amount of Raman light absorption increases massively to the point where comparatively low levels of scattered light above 650 nm escapes the plasma. The hypothesis that the Raman gap is caused by collisional absorption of the scattered rays has been developed independently of our work and recently published, further corroborating our theory [80].

2.2.5 Stimulated Brillouin Scattering

Stimulated Brillouin scattering describes a pump EMW scattering off of an ion mode in the plasma. While the general process describing wave coupling and instability growth is very similar to SRS, but with ions in the place of electrons, the dynamics of the instability are quite different because of the change in mass of ions. The ion acoustic wave has a significantly lower frequency than both of the pump and scattered EMW, typically of the order 10^{13}s^{-1} compared to 10^{15}s^{-1} . Following the same method as at the start of the SRS section, ω_{iauw} can be neglected from the frequency matching expression to prove that SBS can occur when $\omega_0 > \omega_{pe}$ and thus when $n_e < n_c$. The presence of Brillouin scattering is also detected by its scattered light, except the wavelength is approximately that of the incident pump, so is distinct from other LPI signals, however it can be hard to separate from reflected or refracted laser light.

Most theoretical work on SBS has been regarding convective backscatter, which is focused on in this work. Sidescatter had been observed experimentally as far back as 1980 [81] and has more recently been noted as a potential source of SBS in ICF implosions [82]. However there has been relatively little experimental and theoretical work on it because the linear SBS growth rate maximises for backscatter, but it is possible this will change in the same way Raman sidescatter has become more of a concern with considerations over growth across the width of today's wider diameter laser beams.

Owing to the significant difference in the group velocities of the daughter waves and the fact that neither velocity gets as close to zero as with SRS waves, absolute SBS at the critical surface is unlikely. SBS could be undergo nonlinear saturation via two ion decay, where its IAW decays into two separate, non-resonant IAWs, however when testing this in the computational model (in a similar process to how LDI is used in Sec. 2.2.7) it made little impact as convective saturation always dominated in the regimes most applicable to ICF.




Figure 2.11: Four different thresholds for the definition of the strong coupling regime of SBS taken from [10–12]. (a) uses a $5 \times 10^{14} \text{W/cm}^2$ intensity pump, and (b) a $2 \times 10^{15} \text{W/cm}^2$ pump, both with $400 \mu\text{m}$ scale lengths. The black dashed line is at 1, the value that must be crossed for each threshold.

The SBS growth rate used in this work is derived by Forslund *et al* [11]:

$$\gamma_0 = \frac{v_{osc} k_{iaw} w_{pi}}{4\sqrt{\omega_s \omega_{iaw}}}. \quad (2.76)$$

This is equivalent to the SRS WKB growth rate except with electron terms swapped for ion terms, owing to the fact the amplification process is the same with an IAW instead of EPW. For backscattered SBS there is very little published theory on SBS without the WKB approximation. At very short scale lengths the condition for the WKB approximation in Eq. (2.38) does not apply at very high densities ($> 0.97n_c$). The reasons this model can ignore non-WKB SBS theory in that region are threefold. Most rays will not reach such high densities in the plasma as a result of refraction, those that do will have already undergone significant absorption as a result of inverse bremsstrahlung and other LPI, and then the rays will undergo resonant absorption in this region as a result of being at their turning point. Liu *et al.* [10] have developed theory on the growth rate for Brillouin sidescatter outside of the WKB regime.

Additionally, above a threshold pump intensity SBS enters the “strong coupling” regime [52], where the pump fields are so strong they prevent an IAW resonance, instead forming plasma “quasi-modes” [10]. The IAW evolving on a time scale faster than its frequency means the slowly varying envelope used to derive eqs. (2.60) and (2.61) for wave coupling are no longer valid. Figure 2.11 shows two thresholds provided by Liu *et al.* [10] and further

thresholds from Michel [12] and Forslund *et al.* [11] all of which much exceed 1 (the black line) to enter the strong coupling regime for SBS. Michel's threshold is defined when it is $\gg 1$, so for this plot a factor 100 has been used to qualify the inequality.

For both intensities tested significant portions of the pump ray are over threshold according to both of Liu's thresholds. Even at the lower intensity, over half of the pump ray is in the strong coupling according to Forslund. Michel's threshold, however, is never crossed (even if the factor 100 we use to qualify the \gg inequality is decreased) and this illustrates one of the most difficult parts of LPI modelling. Here there are three authors providing different thresholds (and also different growth rates) for the same process. There is a lot of different and untested analytic theory, often with differing levels of intelligibility. Ultimately, in this project SBS is only modelled in the weak coupling regime as there is limited concurrent theory on the strong coupling and its position as a higher order effect made it lower priority. It is however an effect that will be considered in future work given just how often the pump is above threshold.

2.2.5.1 Kinetic Spatial Amplification

Whilst the SRS model used the fluid expressions for its convective gain, these were found not to be appropriate for SBS. The analytically integrated models consistently failed to provide results of a reasonable order of magnitude, likely due in part to the approximation made in Sec.2.2.3.1 that the loss of coupling over space can be captured by the linear gradients in plasma parameters at the resonance. Instead, the kinetic expression for spatial amplification is used, where the spatial amplification rate for SBS is [12]:

$$\Gamma = \frac{k_{iaw}^2}{8k_s} \text{Im}[F_\chi] \left(\frac{v_{osc}}{c} \right)^2, \quad (2.77)$$

where $F_\chi = \chi_e(1 + \chi_i)/(1 + \chi_e + \chi_i)$. The parameters

$$\chi_s = \frac{-1}{2k^2 \lambda_{Ds}^2} Z' \left(\frac{\omega}{\sqrt{2}kv_s} \right) \quad (2.78)$$

describe the susceptibility of a plasma species s , with thermal velocity v_s , in response to a wave with frequency and wavevector ω and k respectively. In a multi-ion plasma, the ion susceptibility is the sum of the susceptibility of each ion species j : $\chi_i = \sum_j \chi_j$. $Z'(\zeta)$ is the derivative of the plasma dispersion [83]. Each scattered ray is amplified according to how

Γ for that ray evolves in space, with the plasma dispersion function capturing the effects of resonance as it peaks for supported electrostatic modes.

The beat frequency is also introduced:

$$\omega'_{iaw} = \omega_{iaw} - \mathbf{k}_{iaw} \cdot \mathbf{u}, \quad (2.79)$$

which describes the Doppler shifted IAW frequency in the frame of the plasma flow. The beat frequency is used in place of ω in the argument of the plasma dispersion derivative in Eq. (2.78).

At the point of resonance, x_R , both Eqs. (2.43) and (2.44) are satisfied. In this steady-state framework it is assumed that the frequencies of all three waves remain constant, however their wavevectors vary with plasma conditions. Thus, away from the resonance the IAW wave that the beating of the EM waves drives has $\mathbf{k}_{iaw}(x) = \mathbf{k}_0(x) - \mathbf{k}_s(x)$. This wave vector varying while the frequency of the IAW remains constant causes the loss of resonance as they no longer satisfy the dispersion relation for an ion mode that can be supported.

Writing out the IAW frequency as $\omega_{iaw} = c_s k_{iaw} + \mathbf{k}_{iaw} \cdot \mathbf{u}$, the beat frequency at a position x in space is then

$$\omega'_{iaw}(x) = c_s(x_R)k_{iaw}(x_R) + \mathbf{k}_{iaw}(x_R) \cdot \mathbf{u}(x_R) - \mathbf{k}_{iaw}(x) \cdot \mathbf{u}(x). \quad (2.80)$$

At the resonance this reduces to $\omega'_{iaw}(x_R) = c_s(x_R)k_{iaw}(x_R)$.

2.2.5.2 Crossed-beam energy transfer

Crossed-beam energy transfer (CBET) is a form of SBS where the instability arises from two overlapping laser beams, rather than a laser beam and a lower amplitude seed from noise within the plasma [84]. CBET is common in both direct and indirect drive ICF. In indirect drive it occurs at the entrance to the hohlraum where many of the beams overlap. It is now controlled to induce energy transfer between the different beams to optimise where energy is deposited in the hohlraum [28]. In direct drive it is particularly problematic. It occurs at the edges of beams where they refract around the edge of the target and into the path of other beams incident upon the target [85]. Significant fractions of energy can be transported away from the target by this process and potentially cause damage to laser optics as well.

In direct drive ICF, the laser's are all tuned to the same frequency, $\omega_0 = \omega_s$ so it may be




Figure 2.12: Wave vector diagram for the two-plasmon decay instability. The solid lines are the three wave vectors, with the EPW subscript being dropped from the daughter waves. The dotted lines are the components of the daughter waves parallel and perpendicular to the pump.

difficult to see from the frequency matching equation, Eq. (2.43) how an ion acoustic mode can be driven. The doppler shift term created by plasma flows in Eq. (2.26) will cancel out the natural IAW frequency term so in the lab frame the driven mode appears to have a zero or near-zero frequency.

CBET is not currently included in Chiron but is a major priority for including in the future. Without a CBET model it is difficult to apply this model to direct drive scenarios because it is so prevalent, and is also difficult to distinguish from normal SBS in the scattered light data.

2.2.6 Two-Plasmon Decay

The third laser-plasma instability that plagues ICF occurs when the pump EMW decays into two EPWs. Unlike with SRS and SBS, TPD can not exist as a one-dimensional instability in the backscatter configuration. To satisfy both EPW dispersion relations as well as the wave matching conditions the instability must be two dimensional. Figure 2.12 illustrates the general form of the TPD wavevectors. As well as an x component parallel to the pump, TPD EPWs have a y component perpendicular to the pump. TPD has one forward pointing EPW, that is labelled with subscript 1, that has a larger magnitude wavenumber than the pump. The EPW labelled with subscript 2 is backwards pointing and has a smaller wavenumber than the pump. Conserving momentum using Eq. (2.44) requires

$$|k_{epw1x}| = |k_0| + |k_{epw2x}| \quad (2.81)$$

$$|k_{epw1y}| = |k_{epw2y}|. \quad (2.82)$$

Different combinations of x and y components can satisfy the wavematching equations and dispersion relations, but Chiron only considers the set of waves that maximise TPD growth rate. The growth rate (Eq. (2.89)) is maximised along a hyperbola in k -space defined by [86]

$$k_{epw1x}^2 = k_{epw1y}(k_{epw1y} - k_0), \quad (2.83)$$

so it is assumed that TPD is growing from waves that satisfy this equation. The mode with the largest growth rate is expected to be dominant because it amplifies and depletes the laser before other modes can. Equation (2.83) provides the additional constraint needed to solve the TPD wave parameters after an extra degree of freedom was created by the two-dimensional wave vectors. Following the same logic used in the SRS section, TPD is also limited to densities below the quarter critical surface, because it creates two EPWs that must have $\omega_{epw} > \omega_{pe}$. TPD is mitigated very heavily by Landau damping, which normally prevents growth when $k_{epw}\lambda_{De} > 0.3$, so is typically not seen at densities below $0.2n_c$.

Figure 2.13a shows the evolution of the EPW wave parameters with density for a fixed temperature of 2.5 keV. Towards the quarter critical surface the magnitude of both wave vectors decreases sharply, leading to the failure of the WKB approximation. These dwindling wavevectors also mean that the EPW group velocities become very small, leading to regions of absolute growth approaching $n_c/4$. The perpendicular component of the wave vectors also tends towards zero in this region, so $k_{epw1} \rightarrow k_{epw1x}$. The wavevector triangle in Fig. 2.12 flattens as the forwards EPW gets closer to being parallel with the pump and the backwards EPW tends to zero. This pattern is also shown in Fig. 2.13b, which plots the angles between each EPW and the pump (as labelled in Fig. 2.12). TPD is often cited in the literature as having daughter wavevectors at $\sim 45^\circ$ to the pump but these results indicate significantly different and varying angles [87, 88].

The presence of TPD in experiments is diagnosed by characteristic light emission at $\omega_0/2$ and $3\omega_0/2$. Two-plasmon decay, unlike SRS and SBS, does not generate scattered light directly, instead generating two EPWs that in turn interact with the laser light or plasma

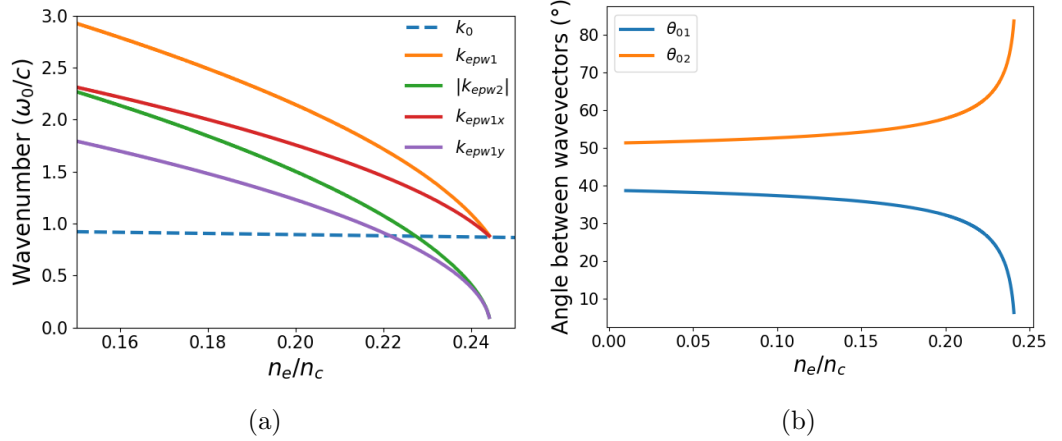


Figure 2.13: (a) TPD wavenumbers normalised to the vacuum wavenumber of the pump in a 2.5 keV plasma. (b) The angles (as defined in Fig. 2.12) between the wavevectors in (a).

waves and generate light at the two frequencies owing to three wave conservation laws. Often in ICF experiments only the $\omega_0/2 (= 2\lambda_0)$ light is measured as it is a similar wavelength to Raman scattered light. Two lobes of light at wavelengths above and below $2\lambda_0$ are detected and these can be created by a number of processes. The redshifted light can be caused by the small k_{epw2} limit of TPD very near $n_c/4$ as shown in Fig. 2.13a. Here $k_{epw1} \approx k_0$ and the beating of these two waves can also drive an EMW with $k \approx k_{epw2}$. The redshifted lobe can also be generated by inverse resonance absorption [89]. The blueshifted light can be generated by Thomson downscattering of the laser light of a TPD plasmon or by the inverse parametric decay instability, whereby an IAW beats with a TPD EPW to amplify a light wave [89]. The scattered light at $3\omega_0/2$ is caused by Thomson upscattering of the laser light and also shows a two lobes of red and blueshifted light. The small k_{epw2} redshifted light and the Thomson downscattered blueshifted light at $\omega_0/2$ and both components of the Thomson upscattered light at $3\omega_0/2$ obey three wave matching conditions so have a frequency $\omega_s = \omega_0 \pm \omega_{epw} \approx \frac{3}{2}\omega_0, \frac{1}{2}\omega_0$. The red and blueshift of the lobes is due to interactions with the two different EPWs which are slightly frequency shifted from $\omega_0/2$, as is demonstrated below. The multitude of mechanisms that can generate this light make it very difficult to model TPD amplitude based on the amplitude of the scattered light that is measured.

The wavematching conditions and dispersion relations can not be satisfied in the config-

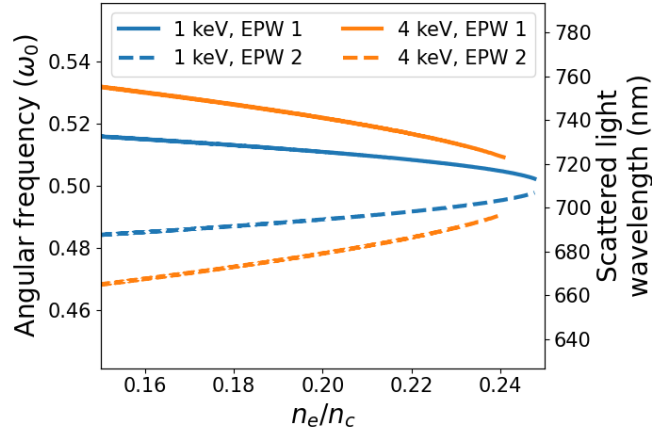


Figure 2.14: Wavelength of the two TPD EPWs as a function the density they are seeded at. The different colours are for two different pairs of EPWs formed in 1 and 4 keV plasmas. The solid lines are for the forward EPW and dashed for the backwards EPWs. The second y axis gives the measured wavelength of light that scatters off the EPWs. Experimentally, it is like that only EPWs at $n_e \geq 0.2n_c$ will be amplified due to the Landau cutoff.

uration in Fig. 2.12 if $\omega_{epw1} = \omega_{epw2}$. Instead there is some shift such that

$$\omega_{epw1} = \frac{\omega_0}{2} + \Delta\omega \quad (2.84)$$

$$\omega_{epw2} = \frac{\omega_0}{2} - \Delta\omega. \quad (2.85)$$

This shift can be quantified by first expressing

$$\begin{aligned} \omega_{epw1}^2 - \omega_{epw2}^2 &= \left(\frac{\omega_0}{2} + \Delta\omega\right)^2 - \left(\frac{\omega_0}{2} - \Delta\omega\right)^2 \\ &= 2\omega_0\Delta\omega. \end{aligned} \quad (2.86)$$

Expressing the same thing in terms of the EPW dispersion relation gives

$$\begin{aligned} 2\omega_0\Delta\omega &= (\omega_{pe}^2 + 3k_{epw1}^2v_e^2) - (\omega_{pe}^2 + 3k_{epw2}^2v_e^2) \\ &= 3v_e^2(k_{epw1}^2 - k_{epw2}^2). \end{aligned} \quad (2.87)$$

Using eqs. (2.81) and (2.82), it can be shown that $k_{epw1}^2 - k_{epw2}^2 = 2k_0k_{epw1x} - k_0^2$ and thus

$$\Delta\omega = 3\frac{v_e^2k_0^2}{\omega_0} \left(\frac{k_{1x}}{k_0} - \frac{1}{2} \right). \quad (2.88)$$

The frequency shift is linearly dependent on temperature and also has an indirect dependence




Figure 2.15: (a) The linear TPD growth rate (Eq. (2.89)) and the threshold for the instability to overcome wave damping (Eq. (2.55)) in blue. In orange is $k_{epw1}\lambda_{De}$. This example has $T_e = 2.5$ keV and $I_0 = 10^{15}$ W/cm 2 . (b) The Landau, collisional and total damping on the forward EPW from TPD using the same parameters as in (a).

on density from the presence of the wavevectors. These dependencies are shown in Fig. 2.14, which has the frequency of two pairs of TPD EPWs as a function of the density they were seeded for plasmas at 1 keV and 4 keV. The TPD frequency shift has been suggested as a method to diagnose plasma temperature near $n_c/4$, given the dependence on density but this might also require some knowledge of the density the EPWs are being generated at to constrain it. Seka [86] finds that the narrowband, redshifted light at $\omega_0/2$ amplified by the in the small k_{epw2} limit has a well-defined frequency shift and so is the most useful as a temperature diagnostic. The other light emitted at $\omega_0/2$ and at $3\omega_0/2$ (except in certain geometries) is broadband due to the spectrum of EPWs that TPD excites, and so is weaker as a temperature diagnostic.

2.2.6.1 Growth rates and damping threshold

The threshold for LPI growth in Eq. (2.55) depends on the product of the damping of each daughter wave. For SRS and SBS the threshold is weakly density dependent, so if the laser intensity is high enough to be significantly over threshold at one point in the plasma, the instability will likely be over threshold everywhere in that plasma. For TPD however, with both waves being EPWs, the Landau damping rate and with that, the threshold, vary over multiple orders of magnitude across a density profile with a constant laser intensity.

Figure 2.15a shows the linear TPD growth rate and the damping threshold as a function of plasma density for a pump intensity of 10^{15} W/cm 2 . In orange is the $k\lambda_D$ parameter for the

plasma. TPD is above threshold at densities over $0.205n_c$, which occurs at $k_{epw1}\lambda_{De} = 0.31$. This aligns with the generally accepted ‘‘Landau cutoff’’ for TPD, which claims the instability can not grow when $k\lambda_D > 0.3$ [90,91]. It should also be noted that Fig. 2.15a shows that at very low densities the TPD growth rate is once again above the threshold, which has dipped dramatically. However, there is no experimental evidence that TPD ever occurs at such low densities.

Figure 2.15b shows the contribution of Landau and collisional damping to the total damping on the forward TPD EPW. Landau damping is the dominant mechanism except for near the quarter critical surface, where collisions become dominant. The Landau damping rate also significantly dips at very low densities and this is why the TPD threshold does the same. Seaton [92] suggests that at large $k\lambda_D$ the linear Landau damping rate derived from fluid theory, Eq. (2.19), is no longer valid. The EPW dispersion relation, Eq. (2.18) is also derived in the fluid limit, which breaks down in these parts of the plasma where density is so low that kinetic behaviour takes over. Given that there is no experimental evidence of low density TPD and the damping rate that implies it might exist is based on theory which may be invalid in that regime, the $k\lambda_D > 0.3$ threshold is also applied. Therefore, TPD is limited to only the high density region where it is over the wave damping threshold.

The linear growth rate for TPD derived under the WKB approximation is [52]

$$\gamma_0 = \frac{|\mathbf{k}_{epw2} \cdot \mathbf{v}_{osc}|}{4} \frac{|k_{epw1}^2 - k_{epw2}^2|}{|\mathbf{k}_{epw1}| |\mathbf{k}_{epw2}|}. \quad (2.89)$$

For wavevectors that satisfy Eq.(2.83) this reduces to [93]

$$\gamma_0 = \frac{k_0 v_{osc}}{4}. \quad (2.90)$$

The dot product involving \mathbf{v}_{osc} in Eq. (2.89) means that TPD grows in the plane of the pump’s polarisation. For example, for Fig. 2.12 to make sense the pump must be polarised in y to maximise the dot product with k_y .

Two-plasmon decay can also exist as a convective instability, where it was recently shown that it requires a slightly different growth rate to be used within spatial amplification ex-

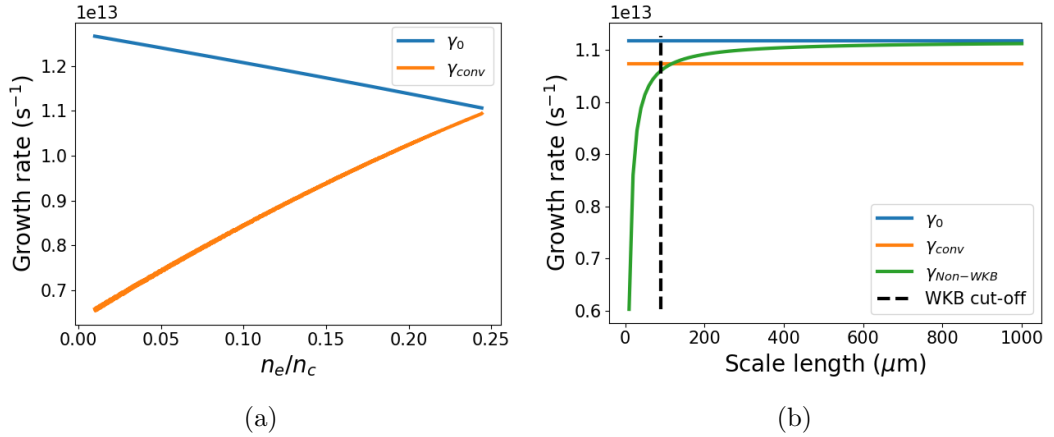


Figure 2.16: (a) The linear TPD growth rate and the modified convective growth rate in a 2.5 keV plasma. (b) The linear, convective and non-WKB growth rates as a function of density scale length at $0.235 n_c$ and 2.5 keV. The black dashed line indicates the scale length below which the WKB approximation (Eq. (2.38)) fails and it becomes appropriate to use the non-WKB growth rate.

pressions like Eqs. (2.69), (2.70) and (2.71) [94]:

$$|\gamma_{conv}|^2 = - \left(\frac{k_{epw1y} v_{osc}}{4} \right)^2 \left(\frac{\omega_{epw2}}{\omega_{epw1}} - \left(\frac{k_{epw2}}{k_{epw1}} \right)^2 \right) \left(\frac{\omega_{epw1}}{\omega_{epw2}} - \left(\frac{k_{epw1}}{k_{epw2}} \right)^2 \right). \quad (2.91)$$

This is a temporal growth rate rather than spatial, it is labelled as a convective growth as Yan *et al.* [94] specifically designate its use in convective gain expressions. This is compared to the linear growth rate in Fig. 2.16a.

For TPD nearer the quarter critical surface, the wavevector gradients get steeper while the magnitudes get smaller meaning the WKB condition in Eq. (2.38) does not hold, so the growth rates given so far are not valid. Liu and Rosenbluth [93] derived the growth rate for TPD in a plasma with a linear gradient without using the WKB approximation:

$$\gamma_{non-WKB} = \frac{k_0 v_{osc}}{4} - \frac{\omega_{pe}}{2k_{epw1y} L_n}. \quad (2.92)$$

This is the maximum linear growth rate minus a term dependent on the scale length.

Figure 2.16b shows this plotted against both previous TPD growth rates as a function of density scale length at $0.235 n_c$. As expected, in the long scale length plasmas it converges with the linear growth rate, while diverging as the scale length decreases. This is important in capturing the fact that very short scale length plasmas suppress LPI growth because the amplification lengths are so short. This graph also shows the cut-off for the WKB condition

that was set at a factor 100 between $|k|^2$ and $|dk/dx|$. In his recent book, Liu [95] also recommends subtracting the EPW damping rates (collisional and Landau) from this growth rate that he derives. Slightly confusingly, he chooses not to subtract any wave damping rates for his absolute SRS growth rate that he derives without the WKB approximation (Eq. (2.75)). This is also in contradiction to how Nishikawa and other authors have described the effect of wave damping on parametric growth rates, which is shown in Eq. (2.54) [65]. For the absolute instability models in Chion, which are detailed in the next section, Liu's subtraction of damping is disregarded and Nishikawa's more commonly cited effects of damping on growth are pursued.

Simulations have shown [67] that convective and absolute TPD modes can occur in the same plasma, with absolute modes occurring at $n_c/4$ and convective modes at lower densities. However, simulation work by Zhang *et al.* [96] has shown that in fusion plasmas with multiple overlapping lasers, the multibeam TPD instability is driven and this sees very low convective gains in comparison to the absolute instability.

Currently the only nonlinear saturation for TPD included in Chiron is the Langmuir decay instability—the secondary decay of a large amplitude EPW to another, non-resonant EPW and an IAW. The model for LDI saturation is described in Sec. 2.2.7. Whilst LDI had been shown to be a key saturation mechanism for single beam TPD [97, 98], multibeam two-plasmon decay has been shown to be more strongly saturated by profile modification than LDI [99]. The multitude of EPWs that are driven at the quarter critical surface beat together, driving plasma out of the region via the ponderomotive force, thus stopping LPI growth at $n_c/4$ by reducing the plasma density there. Mode coupling describes the coupling of TPD EPWs into IAWs normally generated by SBS at the quarter critical surface and is another important saturation mechanism that has been measured in experiments [87, 100]. Developing models for both mode coupling and particular profile modification will be key for creating a more complete picture of TPD growth and saturation.

2.2.7 Langmuir Decay Instability and absolute instability saturation

The Langmuir decay instability (LDI) describes an EPW (aka Langmuir wave) decaying into an IAW and a separate EPW. It is one of the dominant mechanisms for saturating absolute growth of SRS and TPD as the rapidly growing EPW decays via LDI and the new LDI generated EPW is frequency shifted so that it is not resonant with the primary instability.

LDI is a three wave instability that must satisfy dispersion relations for the three waves as well as the wave matching conditions. Riconda and Weber [101] provide a shortcut to solving all of these equations with the approximation that

$$|k_{epw}^{ldi}| = |k_{epw}| + \delta k \quad (2.93)$$

$$|k_{iaw}^{ldi}| = 2|k_{epw}| - \delta k, \quad (2.94)$$

where k_{epw} is for the decaying, initial EPW and

$$\delta k = \frac{2}{3\lambda_{De}} \sqrt{\frac{Z^* m_e}{m_i}}. \quad (2.95)$$

They give the growth rate of the instability to be

$$\gamma_{ldi} = \frac{eE_{epw}}{2m_e v_e \omega_{epw}} \sqrt{\omega_{epw}^{ldi} \omega_{iaw}^{ldi}}, \quad (2.96)$$

where ω_{epw} and E_{epw} are the frequency and electric field amplitude of the parent EPW and the two frequencies in the square root are for the daughter waves. The growth rate depends on the amplitude of the decaying wave, i.e. as the EPW grows larger due to SRS or TPD it decays faster as a result of LDI. This is used to derive the level that LDI saturates an instability at. As an EPW is growing in amplitude as a result of SRS the rate it is decaying via LDI increases until the point where $\gamma_{srs} = \gamma_{ldi}$ is reached. The rate of growth and decay have balanced each other out and the wave can neither grow or decay away from this point without returning back to it. By setting the SRS growth rate equal to Eq. (2.96), the amplitude that the EPW saturates at as a result of LDI is

$$E_{epw} = \frac{2m_e v_e \omega_{epw}}{e} \frac{1}{\sqrt{\omega_{epw}^{ldi} \omega_{iaw}^{ldi}}} \gamma_{srs}. \quad (2.97)$$

Unlike with the convective model, the level this model saturates at is not dependent on the amplitude of the waves that seed the instability. In this case, regardless of the initial seed amplitude, LDI will saturate the EPW at the specific amplitude that causes the LDI growth rate to equate to the SRS growth rate. The implementation of the LDI model and how it works in tandem with the convective saturation model is described in Sec. 3.1.6.

Chapter 3

Computational Methodology

This chapter describes how the theory presented in chapter 2 was developed into computational models capable of predicting LPI levels in three-dimensional simulations. Section 3.1 describes the initial LPI models developed for a single, one-dimensional ray. Pump depletion, integrating absolute and convective saturation, and combining multiple instability models were all addressed in this reduced form of the models.

Section 3.2 outlines how the one-dimensional models are used alongside three-dimensional outputs from radiation-hydrodynamics simulations with physics additions to account for interactions between multiple laser rays.

3.1 A Single Ray

To begin, a single instability happening along a single ray was considered. As the plasma profiles were all manually defined, the simulations had high spatial resolution, often thousands of cells for a ray with a length on the order of millimetres, which differs from the lower resolution provided by the mesh of a radiation-hydrodynamics simulation. Each cell along the ray has the plasma parameters density, temperatures and flow velocity and from these further parameters like gradients, plasma frequency, thermal velocity, and the Debye length can be derived. Next, using the parameters specific to each cell, the wave parameters for an instability seeded there can be calculated, such that both the resonance conditions and required dispersion relations are satisfied.

3.1.1 Wave Parameters

This method differs between each instability. Looking at SRS, there are 4 unknown parameters, the frequency and wavevector of each daughter wave, and there are 4 equations that link these parameters. Given that only Raman backscatter is currently modelled, the wavevector matching condition reduces to

$$|\mathbf{k}_0| = |\mathbf{k}_{epw}| - |\mathbf{k}_s| \quad (3.1)$$

to account for the pump and EPW being parallel with the scattered wave antiparallel. These parameters are solved for iteratively, starting with an initial guess for \mathbf{k}_{epw} and cycling through the equations to end up with a new value for k_{epw} , repeating this until it converges. The iteration order is arbitrary and only depends on which parameter is started with as the guess, but the current one is as follows:

$$k_{epw} = 2k_0$$

Open iteration loop

$$\omega_{epw} = (\omega_{pe}^2 + 3k_{epw}^2 v_e^2)^{1/2}$$

$$\omega_s = \omega_0 - \omega_{epw} \quad (3.2)$$

$$k_s = \frac{1}{c} (\omega_s^2 - \omega_{pe}^2)^{1/2}$$

$$k_{epw} = k_0 + k_s$$

Restart loop with new k_{epw} unless converged.

This process can be further sped up by starting with a better guess for k_{epw} . Vu *et al.* give

$$k_{epw} \lambda_{De} = 0.045 \left(\frac{n_c}{n_e} \right)^{1/2} [T_e(\text{keV})]^{1/2} \times \left[\left(1 - \frac{n_e}{n_c} \right)^{1/2} + \left(1 - 2 \left[\frac{n_e}{n_c} \right]^{1/2} \right)^{1/2} \right], \quad (3.3)$$

which is found by expanding the dispersion relations with consideration for the resonance conditions [102]. Iterating with this as a starting point reaches convergence of SRS wave parameters to a difference of 0.1% within typically less than five iterations.

The SBS wave parameters are found in a similar way to those of SRS, again iterating through the four equations, just using the different dispersion relations for IAWs in a

one or two ion plasma. Whereas with SRS and SBS backscatter has been assumed and the wavevector matching equation can be simplified, TPD is a three-dimensional instability and the relative angles of the wavevectors can not be assumed. The wavevector matching condition is now

$$k_{epw2}^2 = k_0^2 + k_{epw1}^2 - 2|k_0||k_{epw1}|\cos(\theta_{01}), \quad (3.4)$$

where θ_{01} is the angle between the pump and forward pointing EPW wavevectors (see Fig. 2.12), which, x is defined as the direction as parallel to k_0 , can be expressed in terms of the components of k_{epw1} :

$$\theta_{01} = \arctan\left(\frac{k_{epw1,y}}{k_{epw1,x}}\right). \quad (3.5)$$

It is also necessary to find the perpendicular component as it defines the TPD growth rate. Considering the components parallel to the pump $|k_0| = |k_{epw1,x}| - |k_{epw2,x}|$, so $|k_{epw1,x}| > |k_0|$ always.

Finding an iterative solution for the wave parameters comparable to that for SRS and SBS was not successful. Instead, a look-up table was created for each of the wave parameters as a function of density and temperature. The method for creating the table was borne from attempts to solve the equations iteratively. Rather than starting with an initial guess for one variable (in this case $k_{epw1,x}$), it starts with a range of possible values that are iterated through, finding which value best satisfies the wavematching equations and dispersion relations. Satisfaction of the equations is measured by taking two approaches to finding $k_{epw,2}$ from each initial guess and finding which guess results in the two values for $k_{epw,2}$ matching. This was done for each combination of temperature and density. The method is outlined below.




Figure 3.1: The difference between the two calculated values of k_{epw2} for each starting guess of $k_{epw1,x}$ at 4 different pairings of plasma density and temperature. Normalising $k_{epw1,x}$ to k_0 means the x axis is equivalent to α in Eq. (3.6). The black dashed line is at zero, where the values are equal.

Start with some n_e and T_e

Open loop over possible values of $k_{epw1,x}$

$$\begin{aligned}
 k_{epw1,x} &= \alpha k_0 \\
 k_{epw1,y} &= [k_{epw1,x} (k_{epw1,x} - k_0)]^{1/2} \\
 \theta_{01} &= \arctan\left(\frac{k_{epw1,y}}{k_{epw1,x}}\right) \\
 k_{epw1} &= (k_{epw1x}^2 + k_{epw1y}^2)^{1/2} \\
 k_{epw2a} &= (k_0^2 + k_{epw1}^2 - 2|k_0||k_{epw1}|\cos(\theta_{01}))^{1/2} \\
 \omega_{epw1} &= (\omega_{pe}^2 + 3k_{epw1}^2 v_e^2)^{1/2} \\
 \omega_{epw2} &= \omega_0 - \omega_{epw1} \\
 k_{epw2b} &= \left(\frac{\omega_{epw1}^2 - \omega_{pe}^2}{3v_e^2}\right)^{1/2}
 \end{aligned} \tag{3.6}$$

This process is depicted in Fig. 3.1. As values of $\alpha > 1$ are iterated through, the difference between the two values of k_{epw2} decreases. At the point where $k_{epw2a} = k_{epw2b}$ the wave parameters satisfy all of the dispersion relations and wavematching equations so are saved into the look-up table for that density and temperature. The look up table currently contains data for $0.05 \leq T_e(\text{keV}) \leq 10.0$ and $0.01 \leq n_e/n_c \leq 0.25$. When a simulation including TPD

is run, the density and temperature at each cell are linearly interpolated between the two nearest density and temperatures points in the table, yielding the required wave parameters.

3.1.2 Growth and Saturation of a Single Instability

Once the instability's wave vector and frequency have been found, they can be used to seed a three wave coupling anywhere along the pump ray. For a single Raman scattered EMW or TPD EPW, Chiron assumes that the ray has reached convective or absolute saturation based on its parameters at resonance. The dominant mechanism is chosen by whichever saturates the ray at a lower intensity. For Brillouin scattered rays the spatial amplification is modelled by tracking the plasma parameters along the propagating scattered ray.

As discussed in section 2.2.3, a scattered ray will grow convectively over the course of its path as it leaves the plasma, with the growth depending on the evolving plasma parameters along the ray. Given that the Rosenbluth and bounded gain expressions (Eqs. (2.69)) and (2.70)) have been found by integrating the path of a ray to infinity, the evolution of plasma parameters as the scattered ray convects can not be accounted for. Instead the parameters at a single point, which is chosen to be the centre point in space between the seed and the boundary of one amplification region (defined as L_d in Eq. (2.68)), are used to calculate the gain of the scattered ray. This is chosen because the majority of spatial growth happens in this region, i.e. close to the resonance where the wave coupling is strongest. This method leads to a significant increase in computational efficiency because a single calculation is required for the LPI levels. The numerically integrated method using Eq. (2.71) requires not only multiple calculations of the amplification as it is integrated along the ray's path, it also requires many further calculations as the evolution of various wave parameters are tracked in space. See section 3.1.5 for a comparison of the results of the two models for convective SRS, which shows acceptable performance from the faster bounded gain method.

Results were poor with the convective SBS model for all of the above fluid gain methodologies, so the kinetic spatial amplification (Eq. (2.77)) is used. This must be integrated along ray paths, with the methodology similar to the DEplete code developed by Strozzi [2] for Raman and Brillouin backscatter. The code ELPSE [74] numerically integrates the fluid gain for stimulated Raman sidescatter. It is likely that a next step for Chiron will be incorporating numerically integrated amplification methods for SRS as these promise to provide more accurate results in capturing the extent of resonances and possible second

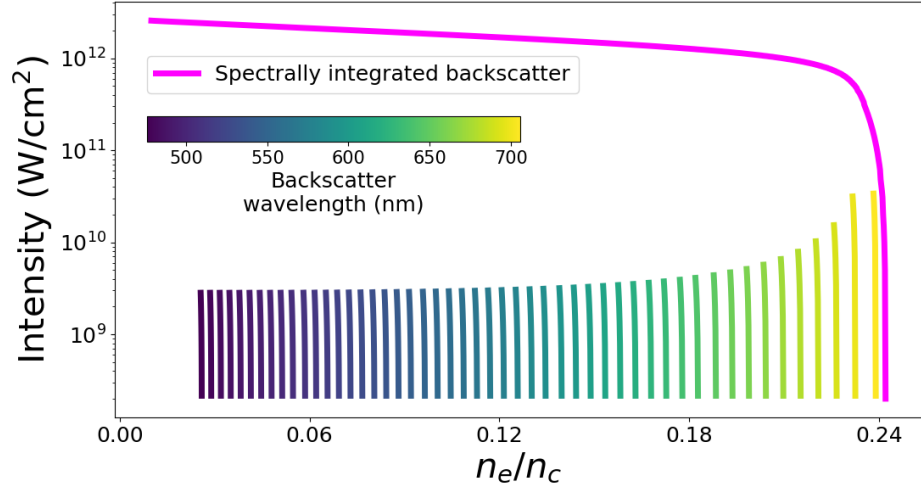


Figure 3.2: Growth and saturation of Raman backscattered rays. Each ray is coloured by its wavelength when it leaves the plasma. The pink line is the spectrally integrated Raman backscattered light, found by summing all of the individual rays (including those not pictured here). This example was for $I_0 = 10^{15}$ W/cm², $T_e = 2.5$ keV and an exponential gradient with $L_n = 500\mu\text{m}$. Rays are saturated convectively and via LDI.

resonances along a ray.

As a pump ray propagates through different densities and temperatures there will be different pairs of daughter waves resonant with it for each instability, as can be seen from the wide spectrum of SRS scattered light data in Fig. 2.7. For one instability, many pairs of daughters waves are launched along the pump ray, with each set of daughter waves' starting parameters known from section 3.1.1. Considering just SRS, Fig. 3.2 shows 50 scattered rays as they amplify and propagate out of the plasma following the approximation that total amplification occurs over one amplification length. Each coloured line is one ray, growing over the short distance L_d from its seed value to the value it saturates at due to convective dephasing or Langmuir decay. The model then propagates the rays out of the plasma while undergoing no further amplification. The colour of each ray corresponds to its wavelength when it exits the plasma rather than its wavelength at birth. The range of the spectrum is comparable to those seen in scattered light data from experiments. The pink line traces the intensity of the spectrally integrated backscatter, which is the total intensity of scattered light leaving the plasma at each point in space.

The convectively saturated amplitude obeys the expression $I_s = I_{seed}\exp(G)$, where G is the bounded gain from Eq. (2.70). Currently, Chiron does not have a physics-based model to predict seed amplitude, but Strozzi [2] details two approaches to model electromagnetic

noise from Thomson scattering and bremsstrahlung emission in the plasma. The amplitude of the broadband noise sources, i_{noise} , are in terms of spectral density such that

$$I_{seed} = \int i_{noise} d\omega. \quad (3.7)$$

The spectral density is the intensity per angular frequency. Therefore, the seed amplitude is found by integrating this over the bandwidth of each seed. This is essential to keeping results invariant to the number of scattered rays launched. Changing the number of scattered rays now will change the bandwidth of each seed keeping the total seed amplitudes across regions of the plasma constant. Due to the lack of appropriate seed models, the approach taken by Hao *et al.* [13] and Bénisti *et al.* [4] in their inline LPI models is followed. They use a constant seed amplitude across the pump with it's amplitude being some fraction of the pump. In a typical example, the total amplitude of all of the seeds along the pump might be

$$\begin{aligned} I_{seed}^{tot} &= \sum I_{seed} \\ &= 10^{-4} I_0 \end{aligned} \quad (3.8)$$

From this, each value of I_{seed} can be calculated using the frequency bandwidth of each scattered ray and the assumption that i_{noise} is constant with respect to frequency.

Although 50 rays are shown in the Fig. 3.2, 1000 rays were actually launched for the simulation, with extras being left out for a clearer image. The number of rays to launch is an important question for the model with a balance needing to be struck; it needs enough rays so that the plasma parameters are sufficiently sampled along the pump, but the fewer rays launched the more computational resource is saved. This can be measured by increasing the number of scattered rays that launched until the results for reflectivity converge.

Figure 3.3 shows how the reflectivity of a single pump ray varies according to the number of scattered rays launched along it for 3 different pump intensities and 3 different density scale lengths. Figure 3.3a shows that there is a clear correlation between the number of scattered rays required for the reflectivity to converge and the total reflectivity. Simulations with higher intensities and longer scale lengths result in larger Raman reflectivities, requiring more scattered rays to give a converged answer.

In this graph, reflectivity, which is defined as the ratio between the scattered and pump wave intensities, is often larger than one, which should be impossible if energy is conserved.

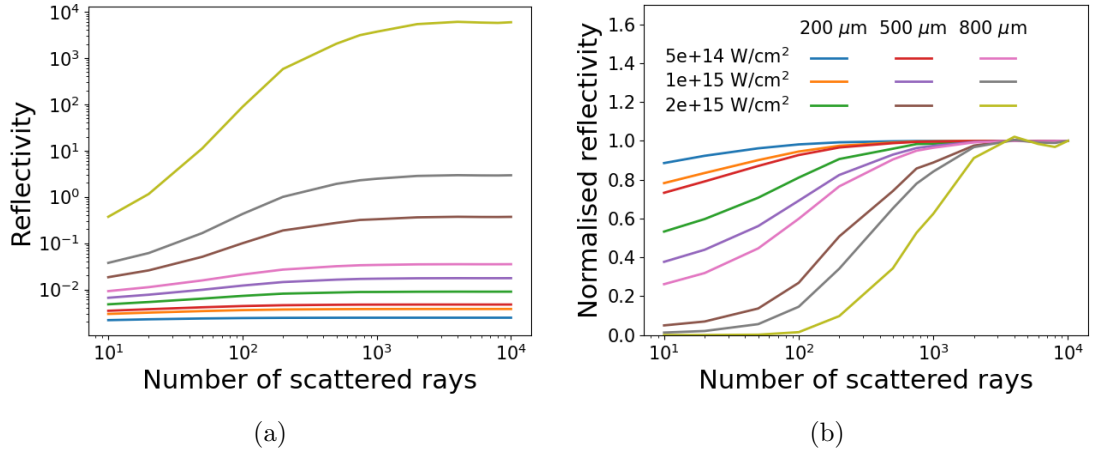


Figure 3.3: (a) Variation of reflectivity with number of scattered rays launched for 9 simulations with different combinations of pump intensity and density scale length. (b) The same data but with the reflectivity of each pair of parameters normalised to the pair’s reflectivity when 10,000 rays are launched. These simulations used the bounded gain to convectively amplify rays and ignored the effects of pump depletion.

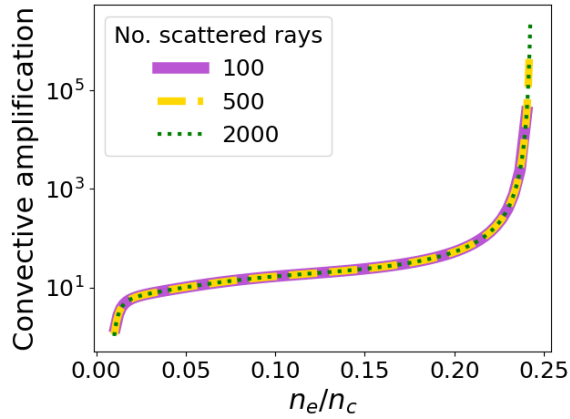


Figure 3.4: The convective amplification for rays seeded at each density for simulations with pump intensity $2 \times 10^{15} \text{ W/cm}^2$, scale length $500 \mu\text{m}$ and varying numbers of Raman scattered rays.

Pump depletion is not included in these simulations, so a constant intensity pump drives all of the LPI growth. This choice is made for the purpose of comparison as it reduces the number of variables between simulations — now identical pump intensity is driving the growth of all of the waves in all of the simulations. Ignoring pump depletion for the purpose of comparison is discussed in more detail in Sec. 3.1.5.

Figure 3.3b shows that the higher reflectivity simulations will require over 1,000 rays before the reflectivity gets to within 10% of its converged value. Simulations resulting in higher reflectivities than those presented here may even start to push the requirement to

10,000 scattered rays per pump ray. When modelling 3D ICF experiments each laser beam will potentially be made up of 1,000s of individual pump rays and there can be up to 192 beams if modelling a NIF experiment. This could end up requiring massive computational expense for the launching of tens or even hundreds of millions of scattered rays to get an accurate picture of the LPI physics from only backscattered SRS at one point in time of a simulation. Factoring in the other two instabilities and the desire to run the code multiple times during a hydrodynamic simulation to capture the time evolution of the LPI and the picture gets even worse for fast computation.

To reduce the necessary number of scattered rays, the spacing between rays was changed. The simulation detailed above had spaced the rays evenly in frequency space such that each had the same bandwidth, and spacing them evenly in density space provided almost identical results. The important parts of the pump to model with high resolution were those where gain was highest, and conversely, where gain was low, the resolution could decrease as it would not affect total reflectivity. For SRS, convective gain roughly increases from low density all the way to the quarter critical surface. Now the rays were positioned in density space such that 20 % of the rays were seeded below $0.1 n_c$, 40% were between $0.1 n_c$ and $0.2 n_c$ and the remaining 40% were above $0.2 n_c$. This managed to reduce the number of scattered rays required until reflectivity was within 10% of its converged value by a factor of ~ 3 , however this was not the multiple order of magnitude reduction we were hoping for.

In truth, the reflectivities were expected to converge at much smaller numbers of scattered rays than they do in Fig. 3.3; needing almost a thousand scattered rays to fully sample the parameters in plasmas of these lengths seemed excessive. Figure 3.4 shows the amplification of each ray based on the density it was seeded at for three simulations that launch different numbers of rays. This graph reveals that the problem arose from sampling at the upper limit of the SRS density range. The convective gain increases exponentially closer to $n_c/4$, so as the spatial resolution of the rays increases, rays are created that see these small regions of much higher gain. In the figure this can be seen by the simulations with more rays having rays at slightly higher density which amplify significantly more than the highest density rays in the other simulations. Section 3.1.6 shows how this problem is resolved, and thus how the necessary number of scattered rays is reduced: absolute saturation by the Langmuir decay instability takes over at the quarter critical surface, preventing the runaway convective gain.

3.1.3 Pump depletion

Throughout chapter 2, when deriving expressions for wave growth rates and saturation, only the evolution of the two daughter waves were considered when modelling their amplitudes. The third coupled equation, which described the pump’s evolution, was ignored as the pump amplitude was treated as constant in the coupled equations for the daughter waves’ amplitudes (Eq. (2.46)). In reality, the evolution of the pump is central to the behaviour of the three-wave interaction. It is well known that at higher pump intensities LPI grow faster, hence why many ICF ignition schemes have been looking towards lower intensity laser pulse shapes [7, 103]. This is clear when examining the growth rates for the instabilities, which all depend linearly on v_{osc} and thus on the pump wave’s field amplitude. The nonlinear relation between the wave amplitudes is one of the reasons LPI are difficult to model — as the daughter waves grow because of the pump fields, they in turn deplete energy from the pump fields and limit their own growth. Solving the time-dependent coupled equations would be one way to track how the three waves’ amplitudes evolve because of each other while stepping in time. Chiron however, works in a quasi steady-state regime in which the balance between growth and damping is assumed to occur, i.e. the instability is said to be saturated. At first, the saturated level of an instability is determined by the initial level of the pump without considering pump depletion, which will inevitably lead to an overprediction of LPI levels.

Figure 3.5 shows the iterative process employed to calculate pump depletion. In this simplified example, only SRS is detailed. The simulation starts with a pump ray of constant intensity propagating through underdense plasma (top panel) and from this profile the amplitude of scattered light due to SRS occurring along the ray is calculated (middle panel). By finding the increase in action density in the scattered SRS waves, this quantity is conserved by depleting the pump by the same amount; this depleted pump profile is in the bottom panel. Thus the problem is clearly outlined: the SRS growth has been predicted by a pump profile of constant intensity, but this growth would actually deplete the pump significantly. The “true” steady-state solution is a pump profile somewhere in between the top and bottom panels of Fig. 3.5, where the SRS growth from this true pump would deplete the initial constant pump to the same true pump profile.

Chiron’s solution, which is shown in Fig. 3.6, is to search for this saddle point-like pump profile by iterating over the LPI calculation depicted in Fig. 3.5. The depleted pump profile from the bottom panel is effectively moved to the top panel, starting the calculation again

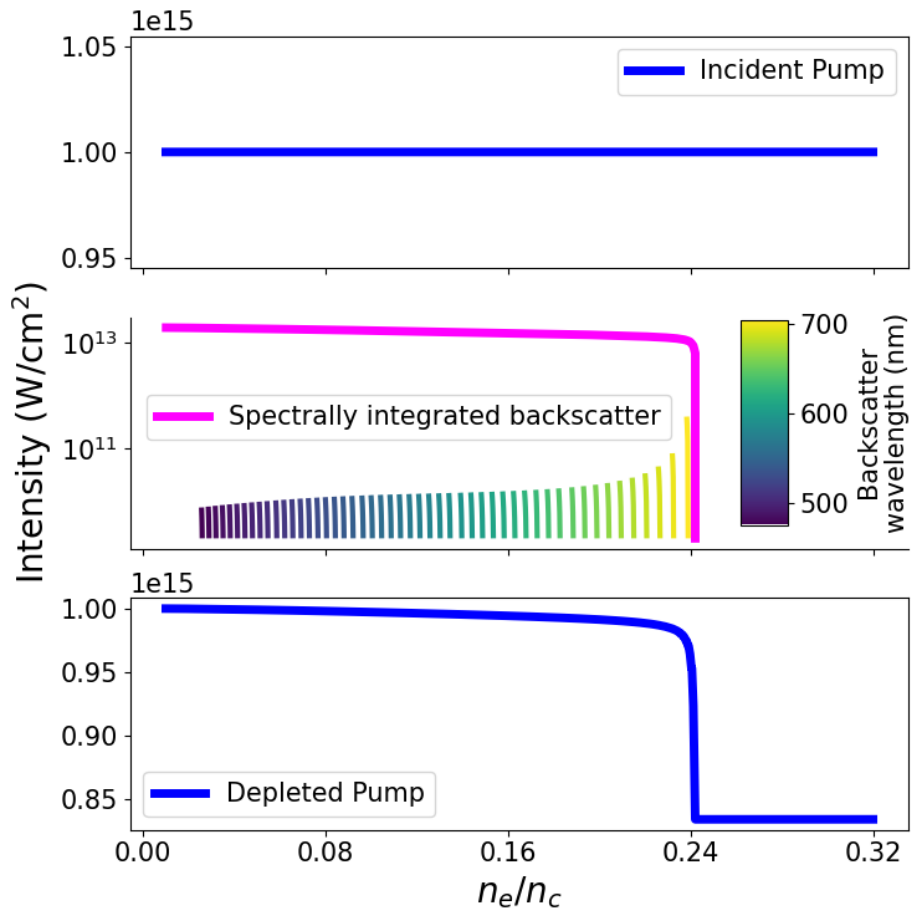


Figure 3.5: Simulation results considering only convective SRS along a single pump ray with intensity $1 \times 10^{15} \text{ W/cm}^2$ in a plasma with scale length $700 \mu\text{m}$. SRS growth rates are calculated from an initial constant intensity pump (top), allowing us to calculate saturated amplitudes of Raman scattered light rays (middle) and then deplete the initial pump by conserving action (bottom). The middle panel shows 50 selected Raman rays with their measured wavelengths colour coded as well as the total intensity of scattered light in pink.




Figure 3.6: Intensity profile of the pump ray before doing LPI calculations (iteration 0) and after each iteration of the LPI calculation. The profile converges to within 1% of the previous iteration after the fourth iteration.

with it. The LPI levels are calculated according to this new pump and the initial incident pump is depleted again, creating a new depleted pump profile. In Fig. 3.6 this is the pump listed as iteration 2, that, as anticipated, falls in between the profiles which are iteration 0 (the initial profile) and iteration 1. Using the new second iteration pump the process is repeated to get a third iteration pump profile and so on until the pump profile converges, i.e. the saddle point is found. In this scenario it only took 4 iterations for the profiles to have $< 0.1\%$ difference between them. The number of iterations taken varies depending on how steeply the pump is depleted.

3.1.4 Defining Reflectivity

In the previous section's simulation only convective amplification of rays was considered and as a result the pump is primarily depleted at the quarter critical surface, where convective saturation is weak. It is notable in Fig. 3.5, that after this first iteration, and by defining reflectivity as the ratio of scattered intensity and pump intensity, the reflectivity is 1.4% and the pump has been depleted by 15.8%. The pump is depleted by conserving action density (Eq. (2.35)), i.e. finding the change in action necessary for the scattered light to reach its saturated level and then subtracting this action from the pump wave. Using $I = v_g \omega A$ and

$\delta A_0 = \delta A_s$, the change in pump intensity as caused by a single scattered ray is:

$$\begin{aligned}\delta I_0 &= v_{g0}\omega_0\delta A_0 \\ &= \frac{v_{g0}\omega_0}{v_{gs}\omega_s}\delta I_s,\end{aligned}\tag{3.9}$$

where the 0 and s subscripts signify quantities belonging to the pump and scattered wave respectively.

Figure 3.7 plots this ratio of group velocities and frequencies for SRS in the conditions used for Fig. 3.5. Closer to $n_c/4$ this ratio spikes. Equation (3.9) shows that for every unit increase in scattered wave intensity, the pump wave will be depleted by a factor of almost 18 times more. For the SRS example presented, the majority of the scattered light was generated near the quarter critical surface so this explains the large, aforementioned difference between reflectivity and depletion.

Experimentally, reflectivity from SRS and SBS is found by comparing the total energy that is measured by scattered light diagnostics to the initial energy in the laser, whereas this model works with waves in terms of their intensity and power. Reflectivity is defined as the ratio of scattered intensity (or power) to the initial pump's intensity (or power), as if there is a diagnostic outside of the plasma measuring these quantities. The next important point to consider with measuring intensity and power of the scattered light is that these quantities




Figure 3.7: The variation of the coefficient from Eq. (3.9) with plasma density. After a scattered EMW at some density has been Raman amplified the pump is depleted by the product of the scattered wave's intensity and the coefficient.

actually vary as the wave propagates out of the plasma despite the amplitude of the wave remaining constant. The intensity of a wave is the rate that the energy it is carrying passes through a surface, so if two waves with equal energy are travelling with different speeds, the one travelling faster has a higher intensity (and power). With electromagnetic waves this is clear from the description of intensity in Sec.2.1.2.5 as $I = uv_g$. Rewriting the energy density for an EMW gives

$$I_{emw} = \frac{v_{g,emw}\epsilon_0}{2}|E_{emw}|^2, \quad (3.10)$$

so the intensity of a wave with constant amplitude varies with its group velocity. This model assumes amplification of a wave saturates locally in space, so as the scattered wave continues to propagate out of the plasma, its amplitude is constant, with only its group velocity changing. The intensity of a wave at some final position is thus

$$\begin{aligned} I^f &= \frac{v_g^f \epsilon_0}{2} |E|^2 \\ &= \frac{v_g^f}{v_g^i} I^i, \end{aligned} \quad (3.11)$$

where the i and f superscripts refer to the initial and final versions of those quantities. For reflectivity, we are interested in the intensity of the scattered wave once it has left the plasma so must have a group velocity equal to c :

$$\frac{I_s^f}{I_0} = \frac{I_s^i c}{I_0 v_{g,s}^i}. \quad (3.12)$$

Figure 3.8a takes the same data from the middle panel of Fig. 3.5, but uses the dashed lines to show how the intensity of rays varies with their group velocity after they have been amplified. The spectrally integrated intensity (pink) illustrates that the reflectivity increases significantly with this consideration. The reflectivity using the initial intensity was 2%, but now it is 10%, in a situation where the pump was depleted by 17%. It is notable in this plot and confirmed by Fig. 3.8b that the scattered waves generated nearer $n_c/4$ (or with higher wavelengths) amplify significantly more as a result of the changing group velocity. At 700 nm the wave has increased by a factor 10 in intensity by the point it has exited the plasma. This is because the initial group velocity for these waves was lowest and so undergoes the largest change as the wave leaves the plasma.

Considering how energy is conserved and relates to reflectivity, this may seem slightly




Figure 3.8: Scattered light data from the same simulation as Fig. 3.5. (a) The colour scaled lines represent the intensity of individual rays, with the colour matching their wavelength upon exiting the plasma. The pink is the total intensity of scattered light leaving the plasma at each point in space. The solid lines show the intensity if we account for parametric amplification due to SRS. The dashed lines represent the increase in the intensity of waves due to the changing group velocity as they travel through new densities. (b) The Raman spectrum when considering the initial intensity I_s^i and the final intensity I_s^f , i.e. the intensities directly after Raman amplification and after Raman amplification and exiting the plasma respectively.

puzzling. The waves that are generated nearer the quarter critical surface are more heavily weighted in the reflectivity calculation. However, by considering the frequency matching condition and the frequencies of the scattered waves, closer to $n_c/4$ the scattered wave frequency decreases, so for every unit of pump energy lost, proportionally less of it is going into the scattered wave than the EPW. For each photon that is Raman scattered at $n_c/4$, where $\omega_s \approx \omega_0/2$, the energy of the scattered photon is 50% of the incident photon, whereas at a lower density where scattered frequency increases, the energy of the scattered photon is higher.

The reason for “inflating” these reflectivities from lower frequency waves is that when applying the conservation of action to deplete the pump, the intensity of the scattered wave has effectively been “deflated” due to its small group velocity. For example, considering Fig. 3.7 and Eq. (3.9), a unit increase in scattered wave intensity near $n_c/4$ led to a factor ~ 15 decrease in pump intensity. Substituting Eq. (3.9) into the expression for wave intensity

exiting the plasma(Eq. (3.11)):

$$\begin{aligned}
 I_s^f &= \frac{c}{v_{g,s}^i} I_s^i \\
 &= \frac{c\omega_s}{v_{g,0}\omega_0} \delta I_0 \\
 &\approx \frac{\omega_s}{\omega_0} \delta I_0.
 \end{aligned} \tag{3.13}$$

So after including the group velocity correction, the intensity ratio between the final scattered intensity and the intensity depleted by that instability is almost identical to the experimental reflectivity in terms of energy, which is based on frequency ratio. When considering the fraction of the absorbed pump energy going into the scattered wave it is clear from the range of possible frequencies that the reflectivity should always be more than 50% of the pump depletion, so in this example with 17% depletion, 10% reflectivity is more suitable than 2%.

3.1.5 Comparing Analytically and Numerically Integrated Convective Amplification

Section 2.2.3 showed a few different approaches to calculating convective amplification that were labelled the ‘‘Rosenbluth’’, ‘‘bounded’’ and ‘‘integrated’’ gain. The Rosenbluth and bounded gain are integrated over space analytically and so the model only has to carry out a single calculation to find the amplitude at which a scattered ray saturates. The integrated model instead integrates Eq.(2.71) numerically as the scattered ray propagates through space. This allows for the accounting of variation in plasma parameters and by using $\kappa(x)$ rather than $\kappa'x$ it more accurately captures the effect of dephasing. The integrated model is thus expected to be the most physically accurate, however it is slower than the other models as a result of calculating the evolving plasma parameters along the scattered rays. In this subsection it is compared with the Rosenbluth and bounded models to define their regimes of applicability.

Figure 3.9 compares the three methods looking at SRS reflectivity for a single pump ray over a range of density scale lengths and pump intensities with each simulation having been iterated until pump profile convergence. The simulation was limited to a density profile that had a maximum of $0.23n_c$, avoiding the densities where convective saturation fails, to make it a more robust comparison of the convective models. As expected, across all three methods, higher intensities and longer density scale lengths lead to higher reflectivities. In all three methods, at intensities below 10^{15} W/cm² and at the shortest scale lengths

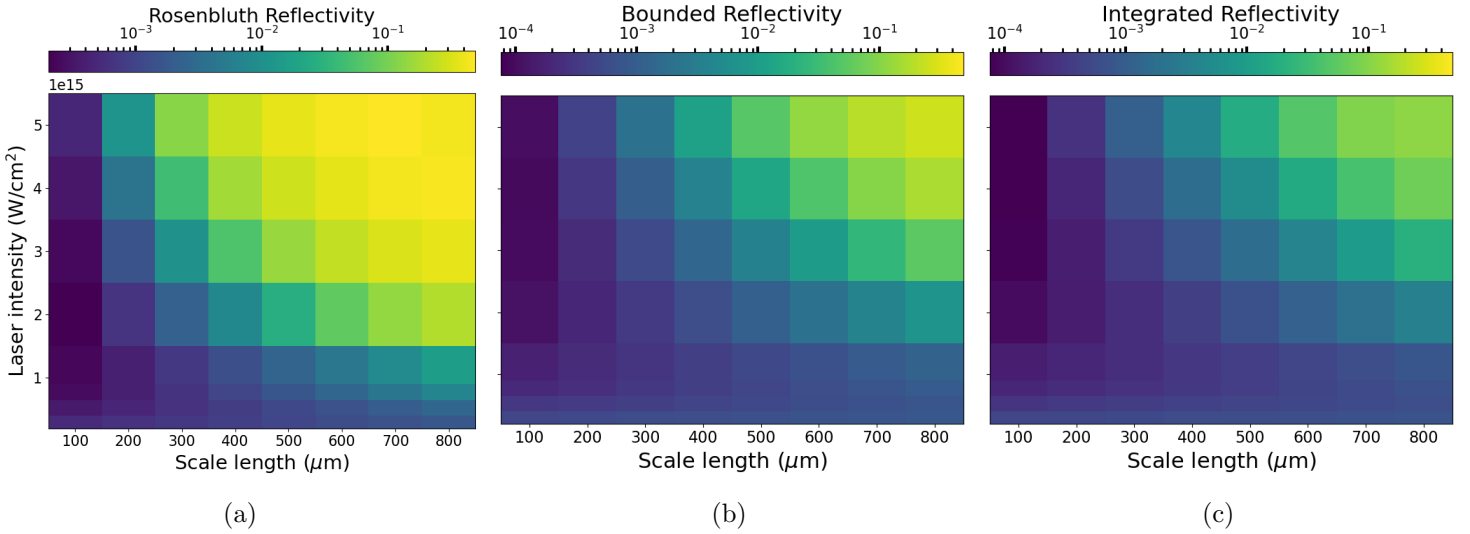


Figure 3.9: Reflectivity defined as $\frac{I_{sc}}{I_0 v_{g,s}^i}$ for a single pump ray with differing initial intensity and plasma scale lengths. Each simulation had $T_e = 2.5$ keV, a density upper bound of $0.23n_c$ and was run for 20 iterations or until convergence. They only modelled convective SRS amplified by (a) the Rosenbluth gain, (b) the bounded gain or (c) the integrated gain.

the reflectivity increases slightly at lower intensities. This is just a result of plotting the normalised reflectivity — in these cases a larger fraction of the pump is reflected, however the absolute value of reflected light is still smaller than the cases with a higher intensity pump.

All three plots have the same colour scale to make comparison easier. Comparing the Rosenbluth reflectivity with the integrated, towards the top right corner there is a difference as large as $0.4I_0$. The difference between them drops below $0.1I_0$ at $400\mu\text{m}$ for a 3×10^{15} W/cm² pump. For a pump intensity of 1×10^{15} W/cm² the difference between them is always below $0.02I_0$, which suggests the Rosenbluth gain is a good approximation in this regime. Part of the reason for that is the absolute values of reflectivity are so low that the differences between them are naturally small, while the fractional difference is still significant. However, once the reflectivity has dropped below $0.01I_0$ it means there is little SRS occurring and so the relative differences between the models is unimportant. The bounded gain performs much better in comparison to the integrated model. Apart from the results in the three uppermost right-hand squares, they are never more than $0.08I_0$ apart.

The simulations from the same three upper-right squares did not converge when using the Rosenbluth gain as a result of extremely large instability growth. The non-converging pump problem and solutions to it are detailed in Sec. 3.3.2. For simulations further towards

the top right corner, those using the bounded gain converge significantly faster than those using the integrated gain, despite the bounded gain simulations having higher reflectivities. Some simulations not pictured here with higher reflectivities found that the bounded method does converge whilst the integrated method fails to converge. The fact the bounded method converges in regions of high gain where the integrated method does not is a positive for the method, implying better precision in these regimes. When combined with the fact its runtime per iteration is faster, the fact that the bounded method converges in many fewer iterations than the integrated method means that it is significantly faster.

Iterating until convergence may not providing the most robust comparison of the different amplification methods. For a steady-state solution, the amount of SRS occurring further along a pump ray is partly dependent on the amount of SRS before it as this depletes the pump and reduces the growth rate. In theory, after iterating two different scenarios to convergence and finding they have a similar reflectivity, examining spectra could show significant differences, where one undergoes much more SRS at lower densities, and no longer has enough energy to produce much SRS at high densities, whereas the other example doesn't undergo as much low density SRS but does undergo high density SRS.

Figure 3.10a shows the spectra produced by the three methods for the simulation with $5 \times 10^{15} \text{ W/cm}^2$ and $600 \mu\text{m}$ as this was one of the highest gain examples that converged for all three methods. The reflectivities of the three methods are $R_{ros} = 0.54$, $R_{bnd} = 0.11$, $R_{int} = 0.07$. Despite the reflectivities being quite different this is still a useful example to show how the shape of spectra and location of pump depletion can vary. The Rosenbluth spectrum sees high gain initially for shorter wavelength SRS and then the reflectivity decreases for higher wavelengths (higher densities) as a result of pump depletion, which is depicted clearly in Fig. 3.10b. Comparing the reflectivity from wavelengths over 650 nm from this simulation, it appears that the bounded gain results in higher SRS levels than the Rosenbluth gain, when really it is just because the pump intensity in this part of the plasma is significantly higher for the bounded simulation. By converging the simulations to include pump depletion the different methods are not being compared with identical parameters, because the pump intensity differs locally in density and temperature space.

One way to try to compare the differences between the methods across the whole plasma density range is to ignore pump depletion so that a constant intensity pump drives all of the SRS. This will of course give non-physical results when looking at reflectivity as the energy

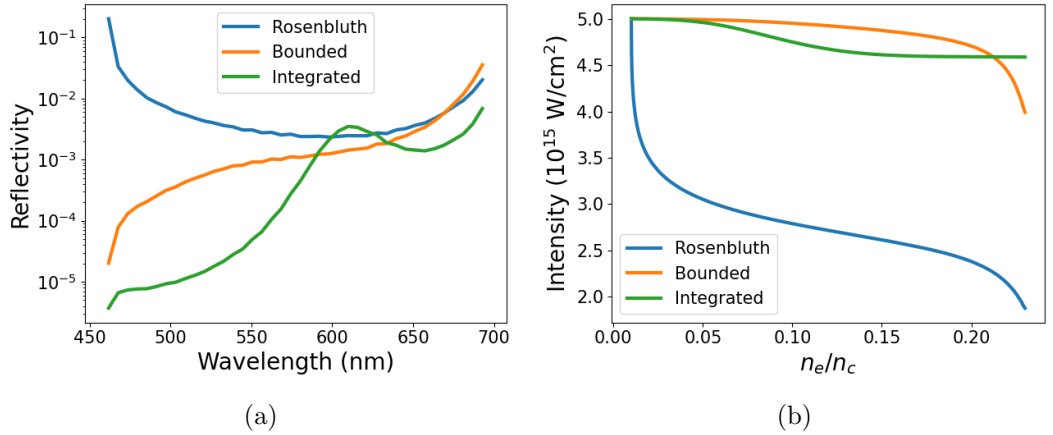


Figure 3.10: (a) Spectra from the simulation in Fig.3.9 with $I_0 = 5 \times 10^{15}$ W/cm² and $L_n = 600 \mu\text{m}$. (b) Pump profiles from the same simulation.

of the pump is no longer conserved, but it is a useful approach to compare the models with the same pump intensity across the whole spectrum.

Figure 3.11 plots reflectivity for simulations that ignore pump depletion, only running the LPI calculations for a single iteration and making no corrections to reflectivity as a result of pump evolution. The Rosenbluth reflectivity reaches values above I_0 at low intensities and scale lengths because its reflectivity is not limited by pump depletion. This method is now many orders of magnitude off both the bounded and integrated reflectivities, indicating that it results in much much higher gain across the plasma, which was not captured in the converged results. The bounded and integrated reflectivity are plotted with the same colour scale. They are more easily compared in Fig. 3.12a, where the difference between them has been plotted in each square. The red line shows the boundary where the difference is more than $0.1I_0$. Given that both models are derived in the fluid limit their regime of applicability is already limited to $\lesssim 3 \times 10^{15}$ W/cm². Thus, in the regimes where the theory used to derive the models is applicable, they give similar results, and the bounded method can be assumed to be a good approximation to the more complete integrated method and extrapolated to be used in larger scale, three-dimensional simulations with many thousands of rays.

Figure 3.12b again shows the spectrum for 5×10^{15} W/cm² and $600 \mu\text{m}$, but this time after a single iteration (c.f. Fig. 3.10a). As expected, the spectrum from the Rosenbluth gain does show an increase in intensity at higher densities compared to Fig. 3.10a as pump depletion is no longer affecting it.

Across the wavelength range the bounded and integrated spectra are similar apart from

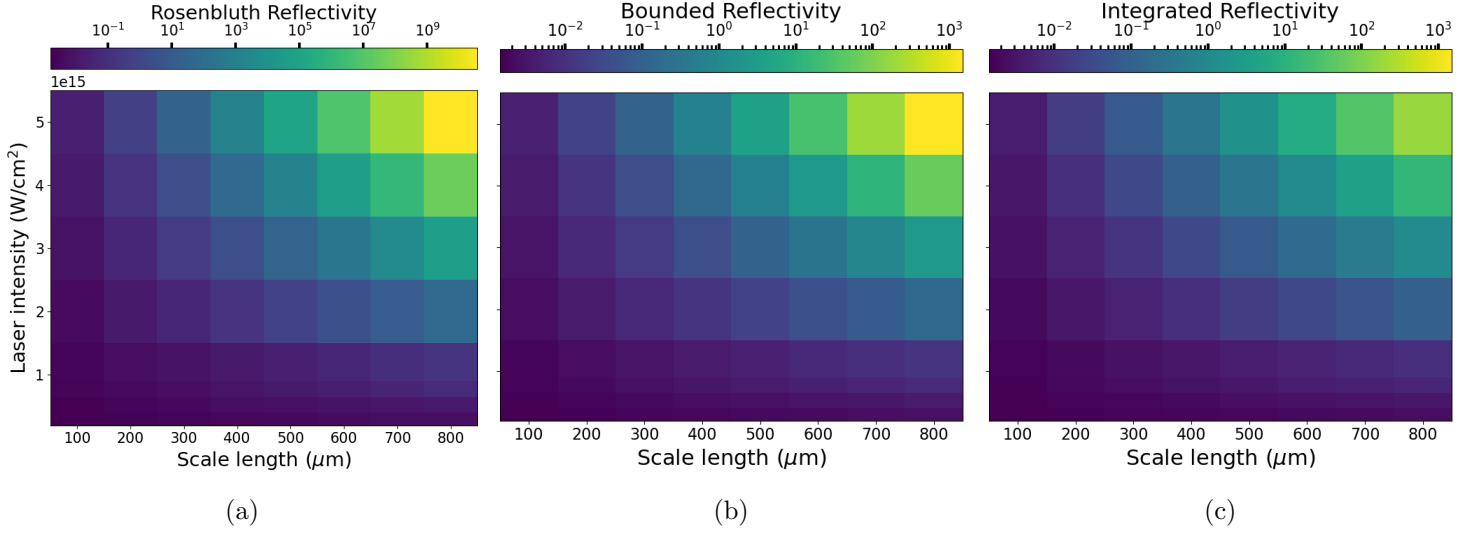


Figure 3.11: Reflectivities for the 3 gain methods along a single pump ray of differing intensity in plasmas with differing scale lengths and constant temperature of 2.5 keV. The only difference between these simulations and Fig. 3.9 is that those simulations had the LPI calculation iterated over up to 20 times to include the effects of pump depletion, whereas these completely ignored pump depletion by using a constant intensity pump throughout the plasma and not iterating.

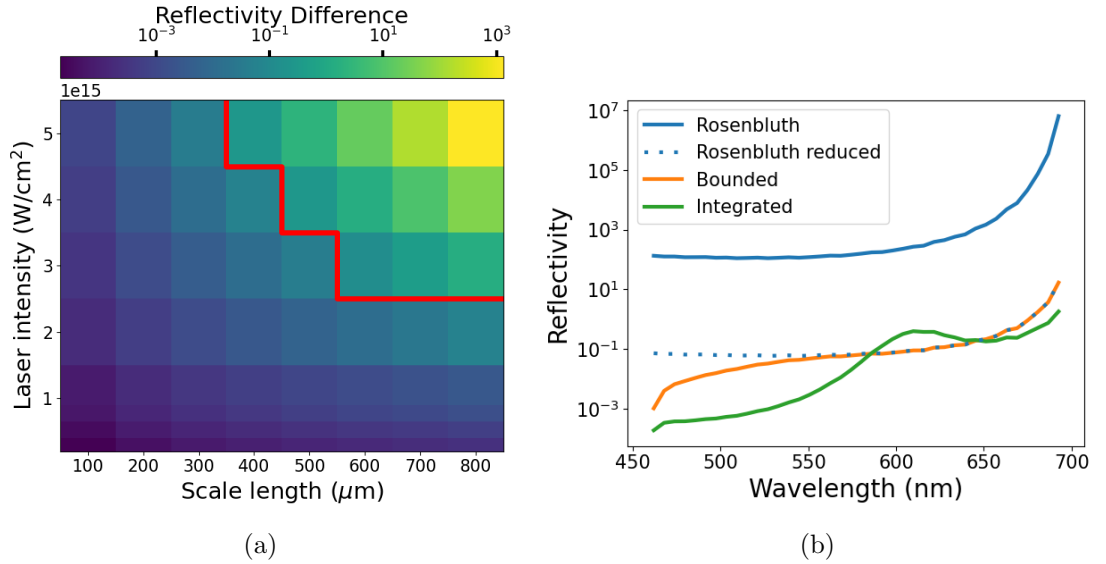


Figure 3.12: (a) The absolute difference in bounded and integrated reflectivity from the simulations that ignore pump depletion in Fig. 3.11. The red line forms the boundary between simulations that had a difference of more or less than $0.1I_0$. (b) The spectra from the simulation with pump intensity $5 \times 10^{15} \text{ W/cm}^2$ and density scale length $600 \mu\text{m}$. The dotted blue line changes the integration limits of the Rosenbluth gain and is explained in the text.

a small hump in the amplification of rays at around 600 nm in the integrated spectrum. The dip this creates at ~ 650 nm is perhaps an additional factor contributing to the Raman gap (see Sec. 2.2.4). The Rosenbluth gain is many magnitudes larger than the other two, but it does follow a similar shape to the bounded gain as a result of them both being derived the same way, the only difference being their integration limits. The bounded gain tails off at low densities in the same way the integrated gain does, which is an important feature as experimental signatures at wavelengths below 500 nm are rarely seen. This feature is not present in the Rosenbluth spectra because, while the bounded gain integrates over the length of the scattered ray, the Rosenbluth gain is integrated forwards all the way to $+\infty$.

Besides the tailing off at low wavelengths for the shorter length scattered rays, there seems to be a fairly constant factor between the Rosenbluth and bounded spectra in this plot. This is because the Rosenbluth gain starts its integration at $-\infty$, whereas the bounded gain started it at 0, where 0 is the location of the resonance. The spectrum labelled Rosenbluth reduced on the graph is created by instead integrating from 0 to $+\infty$. It is identical to the bounded gain at longer wavelengths where the scattered ray lengths are many times their amplification lengths (see L_d in Eq. (2.68)), and differs for shorter rays where integration limits matter.

The integration in Chiron was chosen to be started at 0 rather than $-\infty$ as this is the location that a seed is amplified. It isn't clear why the Rosenbluth gain is integrated from $-\infty$ and also figures in Michel's [12] and Strozzi's [104] works also depict rays starting at $-\infty$ and amplifying a small amount before they reach their resonances. Ultimately, for this work, which currently lacks self-consistent, physics-based models for LPI seeds it is somewhat arbitrary. When a seed amplitude is considered in this work, it is envisioned as the amplitude of the wave at the point of resonance so the integration from 0 to δx fits. If instead the seed model was representative of rays that start amplifying many amplification lengths before the resonance then that seed should be integrated from before 0.

3.1.6 Nonlinear Saturation by the Langmuir Decay Instability

Throughout the convective saturation comparisons between the three models in the previous section the plasma density was always capped at $0.23n_c$ to remove the part of the plasma where convective dephasing often fails to saturate SRS and thus compare the models only in regimes they are applicable. Near the quarter critical surface, when convective saturation

fails, SRS and TPD are saturated by other mechanisms, with the most prevalent (and the only one included in Chiron) being the nonlinear decay of an EPW via the Langmuir decay instability as described in Sec.2.2.7. Equation (2.97) describes the amplitude an EPW saturates at due to LDI. Notably, it has no seed dependence, so, unlike with the convective model, the seed amplitude cannot be used as a way to include scattered ray bandwidth and ensure results are invariant to changes in the number of scattered rays that are launched. Instead, we add a term to the LDI saturated EPW amplitude:

$$E_{epw} = \frac{2m_e v_e \omega_{epw}}{e} \frac{1}{\sqrt{\omega_{epw}^{ldi} \omega_{iaw}^{ldi}}} \gamma_{srs} \frac{\delta\omega_s}{\Delta\omega_s}, \quad (3.14)$$

where $\delta\omega_s$ describes the bandwidth of the specific scattered ray and $\Delta\omega_s$ is the total bandwidth of SRS rays along the pump ray. This expression is now invariant to changes in the number of scattered rays.

Multiple methods were attempted to determine when Langmuir decay saturation came into effect. Some were based on the condition for an absolute instability to occur, such as using the threshold in Eq.(2.56) plus an expression by Michel [12] for when absolute growth can occur or an expression from Drake *et al.* [75] for the width of the absolute growth region




Figure 3.13: Spectra from two simulations with (a) $I_0 = 2 \times 10^{15} \text{ W/cm}^2$ and $L_n = 500 \mu\text{m}$, and (b) $I_0 = 4 \times 10^{15} \text{ W/cm}^2$ and $L_n = 400 \mu\text{m}$. Both have switched off the effects of pump depletion for clarity. Each simulation shows the reflectivity for rays based on whether the saturation is convective or LDI, with the solid lines depicting the dominant saturation mechanism at each wavelength. The reflectivity is defined here as I_s^i/I_0 using a scattered ray's initial intensity after Raman amplification rather than final intensity after exiting the plasma because these are the values the model actually compares when selecting the dominant saturation mechanism.

in space. The method settled on, when used in tandem with the bound (or Rosenbluth) convective method, is to evaluate the amplitude each scattered ray would saturate at if convectively damped or if saturated by LDI. The saturation mechanism that results in a smaller amplitude wave is chosen as the dominant mechanism, as logically this process will prevent any further growth. This technique is demonstrated in Fig. 3.13a where we plot spectra for rays that have been saturated either convectively or by LDI. As expected, at higher wavelengths (nearer $n_c/4$) convective saturation becomes less effective. In this region, saturation via LDI becomes the dominant mechanism, limiting the growth of the scattered ray more strongly than convection.

Figures 3.14a and 3.14b show the difference adding LDI makes to the Raman model using the bounded convective gain. These simulations had no density cap (like the previous section) to examine the role of LDI at the quarter critical surface. These simulations were also run until convergence rather than just a single iteration, thus including the effect of pump depletion. The aim is to compare the final results for reflectivity the model gives with and without LDI rather than giving equal importance to amplification across the spectrum. The simulations without LDI reach reflectivities approaching 60% with the pump depleted by almost 90% in the highest gain examples, which is approaching the maximum for a simulation which includes the effects of pump depletion. This is expected for simulations which have no way to cap Raman amplification at the quarter critical surface, because convection fails to saturate the rays. Travelling from left to right and bottom to top on this graph reflectivity reaches $>30\%$ quickly with significant jumps between squares. Beyond this the rate of increase in reflectivity slows down, however this is just the nature of conserving energy by including the effects of pump depletion in the simulation — there is less and less pump energy left to be scattered away.

Figure 3.14b shows that including LDI instead caps the Raman reflectivities at 6% and the pump depletion at 9%. A notable feature here is reflectivity becoming constant at the largest intensities and scale lengths. This reflectivity is normalised to the pump intensity so this means that the scattered intensity does still increase with pump intensity, but that it now increases linearly. This is illustrated by Fig. 3.14c which shows the total scattered intensity exiting the plasma without the pump normalisation. Here the increases in scattered intensity with pump intensity are clear, however the simulations with intensities and scale lengths that resulted in constant reflectivity in Fig. 3.14b show constant scattered intensity

with respect to increasing scale length. This is not a feature that would likely be seen with the convective model, because convective amplification depends on the gradients in space and the amount of space that a ray amplifies over. The LDI model does not have any dependence on gradients or amplification length and so this behaviour emerges. The behaviour only emerges from this model at high intensities and scale lengths because this is where the LDI becomes the dominant instability over the majority of the pump ray due to convective gains becoming much larger.

Figure 3.13b shows an example spectrum for a simulation where LDI saturates all of the scattered rays launched from $n_e > 0.07n_c$. The increased pump intensity from Fig. 3.13a has led to a significant increase in the level that the convective model saturates the instability at, so LDI saturation dominates at lower densities. In these cases where LDI saturates the majority of the rays, the expected scattered intensity will be the same regardless of scale length. These are also the simulations that see constant reflectivity regardless of pump intensity, which again comes about as a result of LDI saturating most of the scattered rays. The amplitude that LDI saturates the daughter waves increases linearly with the pump intensity. Convective saturation grows exponentially with pump intensity, thus for any simulation where LDI is not dominant across the majority of scattered rays, i.e. all of Fig. 3.14a and the lower left parts of Fig. 3.14b, reflectivity increases with pump intensity.

Rosenberg *et al.* [73] found that the levels of absolute SRS scale linearly with laser intensity following a thorough investigation of SRS on the NIF. The nature of Eq.(3.14) which defines LDI saturation also means that this SRS saturation mechanism will scale linearly with intensity. In reality however, it is likely that there will be other, potentially more dominant, mechanisms that saturate absolute SRS such as density profile modification and mode coupling as discussed in Sec. 2.2.6.1.

Throughout this section, the new saturation mechanism has been termed as being via LDI rather than absolute, because, as is shown in Fig. 3.13b, it can saturate LPI far away from the region where growth would be considered absolute. While being the dominant saturation mechanism of absolute SRS and TPD, Langmuir decay can still happen at lower densities, it is just less effective. Ideally, LDI would be included in the convective model as an additional damping rate, however this has proven difficult due to the nonlinear nature of LDI growth. Specifically that it depends on the amplitude of the EPW that is growing due to SRS and depleting due to LDI.

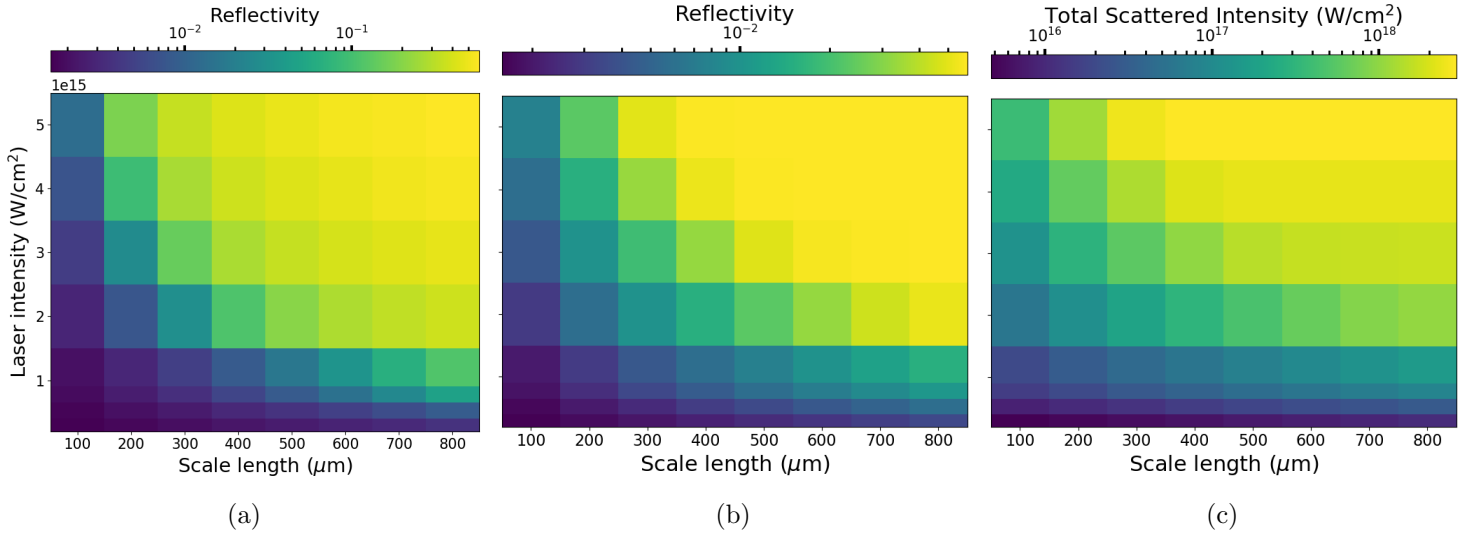


Figure 3.14: Raman scattered light data from simulations to investigate the effect of our LDI model on the SRS model utilising the bounded gain for convective saturation. Each simulation was for a single pump ray in a 2.5 keV plasma with a maximum density of $0.25 n_c$ and includes the effects of pump depletion. (a) Allowed only convective saturation using the bounded gain while (b) used both this and saturation via the Langmuir decay instability. (c) Shows results from the same simulations as in (b) but this time plotting the total intensity of scattered light exiting the plasma without normalising to pump intensity.

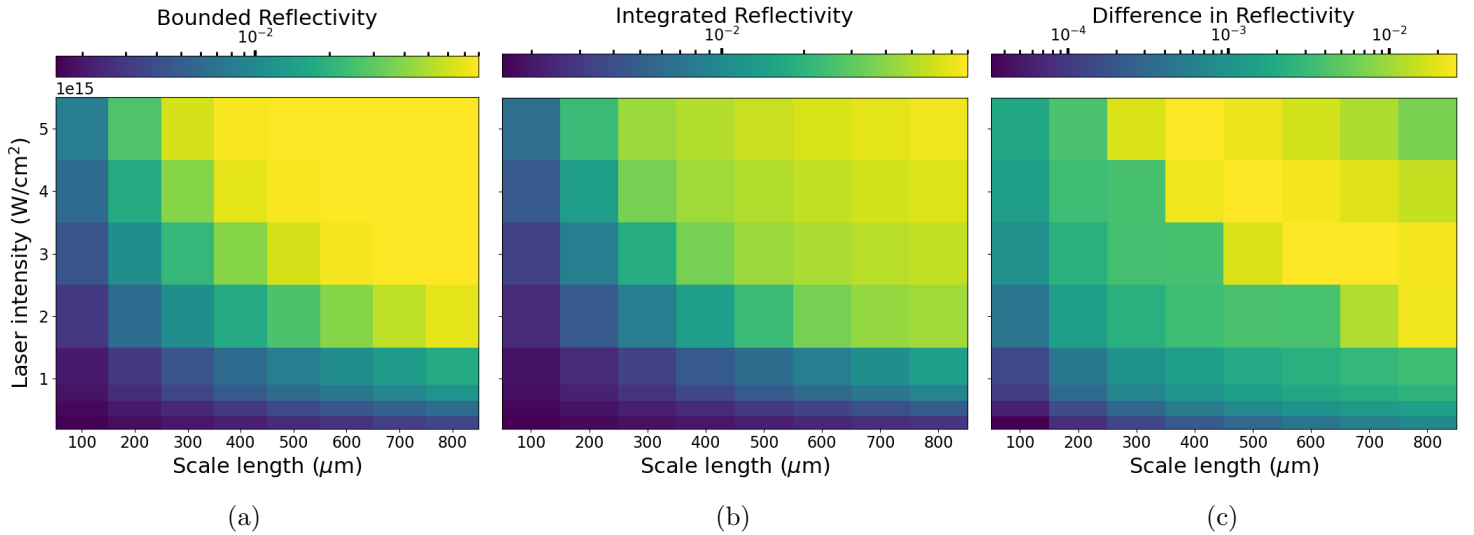


Figure 3.15: Raman reflectivities from simulations using the (a) bounded and (b) integrated gain methods for convective gain combined with Langmuir decay saturation as well as (c) the difference between the two sets of data. The simulations ran for 1 iteration only, ignoring the effect of pump depletion.

Despite a desire to compare the integrated and bounded gain models with the inclusion of LDI, a unique implementation of the LDI model for the integrated method was not found. One attempted approach was to implement LDI as an additional EPW damping rate based on an approximation of EPW amplitude. Ultimately, the comparison between the two models in Fig. 3.15 used an identical LDI implementation between them, so for the integrated model we again saturate a scattered ray at whichever method results in a lower amplitude, either the integrated gain or using the same LDI calculation in Eq. (3.14). This is somewhat unsatisfactory as when LDI is dominant the rays are no longer amplified in space at all, thus not considering the variation of parameters along the scattered ray. It also weakens the comparison because both methods have an identical LDI model. The comparison was done without iterating for pump convergence, instead just running the simulation once, making it more comparable to the results without LDI in the previous section.

Figure 3.15a shows the bounded reflectivity and is identical to Fig. 3.14b, only placed here again for easier side-by-side comparison. Figure 3.15b is plotted with the same colour scale that shows that it still has lower reflectivities than the bounded method. Given that both methods return identical results for LDI saturation of rays, the reason the integrated method gives lower reflectivities now is that its levels of convective saturation are lower than the bounded methods so these dominate over LDI for more scattered rays. The difference between the two sets of simulations in Fig. 3.15c is always below $0.03I_0$, which is small in comparison to the values $> I_0$ in Fig. 3.12a. When moving in the upper right direction on the graph the difference increases, which is similar to previous plots comparing the two methods. However, in the top right corner it decreases again, as a result of LDI becoming dominant in simulations using either method and thus bringing the results closer together.

An issue with the convective modelling of SRS presented in Sec. 3.1.2 was the large number of scattered rays needed to accurately sample the profiles along the pump. This problem is predominantly caused by any region where the convective gain spikes and is most often seen at the quarter critical surface. By implementing the LDI model there is now a nonlinear saturation mechanism that prevents this runaway convective gain and should significantly decrease the number of scattered rays needed to provide accurate results. Figure 3.16a plots the results of the same simulations seen in Fig. 3.3 except this time the LDI model is included. The difference is stark. Convergence to within 10% of the converged reflectivity now happens within 50 rays and within 100 rays the results are almost identical to those

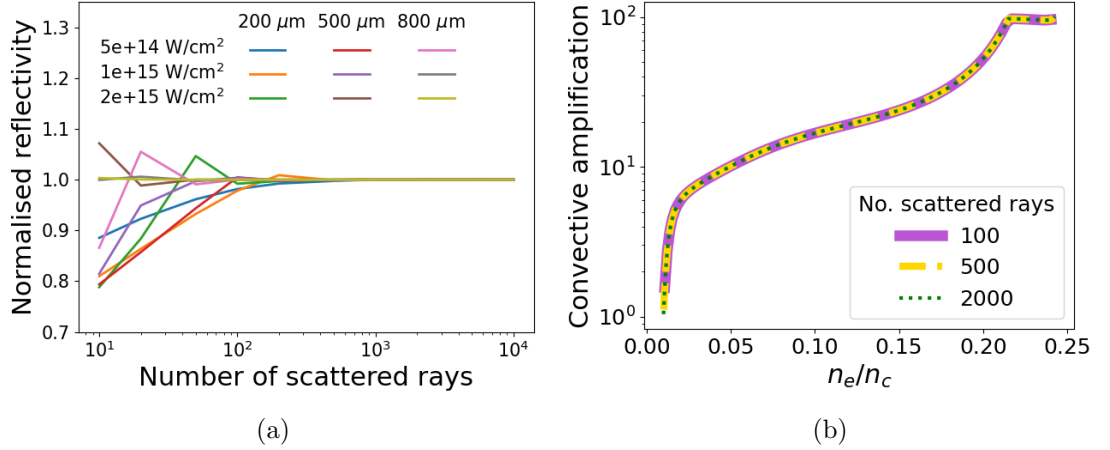


Figure 3.16: (a) Reflectivity variation with number of scattered rays launched for 9 different combinations of pump intensity and density scale length. Each pair of parameters' reflectivity is normalised to the pair's reflectivity when 10,000 rays are launched. The simulations presented here are identical to those in Fig. 3.3 except for the inclusion of the Langmuir decay instability. (b) The convective amplification for rays seeded at each density for simulations with pump intensity $2 \times 10^{15} \text{ W/cm}^2$, scale length $500 \mu\text{m}$ and varying numbers of Raman scattered rays. These simulations are identical to those presented in Fig. 3.4 except for the inclusion of LDI.

at 10,000 rays. There is also no longer a trend for certain parameter pairs of intensity and scale length (those that caused higher gain) requiring more rays to converge as there was previously, meaning the limits on ray number that are derived from this graph are likely applicable to all intensities and scale lengths. Figure 3.16b similarly repeats the simulation carried out to create Fig. 3.4, but this time with the inclusion of an LDI model. Previously, the larger number of rays allowed higher sampling resolution and thus higher gain at the upper density limit of the pump ray. The new results show the restriction on amplification provided by LDI saturation. Increasing ray resolution here will not lead to an increase in reflectivity, hence the data in Fig. 3.16a converges at much smaller numbers of scattered rays and without any dependence on the magnitude of reflectivity in the simulation.

3.1.7 Multiple Absorption mechanisms

Moving on from the examples of just SRS occurring along a ray, this section discusses how the model works with all three LPI and IB occurring along one pump ray. Each instability is modelled in the same way: a number of instability waves are generated that amplify according to the appropriate saturation physics and then the pump is depleted by each occurrence of an instability. Each iteration of our LPI calculation now proceeds like:

1. Use the same pump profile found at the end of the previous iteration to calculate the growth rates of all three instabilities
2. Calculate the saturated levels of SRS
3. Calculate the saturated levels of SBS
4. Calculate the saturated levels of TPD
5. Deplete the pump by subtracting the intensity lost to SRS, SBS and TPD from the initial (iteration 0) pump profile
6. Compare to the pump profile at the end of previous iteration and if not converged return to step 1, using this pump profile to determine LPI growth rates.

The choice to use one pump profile per iteration to calculate the growth rates of all of the instabilities prevents this model from favouring one instability over the others. For example, if the SRS calculation happened first, depleted the pump, and then used that depleted pump for the SBS calculation the model would always have SRS seeing a larger pump intensity than SBS. Physically, this would imply that the instability that is considered first has effectively saturated before the other instability has even started growing. By using the same pump for each instability no preference is given to any instability such that the order of steps 2, 3 and 4 are arbitrary. Figure 3.17a shows the results of simulations for all 6 possible permutations of the order of steps 2, 3 and 4. As anticipated, the order makes no difference as the fraction absorbed by each instability is constant between each scenario. The high levels of SBS in this simulation and the fact that 40% of the pump is not absorbed are a result of no inverse bremsstrahlung being included.

Another important feature to note from the LPI calculation is that once the amount each instability depletes the pump by is found, this is subtracted from the original pump profile that the simulation started with, rather than the pump profile from the last iteration. Doing it this way we are looking for the steady-state solution: the resulting pump profile that balances with the amount of LPI that can grow.

Modelling inverse bremsstrahlung is functionally different from LPI. Whereas the instabilities absorb some amount of the pump depending on the pump profile, linear IB absorbs a fraction of the pump power independent of what that power is. Combining these models raises some questions about the order of LPI and IB and which pump profile (depleted by

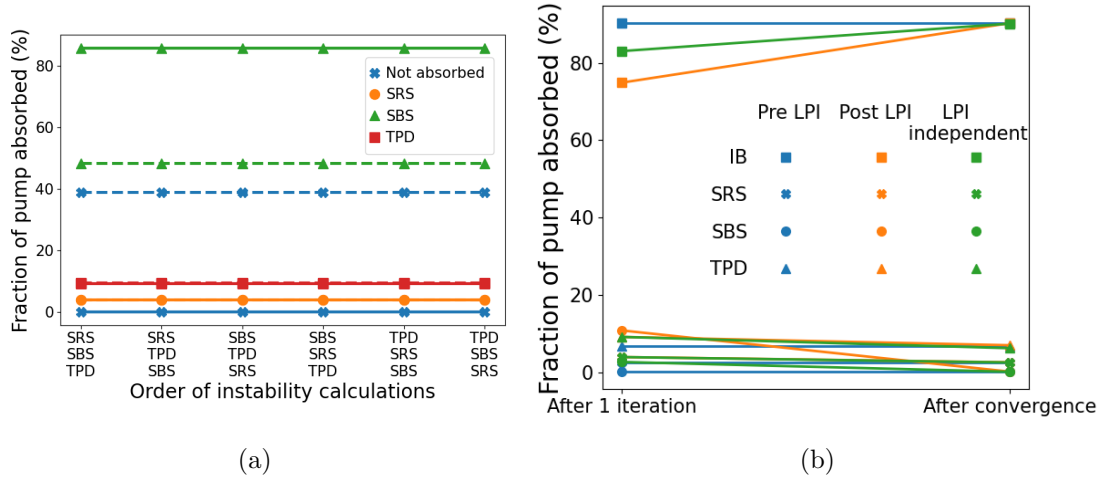


Figure 3.17: (a) Fraction of the pump absorbed by each instability in simulations with inverse bremsstrahlung turned off. The solid lines are from a simulation that ignored pump depletion (lasted one iteration) while the dashed line is a simulation that included pump depletion (iterated until pump profile convergence). (b) The fraction of the pump absorbed by each absorption mechanism for three different implementations of the IB model after 1 iteration of LPI calculations and after converging iterations. For example, the blue markers show the fraction of the pump absorbed by the pump when the pre-LPI method was implemented. All simulations had $I_0 = 2 \times 10^{15} \text{ W/cm}^2$, $L_n = 500 \mu\text{m}$ and $T_e = 2.5 \text{ keV}$ and modelled a ray propagating in one direction through a plasma extending up to its critical density.

LPI or not) the IB calculation should be applied to. Three approaches were developed for how to implement the IB model in tandem with the LPI models that are here called pre-LPI, post-LPI and LPI independent. The pre-LPI model calculates IB first during each iteration and then uses the IB depleted pump to calculate the LPI growth rates. The post-LPI model applies IB absorption to the pump profile after it has already been depleted by the LPI. The LPI independent model functions the same way as the LPI models in that IB is modelled using the pump profile at the start of an iteration while the LPI use the same profile for their growth with depletion by all mechanisms coming after this.

Figure.3.17b shows the fraction of the pump absorbed by each absorption mechanism for the three different IB implementations. After 1 iteration the results are not consistent between implementations as the the LPI and IB see different pump profiles based on when they are calculated. However, after iterating until convergence, all three implementations result in the same absorption fraction for each of the four mechanisms, showing how robust this method is and that it ultimately doesn't matter which IB implementation we use. The model currently uses the post-LPI method for IB absorption.

3.1.8 Scattered Light Diagnostics

Chiron has two key outputs for diagnosing simulations. The first is the ability to see where in the plasma a laser ray was absorbed and by which process, allowing the creation of data that show the spatial deposition of laser energy by IB and the different LPI processes. This is useful for understanding what is going on in the simulation, and can be used in comparison with other simulation work that also provides this insight. The spatial deposition of energy by absorption process is not something that can be measured in an experiment, so this diagnostic can not be used to benchmark the code against experiments.

The second diagnostic, which can be used to compare Chiron simulations to experiments (or other codes), is measuring the scattered light generated by each instability. For SRS and SBS scattered light is directly amplified by the instability, so the amplitude is calculated automatically. Chiron does not currently trace the paths of the scattered rays, instead assuming they propagate directly back along the path of the pump ray. This is a particularly bad approximation for SRS, where the large difference in frequency between pump and scattered ray will lead to them refracting differently through the plasma. The diagnostic is currently in its infancy and tracing the scattered rays is its next step. With this, synthetic diagnostics will be able to be created to more accurately compare with experiments. The scattered light also undergoes collisional absorption as it exits the plasma.

The scattered light that results from TPD is from multiple secondary processes as discussed in Sec. 2.2.6.1. Chiron calculates a value for the amplitude of the two EPWs that are generated by TPD, but it is then difficult to predict the expected scattered light spectrum that they generate. The results presented in Sec. 4.3.2, instead assume that the relation between EPW amplitude and scattered light amplitude is linear, and so sets the amplitude of the scattered light equal to the amplitude of the EPW that generated it. Then it applies collisional absorption to that scattered light amplitude to look for the relative decrease in amplitude between the two lobes.

This was first attempted with TPD to try to just compare the relative amplitudes of the two scattered light lobes, with no attempt to make conclusions based on the absolute levels of scattered light, because of the various inaccuracies and approximations involved. However, given the fact that different mechanisms are responsible for generating the two lobes of light, even the assumption that their amplitudes would scale similarly with the EPW amplitudes may be incorrect.

Two-plasmon decay scattered light is emitted at a range of directions due to the EPW orientations, so being able to trace these rays, whose directions are known, would be useful to work out when TPD scattered light is expected to be measurable by a detector at a fixed location in a chamber.

3.2 Moving to Three Dimensions

3.2.1 Ray Tracing and Radiation-Hydrodynamics Codes

The LPI modelling software currently takes outputs from the one dimensional radiation hydrodynamics code HYADES and its two dimensional version H2D [105]. Future plans for this project are for the code to run as a part of the radiation-hydrodynamics code Odin. Radiation-hydrodynamics describes the study of the dynamical response of a fluid to interaction with electromagnetic radiation through absorption, scattering and emission. Radiation-hydrodynamics codes are therefore ideal candidates for simulating ICF experiments as they can model the plasma response to the driving lasers. As the beams propagate through the target, the Navier-Stokes equations describing the conservation of mass, energy and momentum are solved for each fluid element to describe its temporal evolution [106].

The two most fundamental methods for solving the conservation equations are the Eulerian and Lagrangian formulations [107]. In the Eulerian method, the volume of each fluid element remains constant while the total mass of fluid in each element varies. This effectively means that in a simulation there is a fixed mesh that the plasma flows through. The Lagrangian method, which HYADES uses, instead keeps the mass of each cell constant allowing their volume to vary instead. Each cell of the mesh is effectively locked to some amount of fluid throughout the simulation, changing shape and position with it. Both methods have advantages and drawbacks, Eulerian methods overestimate diffusion of matter as a result of treating the matter in a cell as being uniformly distributed. Lagrangian methods struggle with the mesh tangling with itself for complicated plasma flow, particularly rotational flows. The merits of both are discussed in more detail in Colvin and Larsen [108], and Castor [109].

Chiron only requires the details of the mesh and the rays that have been traced through it after a timestep so will work with both Lagrangian and Eulerian frameworks. This will allow adaptations of the model to be used across a variety of codes in the future including those that use more complex methods like Odin, which is an arbitrary Lagrangian-Eulerian

code [110].

In a HYADES or H2D simulation, a laser beam is made up of many rays distributed in space over the area of the laser focal spot and each with their own power to fit the beam’s spatial power distribution. Each ray is then traced in two or three dimensions and its energy deposited according to the inverse bremsstrahlung package with no way to model LPI physics along the rays. The ray tracing in hydrodynamics codes will typically use very large numbers of rays for an ICF simulation, often in excess of one million rays. Chiron applies LPI calculations to each individual ray so it is preferable to work with significantly fewer rays for computational speed, as I believe the same LPI physics can be captured with reduced ray numbers. Therefore, we choose to use a ray tracer written in Python by R. Scott, through which the number of rays can be redefined for Chiron to use. The ray tracer takes the hydrodynamic mesh output by HYADES after a timestep and launches beams with significantly fewer rays through it. The requisite number of rays is discussed in Sec. 3.2.3.

3.2.2 3D Mesh, multiple rays

After tracing the laser rays through the plasma mesh, there are two types of parameters: plasma parameters like electron and ion density, temperature and flow velocity, which are defined on the three-dimensional hydrodynamic mesh; and the ray parameters, like power, position and length that are defined for each step a ray takes through the mesh. Most important of these ray parameters are the indices of each cell of the mesh that the ray is in after each step. The aim was to develop the previously described models for single rays (Sec. 3.1) to be used along each ray of the 3D simulations with multiple rays. These models required the parameters along the path of a ray to determine LPI levels generated by that ray, so each fluid parameter is mapped onto each ray. For each ray, the indices of the cell it is in after each step is used to find the density and temperature that it sees in that step, generating three-dimensional arrays for each parameter along rays. There is one dimension for each laser beam, one for each ray within a beam and then one for the parameter that that ray sees in each cell that it steps through.

The ray tracer models the propagation of infinitesimally thin, so called “eikonal”, rays which have a power but not the intensity that is needed to find LPI growth rates. Therefore, intensity of a ray in a cell as

$$I = \frac{Pdx}{V}, \quad (3.15)$$

where P is the ray's power, dx the path length of the ray in the cell and V is the volume of the cell. The length of the ray in the cell divided by the volume of said cell effectively give the cross sectional area that the ray sees in the cell.

This approach is very prone to noise and requires large numbers of pump rays per hydrodynamic cell to reduce the noise created by fluctuations in dx and V . An alternative method, used in the CBET post-processor BEAM-CROSSER [111], is to use neighbouring rays. For every pump ray that is initialised, three neighbour rays (or two for a two-dimensional ray trace) are initialised surrounding it in an equilateral triangle perpendicular to the initial direction of propagation. These three rays are also traced up to the phase of the main ray and the area of the triangle that they form is used to find the intensity of the main ray. This method would allow significantly fewer pump rays to be launched per simulation as it is less prone to noise created by the hydrodynamic mesh. Recent discussions with researchers at Imperial College London, where they have changed from calculating intensity via cell volumes and ray path lengths, as in Eq. (3.15), to the neighbouring ray method for their CBET model has further confirmed that this is an approach that will be pursued by Chiron in the future.

Next the model had to be made invariant to the number of pump rays that we launched. This problem is illustrated by considering a simulation where there is a single ray in a cell with a power P_0 , or the same simulation but with higher ray resolution such that the single ray has now been broken up into ten rays in the same cell each with power $P_0/10$. In the single ray case the intensity is I_0 so the convective gain $G \propto \gamma_0^2 \propto v_{osc}^2 \propto I_0$ and thus $I_s^{(1)} \propto \exp(I_0)$. The superscript in parentheses refers to the number of pump rays in the example. For the ten ray case, the scattered intensity from each ray is $I_s^{(10)} \propto \exp(I_0/10)$ and then the total scattered light generated by those ten rays is $I_{s,tot}^{(10)} \propto 10 \exp(I_0/10)$. Clearly, the two simulations give different results as $\exp(I_0) \neq 10 \exp(I_0/10)$. Another simple way to think about this problem is if the intensity I_0 is just over the damping threshold to drive the instability, the first simulation would see LPI growth, however in the version where the intensity is split by 10 rays, they will all be under this threshold and no LPI will grow.

The approach taken by Chiron is instead to drive any LPI within one cell by the summed intensity of all of the rays in that cell, while still launching LPI waves off each individual pump ray. By using an intensity per cell the value is now invariant to ray resolution, however by launching scattered rays off each pump there is a possibility of significantly overpredicting

the LPI levels. In the one versus ten rays in a cell example the intensity in that cell is I_0 in both cases. However, if each pump ray launches one scattered ray in this cell, the single ray scenario still launches a scattered ray with $I_s^{(1)} \propto \exp(I_0)$, but each of the ten rays in the other scenario launches a scattered ray of the same intensity, so in total launches $I_{s,tot}^{(10)} \propto 10 \exp(I_0)$. The model multiplies the intensity of each scattered ray by the fraction of intensity that that pump ray contributed to the total cell based intensity which is driving the instability growth. For example, the intensity of each of the ten scattered rays that are launched off each pump ray is $I_s^{(10)} \propto \frac{1}{10} \exp(I_0)$. More generally, in this scenario where all of the pump rays are parallel, the scattered intensity off a pump ray r is

$$I_{s,0} \propto \frac{I_{0,0}}{I_{cell}} \exp(I_{cell}), \quad (3.16)$$

where $I_{cell} = \sum_{r=0}^n I_{0,r}$ is the total intensity in a cell with $n + 1$ pump rays. The first subscript, 0, is used to refer to the pump rays that drive an instability (rather than scattered rays with subscript s), whilst the second subscript number is for the specific pump ray that under consideration. Pump ray 0 is the one that is generating the scattered ray that is being modelled, while the other n pump rays will interact with ray 0 in the cell. Looking at Eq. (3.16) it is clear that when summing over all of the scattered rays launched by multiple pump rays, $I_{s,tot} \propto \exp(I_{cell})$, thus this expression is invariant to the number of pump rays.

Equation (3.16) only considers rays that are all parallel to each other and so can be assumed to all be driving the growth of the same pair of daughter waves, however the picture is further complicated when considering simulations with multiple beams and pump rays that cross over at non-zero angles. Remembering the wave vector matching equation, the orientation of the wave vectors is important for resonant amplification and must conserve momentum. There has been significant experimental and computational work to show that multiple beams can drive common plasma waves to amplify LPI [112, 113]. In the instance where there are two pump rays in a cell that have some non-zero angle between them, the model combines their intensities in a way to account for the amount of each pump ray that is parallel to the other and can drive its corresponding daughter waves. For example, if the angle between them is a small 5° , the rays are roughly parallel and will drive similar waves together, whereas if the rays are at 85° , they are almost perpendicular, and will not drive each other's different orientation daughter waves. The intuitive place to start when

considering how two vectors combine is with the dot product and thus the cosine of the two vectors. So for a pump ray in a cell, the intensity that drives the instability on that pump ray is

$$\begin{aligned} I_{cell,0} &= \sum_{r=0}^n I_{0,r} \cos(\theta_{0r}) \\ &= I_{0,0} + \sum_{r=1}^n I_{0,r} \cos(\theta_{0r}), \end{aligned} \quad (3.17)$$

where $I_{cell,0}$ is the intensity in the cell that the pump ray being considered sees and θ_{0r} is the angle ray r makes with this ray. Thus a different cell based intensity is calculated for each ray in a cell based on its angular relation to the other pump rays in the cell. Alternatives to using the cosine of the angle between rays has been considered, because $\cos(45) \approx 0.7$ and $\cos(85) \approx 0.09$. Therefore, in the case where a second ray is at 45° to our pump it contributes 70% of its intensity towards driving that instability and would even contribute 9% at 85° . Functions like $1 - \tan^2(\theta)$ have been suggested that are still 1 when $\theta = 0$, but decrease at a faster rate, going to zero at $\theta = 45^\circ$. This choice does not have a physical or mathematical basis so is not currently used in the model, which combines intensities using the wavevectors' cosines. With this new definition for the individual cell based intensity seen by each pump ray the our final, general expression for the intensity of a scattered ray generated by a pump ray r is:

$$I_{s,0} \propto \frac{I_{0,0}}{I_{cell,0}} \exp(I_{cell,0}). \quad (3.18)$$

3.2.2.1 Running a Simulation

The calculation of $I_{cell,0}$ for each individual ray is a slow and complicated process. The model already looped over all of the rays present in a simulation to find the plasma parameters along each ray. Looping over every step of every ray had been avoided by using NumPy array operations to calculate parameters in all of the cells a ray passes through. Unfortunately, now a loop over the steps of each ray had to be opened anyway. First, a container called rays per cell (RpC) is created that is the shape of the plasma mesh and records which rays pass through each cell of the plasma and at which of their steps this happens. If any rays are present in a cell of the plasma, that cell of RpC becomes a Python dictionary whose keys are the laser beams present, b . The values paired with each laser beam are each individual dictionaries, each with a key that is the ray number r (in beam b) that is in the cell and a

value that is the step of the ray r in which the ray is in the cell. With the RpC array one can look into any cell in the mesh and find when (which step) any of the rays pass through it.

An array called the intensity mesh or I_{mesh} is also created. It is the shape of the plasma mesh, and in each cell the intensity of all of the rays that pass through that cell is summed. The intensity mesh is important for testing the LPI calculations for pump profile convergence and will be further discussed below.

Next, all three loops over the beams, rays and steps are exited and need to be restarted for the calculation of $I_{cell,r}$, the intensity each ray sees in each cell. All of the information in RpC for every ray in every cell was needed before the $I_{cell,r}$ calculations could start. Once again, it loops over each step of each ray of each beam, first adding the intensity of this pump ray to its $I_{cell,r}$ value and then searching RpC for other rays that also cross through this cell. If any other rays are present in that cell, the angle between they make with the current pump ray is found using their velocity vectors and then their intensity is added to $I_{cell,r}$ as described by Eq. (3.18).

The $I_{cell,r}$ array now has an intensity value that each ray sees in each cell it enters based on the combined intensity of that ray and others in the cell. This determines the growth rate of LPI along that ray.

Next, LPI calculations are started with the new set of plasma parameters in ray dimensions (rather than mesh dimensions). First other plasma parameters are derived from those that the ray tracer output, such as the plasma frequency, the Debye length and the electron thermal velocity. Gradients in parameters like density, temperature and flow velocity are numerically differentiated. These gradients are taken along the paths of rays, resulting from the plasma profiles in space that each individual ray experiences. The wave parameters are calculated as described in Sec. 3.1.1 — iteratively solving the simultaneous equations for SRS and SBS, and interpolating through the TPD parameter look-up table.

The model is now ready to carry out the LPI calculations along each ray. For each ray, the seeds are distributed along the ray depending on the number of daughter rays chosen. Often the number of seeds is high enough compared to the hydrodynamic resolution for there to be multiple seeds per cell that the pump ray is in. For each instability, all of the variables are linearly interpolated in space along the path of the pump ray. For example, $n_e^{ray} \mapsto n_{e,srs}^{interp}$, and the *srs* subscript is needed as each instability will have a different distribution of seeds

along a ray and so different interpolated arrays. The linear interpolation follows the standard expression:

$$n_{e,srs}^{interp}(\text{seed}) = \frac{n_e^{ray}(\text{step} + 1) - n_e^{ray}(\text{step})}{x^{ray}(\text{step} + 1) - x^{ray}(\text{step})} (x(\text{seed}) - x^{ray}(\text{step})) + n_e^{ray}(\text{step}). \quad (3.19)$$

The superscript *ray* refers to our parameters in ray dimensions i.e. spaced in the steps the ray takes through the mesh. We have decided to put a pair of daughter waves at $x(\text{seed})$ the exact location in space of the seed on our pump ray. Step and step+1 are the two ray steps either side of the seed position such that $x^{ray}(\text{step}) < x(\text{seed}) < x^{ray}(\text{step} + 1)$.

With interpolated parameters calculated, the model checks if the daughter waves are in a region where the WKB approximation is valid using Eq.(2.38) and thus determine the appropriate growth rates. The waves generated by each seed are compared to the damping threshold (Eq.(2.55)) and, if switched on, the absolute instability threshold (Eq.(2.56)), however this is rarely used as explained in Sec.3.1.6. The saturated amplitude of this scattered ray is then found by comparing the convective saturation to LDI saturation in the case of SRS and TPD or just using convective saturation for SBS. Finally, the pump is depleted by conserving action based on the saturated amplitude of the daughter waves and this is then repeated for every single instance of the instability that is seeded along the pump ray. This process is then repeated for the other two instabilities that are all also consider.

Next, the model “uninterpolates” depletion. For each instability, new interpolated sets of ray parameters based on the locations of the seeds were generated and the pump was depleted in space according to these interpolated arrays. This leaves one depleted pump profile for each instability with different spatial resolutions. To consider the total depletion of our pump by all absorption mechanisms and map all of the pump rays back onto the original mesh, the interpolated pump profile must be uninterpolated, i.e. mapped back to the ray dimensions, which has one step per cell of the mesh. With a pump profile that has been depleted by all three instabilities, the inverse bremsstrahlung model is applied to get a final pump profile after an iteration of the model.

Once this has been repeated for every ray, a new intensity mesh is created using all of the newly depleted rays. The new mesh is compared to the intensity mesh from the previous iteration and if the difference is sufficiently small ($\leq 0.1\%$) the simulation to have converged on a steady state solution. If not, a new iteration is started, calculating new values of $I_{cell,r}$




Figure 3.18: The fraction of laser energy absorbed by each mechanism for simulations with increasing number of rays making up the beam. This simulation was of a plasma with a cartesian mesh with a maximum density of $0.75 n_c$ and scale length of 400 nm. The beam propagates parallel to the density gradient in the x direction while the mesh only has one cell in the y and z directions.

from this newly depleted intensity mesh and using that to drive LPI along each pump ray, iterating until convergence is reached. This is effectively the same method for handling nonlinear pump depletion as described in Sec. 3.1.3, except extended to 3D and using I_{mesh} rather than the profile of a single ray. Section 3.3.2 details solutions to the cases where the intensity mesh does not converge.

3.2.3 Convergence tests

Sections 3.1.2 and 3.1.6 showed that the method was invariant to the number of daughter rays launched along a pump ray and looked for the minimum number needed for the simulation's reflectivity to converge. This section explores whether the method, particularly the cell based intensity that each ray sees, is invariant to the number of pump rays launched with the ray tracer.

The first test, as pictured in Fig. 3.18, was for a single laser beam propagating parallel to a density gradient through a cartesian mesh. This example had 200 cells in the x direction that the beam propagates through, but only had a single cell in the y and z directions perpendicular to the direction of propagation. This keeps this example as simple as possible by avoiding complications that arise from refraction and crossing cells in multiple dimensions because there is the exact same number of rays in each cell. Importantly, in this scenario the methods described in Sec. 3.2.2 makes the code invariant to the number of pump rays.

Figure 3.19 shows results from a simulation with a single laser beam incident on a spherical mesh where the laser focal spot covers an area of multiple cells and the beam will also refract out of the plasma. This is to explore whether the model still converges when rays are moving between cells angularly causing variations in the number of rays in each cell. Following the discussion in section 3.1.5 regarding using results after only 1 iteration and ignoring the effects of pump depletion on LPI, data is included from both a single iteration simulation (a) and a converged simulation (b). When allowing the simulation to iterate over the pump depletion calculation until convergence the results converge at 1600 rays. At lower rays numbers SRS is present in the simulation, but it does not occur in the simulations with more than 1600 rays.

The single iteration calculation enters a strange regime at this same ray number, and finally converges at 2800 rays. Between 1600 and 2800 rays the results fluctuate between SBS and SRS taking extra energy away from IB. When discussing reflectivities in section 3.1.5 it made sense to use a single iteration calculation as to use a constant intensity pump and thus reduce variables between the comparison. However, when looking at pump depletion there is a limit on how much energy can be depleted from the pump results still have to be capped and potentially ignore absorption happening at higher densities. Thus, a single iteration calculation here may not be as effective as previously and may in fact be the cause for these fluctuating results between 1600 and 2800 rays depending on the location the pump has been fully depleted and further LPI calculations are ignored across the differing ray number simulations.

The results from our simulations are not always going to show convergence if there are 1600 rays in a beam in every single simulation possible. The actual parameter that is important for invariant results is the number of rays present in each cell of the mesh. Enough rays need to be in the cells of interest to reduce the noise from effects like rays moving angularly between cells. For the 1600 ray simulation, at the outer edge of the mesh there were 80 rays in the cells around the beam edge and > 170 rays per cell in its centre. At $n_c/4$ there were > 60 rays per cell in the centre of the beam. When setting up simulations it is difficult to know how many rays per cell there will be in higher density regions of the plasma where the LPI calculations are important, but the values at the outer layer of the mesh are easier to find. Moving forward, having 170 rays per cell at the mesh's outer layer is used as a minimum to guide the set up of simulations.




Figure 3.19: The fraction of laser energy absorbed by each mechanism for simulations with a single beam made up of increasing numbers of rays incident upon a spherical plasma with a scale length $\sim 500 \mu\text{m}$ at $n_c/4$. (a) Shows the results after only running the code for 1 iteration, whilst (b) is the results after iterating until pump profile convergence.

3.3 3D Problems + Fixes

3.3.1 Mesh effects on parameters

Significant issues are encountered when moving the simple one-dimensional models into three dimensions and relying on the plasma mesh to define parameters along each ray. The intensity calculation described in Eq. (3.15) depends on both the distance a ray propagates through a cell and the volume of the cell. This creates two problems with the intensity of each ray that are depicted in Fig. 3.20a: a general upwards trend towards the ray's turning point and constant noisy troughs in the profile. The increasing intensity is caused by the volumes of a spherical mesh getting smaller as a ray propagates inwards, inflating the value of intensity taken from Eq. (3.15). This was part of the motivation to use a cell based intensity for finding LPI growth rates — as the ray propagates into smaller cells there are fewer rays present in these cells and so the total intensity across cells remains fairly constant as is demonstrated in Fig. 3.20a.

The sharp fluctuations in the ray intensity are caused by variations in dx in Eq. (3.15) — the path length of a ray in a cell — which are illustrated by the ray path through cells in a polar mesh in Fig. 3.20b. There are three cells that are roughly the same volume (the differences in this schematic are exaggerated) and a similar path length for the ray in cells 1 and 3. While stepping angularly through cells, the ray passes very briefly through cell 2,




Figure 3.20: (a) Intensity profile of an individual ray up to its turning point and the intensity that ray sees in each cell it passes through as described by Eq. (3.17). This simulation had ~ 150 rays in the cells at the outermost layer of the plasma. (b) Schematic of a ray passing through three cells with significantly different path lengths in each cell as it propagates into a spherical mesh. The scale here is exaggerated for illustrative purposes, most simulations would have higher angular and radial resolution.

in which it has a very small intensity compared to its intensity in cells 1 and 3 due to the very short path length. This happens often when rays step angularly through the mesh and is the cause of the many sharp troughs in the ray's intensity in Fig. 3.20a. Moving to cell based intensities was also anticipated to prevent this as considering hundreds of rays in a cell reduces the impact of each individual ray. The cell based intensity that this ray sees does still show noisy fluctuations, but these are of much smaller relative magnitude to the ray's intensity, so this effect is largely reduced if not entirely solved.

The method for pump depletion via LPI and IB considers the depletion of each ray individually rather than depleting the amount of intensity in each cell. This creates difficulties when using intensity profiles like those depicted in Fig. 3.20a, because they fluctuate so much. It is particularly problematic with the IB calculation which depletes the laser by a fraction of its total intensity. This raises the questions: how can one take a fraction of the ray's intensity away if it is constantly increasing artificially? What does it mean for the ray if in one of the intensity troughs, the various absorption mechanisms deplete the ray intensity to zero despite having only taken a fraction of the non-trough intensity? With these questions in mind it has instead been chosen to treat the rays using their power when considering pump depletion.

The model reads in ray powers from the ray tracer and still calculates ray and cell

based intensities as these are needed to define the LPI growth rates, but everything else is calculated in terms of power. The waves generated by LPI are calculated in terms of power, so the expression for convective saturation becomes $P_s = P_{seed}\exp(G)$. Action is conserved using these changes in power to deplete the power profile of the pump ray. Each pump ray starts with a constant power across its path so the problems of depleting from a fluctuating intensity profile are avoided. This method does make one significant approximation when conserving action between the daughter wave and pump wave as in Eq. (3.9). By swapping the intensities of two different waves for their powers it is assumed that the cross sectional area of the two waves are equal. Building on this, in the future it is anticipated that a more sophisticated method for calculating ray intensities will be implemented and with that there will be no need to convert between intensity into power.

Further problems arise when taking the gradients of profiles along rays as the spatial derivatives are calculated using values per cell, and as illustrated in Fig. 3.20b, the spatial derivative dx can vary steeply. The biggest effect this has on results is affecting the density gradient, because the convective gain exponent $G \propto 1/\kappa' \propto (dn_e/dx)^{-1}$. Figure 3.21a shows the density profile of a ray propagating through a plasma with a manually defined spherical mesh that has a constant density across every radial layer of the mesh. The density that the ray experiences varies smoothly apart from the times the ray steps angularly rather than radially through the mesh as is depicted in Fig. 3.20b as the ray steps between cell 2 and 3. The inset on Fig. 3.21a shows that when the ray steps angularly, it has one step where density remains constant before increasing again as it continues to step radially. Plotting the density gradient in Fig. 3.21b there are regular spikes and troughs every time a ray steps angularly. The inset zooms into two of these discontinuous parts of the profile and shows that there are three steps in the density gradient as part of each discontinuity.

A zoom in on one of the gradient discontinuities is pictured in in Fig. 3.22a. The aim was for the model to detect these discontinuities automatically. The solution was to take the second order spatial derivative of the density profile and search for a 4 point sequence, which has been marked with red crosses on the bottom plot, where the first and second points must have opposite signs, the second and third points must have the same signs and the third and fourth point must have opposite signs, i.e. when:




Figure 3.21: (a) Density profile of a ray propagating into and out of a spherical plasma mesh which has a smooth density profile that is constant on each radial layer of the mesh. The inset zooms in on a few points of the profile where the ray steps angularly through cells. (b) Density gradient of the same ray with an inset again zoomed in on a similar region to examine behaviour when the ray steps angularly.

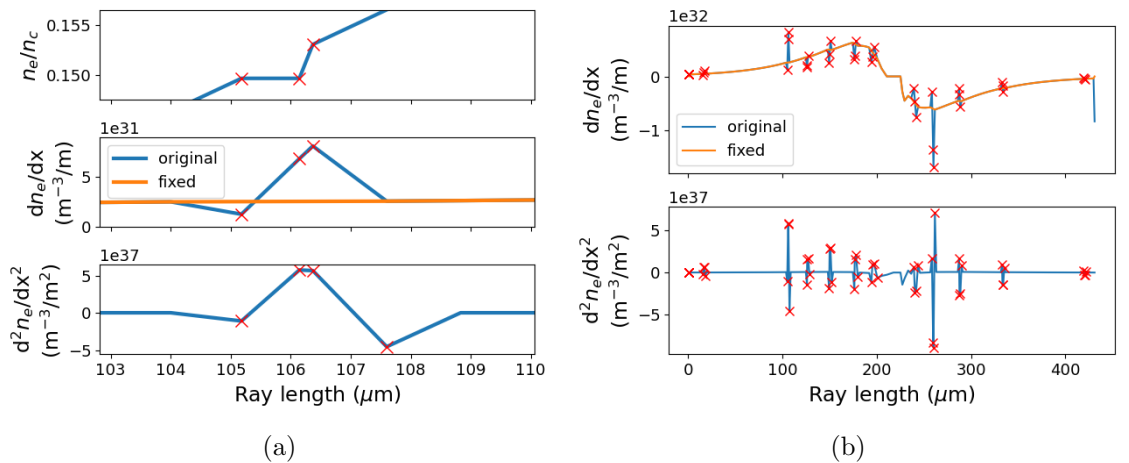


Figure 3.22: (a) The density profile (top) and its first (middle) and second (bottom) order derivatives of the ray from Fig. 3.21 zoomed in to examine one instance where the cell steps angularly and the gradient fluctuates sharply. The orange line labelled fixed shows the gradient after interpolating when an angular step is detected. (b) The two derivatives of the density profile over the entire course of the ray, with red crosses highlighting every time the ray steps angularly and the orange profile again showing the fixed gradient once these fluctuations have been removed.

$$\begin{aligned}
\frac{d^2 n_e}{dx^2}(x) \frac{d^2 n_e}{dx^2}(x+1) &< 0, \\
\frac{d^2 n_e}{dx^2}(x+1) \frac{d^2 n_e}{dx^2}(x+2) &\geq 0, \\
\frac{d^2 n_e}{dx^2}(x+2) \frac{d^2 n_e}{dx^2}(x+3) &\leq 0,
\end{aligned} \tag{3.20}$$

is satisfied. In this scenario, the first three points of the sequence, as marked by crosses in the n_e and dn_e/dx plots, are those where the gradient steeply fluctuates. A new density profile, labelled as “fixed”, is created and is identical to the original gradient, except that it ignores the three discontinuity points and instead linearly interpolates the values of dn_e/dx either side of this region. Figure 3.22b shows the whole density gradient and its spatial derivative, with red crosses marking the points the code has automatically found where the condition in Eq. (3.20) is met. Ignoring the very first and last steps of the ray, the method successfully identifies every time there is an angular step through the mesh leading to the fluctuations in density gradient and the fixed density gradient thus shows a much smoother result. For more realistic problems where plasma profiles come from the outputs of a radiation-hydrodynamics code, it is rare that the density profile will be quite so idealised because profiles are much noisier and will not be constant around each radial shell of a mesh. However, when combined with a windowed average as a smoothing technique to reduce noise, this technique is still useful as two cells next to each other in the same radial layer will have very similar if not identical density. Thus the same problems will occur in the density gradient, and can still be found and removed by examining $d^2 n_e/dx^2$.

3.3.2 Failure to converge

When a pump ray is depleted by large amounts over short distances the model will sometimes fail to converge. Figure 3.23 shows results from simulations of a single pump ray being depleted by convective SRS in a homogeneous plasma. In a homogeneous plasma, there is no decoupling of waves as the wave vector remains constant so for simplicity in this simulation a single scattered ray is launched and the pump depletion is calculated at the end of the region of plasma ignoring any spatial detail in the depleted pump profile. This simulation used the spatial growth rate used in Eq. (2.71) to determine the amplification of the seed which was launched from the opposite side of the plasma so that it could amplify over its entire length. As was demonstrated by Fig. 3.6, the code incorporates the nonlinear effects of

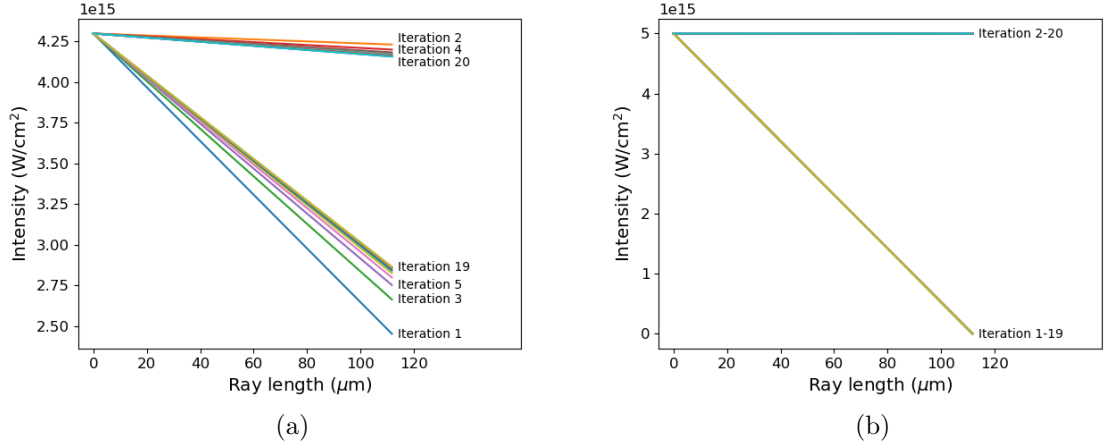


Figure 3.23: Pump profiles from simulations using (a) 4.3×10^{15} W/cm² and (b) 5×10^{15} W/cm² intensities, in a homogeneous plasma with $n_e = 0.1 n_c$, $T_e = 3.5$ keV and a seed amplitude $I_{seed} = 10^{-4} I_0$. Neither simulation converges so the pump profiles after the first 20 iterations over the LPI calculation are presented.

pump depletion by iterating over the LPI calculation with the most recently depleted pump profile used to drive the instability growth until the pump profile converges. Figure 3.23a shows an example where after iterating over the calculation, the pump fails to converge, instead tending towards two distinct profiles that it alternates between after each iteration. The higher pump intensity and thus higher LPI gain in Fig. 3.23b create a more extreme version of this problem where the pump profile oscillates between entirely depleted or not depleted at all after each iteration.

In the scenarios where the simulations result in these pump profiles that do not converge and instead alternate between 2 different profiles, it is reasonable to assume that the steady-state solution lies somewhere in the centre of the alternating profiles. To try to force the model towards a converging solution, a guess is made after iteration 20, labelled as such in Fig. 3.24, which took the average of the pump profiles after iteration 19 and 20. Iteration 21 then used this new averaged profile to determine LPI growth, anticipating that it might have pushed the simulation towards converging upon the solution. This graph plots the pump profiles until the sixtieth iteration, instead finding that the pump returns to the two alternating values it reached after 20 iterations. This behaviour is replicated when repeating this process with the higher intensity example in Fig. 3.23b.

Estimating the steady-state pump profile was attempted again, but this time with a different guess after iteration 20 that was in the gap in Fig. 3.24 between the first guess




Figure 3.24: Results from a simulation using the same set-up as in Fig. 3.23a. Now the pump profiles after iteration 19 and 20 are averaged to create a new guess at the answer in between them, and the simulation is continued from here for 60 iterations in total.

and iteration 21. Once again, the following iterations led to two separate alternating pump profiles as has already been demonstrated. An important difference was that this time it took more iterations to return back to the alternating profiles because the rate at which the profiles diverged after the guess was much slower, i.e. the difference between each odd iteration after iteration 21 and each even iteration after iteration 22 were initially much smaller. This indicated that there may be a saddle point-like solution where the code may stabilise and stop returning to the alternating pump profile result. By manually altering that guess after iteration 20 a relatively stable solution to this problem was found. The solution was not completely stable as if more iterations were ran it did eventually diverge to the alternating solution, however those initial iterations after the guess were fractions of a percent away from each other — close enough to suggest that it was in the vicinity of the saddle point-like solution to this simulation. The automation of this process for both example simulations is depicted in Fig. 3.25.

In figure 3.25a the same case as in Fig. 3.24 is solved. The method starts by again making a guess in between the result from the last two iterations, 19 and 20. Using this new guess as the pump to define the growth rates, the LPI calculation are run again and result in the profile labelled iteration 21. The difference between the profile after the iteration 20 guess and iteration 21 is checked to see if it meets the convergence criteria (within 0.1% of each other) and if not, as is the case here, a new guess is made in between these two profiles as




Figure 3.25: The pump profiles from the two respective simulations in Fig. 3.23 but now with an automated method that continually guesses in between the diverging pumps to search for the saddle point-like solution.

to try to force convergence upon the saddle point. The next guess and following iterations are depicted in the inset. From the guess after iteration 21, the LPI calculation is run for a 22nd iteration to get a new profile, which also hasn't met the convergence criteria. Thus a new guess is made after iteration 22 and the calculation is run once more, this time finding that this guess and the profile after iteration 23 do meet the convergence criteria and a steady-state solution has been reached.

This process is repeated for the higher intensity case in Fig. 3.25b, and again finds success. This simulation takes an extra 9 iterations to find the solution because the unknown region that is being searched is much larger (the gap between the profiles after iteration 19 and 20).

Simulations which changed the initial guess were also performed to explore whether this value, currently selected arbitrarily, was having an impact on where the code finds the saddle point to be. It was verified in both of the examples in Fig. 3.25 that the initial guess had no bearing on the final value of the saddle point, but did affect the number of iterations it would take to reach it.

A crude guess when faced with a simulation that fails to converge like these two, would be to assume that the solution is the average of the two alternating pump profiles. However, in Fig. 3.25b the pump is only depleted by $\sim 20\%$ rather than 50%. This method is effective as in this simulation with increasing intensity until $\sim 7 \times 10^{15} \text{ W/cm}^2$, where after the initial averaged guess the next iteration returns to the fully depleted pump seen in the fig. 3.23b. The LPI gain is so high that even the averaged pump guess can not stop it from fully

depleting. This may be solvable by getting the model to try different guesses which are smaller than the average pump profile in these scenarios. This hasn't been implemented as, while it may allow for an interesting development of the physics understanding, this one-dimensional homogeneous problem is quite limited in how it can be extended to three-dimensional simulations with multiple instabilities per pump ray whose intensity interacts with other pump rays in the same cell.

The physical significance (if any) of this result where the pump ray converges upon two alternating profiles is something that has been considered. It is not uncommon to see “bursting” scattered light from SRS and SBS in experiments where the reflectivity fluctuates rapidly in time. Cristoforetti *et al.* [114,115] record this in multiple experiments exploring SRS and TPD, and Škorić *et al.* [116] had further developed theory and successfully modelled this behaviour for SRS. These behaviours are always noted on roughly picosecond timescales and result from examples with very high laser intensities, often approaching 10^{16} W/cm² as this drives kinetic effects like electron trapping in plasma waves that are necessary to create the bursts [117]. It had been suggested that the nonconvergent behaviour this model displays is a manifestation of bursting LPI as it is driven by high LPI levels and that the saddle point-like solution may be a time averaged result of a bursting instability. Unfortunately, without any feedback between the plasma and laser in our steady-state model it is unlikely that this behaviour is a representation of bursting LPI and is likely just a computational artefact.

Figure 3.26 contains an example of this problem from a three-dimensional simulation. When considering the convergence of the pump, the total intensity in each cell of the mesh is now used rather than considering individual rays. In this simulation there was high gain in a few regions and I_{mesh} failed to converge in these cells. The absorption profiles of one of the rays that passed through these cells is plotted to illustrate the problem. As explained in Sec. 3.2.2, new interpolated pump profiles are generated for the absorption caused by each instability on a ray — these are the SRS and SBS profiles on the graph. These two profiles are then uninterpolated onto the original pump profile with spatial resolution defined by the plasma mesh. Finally, the IB model is applied to the pump ray and its depletion for this iteration is complete. In this simulation, after the ninth iteration the ray is almost entirely depleted by SRS that occurs around $0.22 n_c$. Although the SBS model has found that SBS will also significantly deplete the pump after $0.3 n_c$, the pump has already been depleted by the time it reaches here and no SBS can occur. The 10th iteration then starts with the

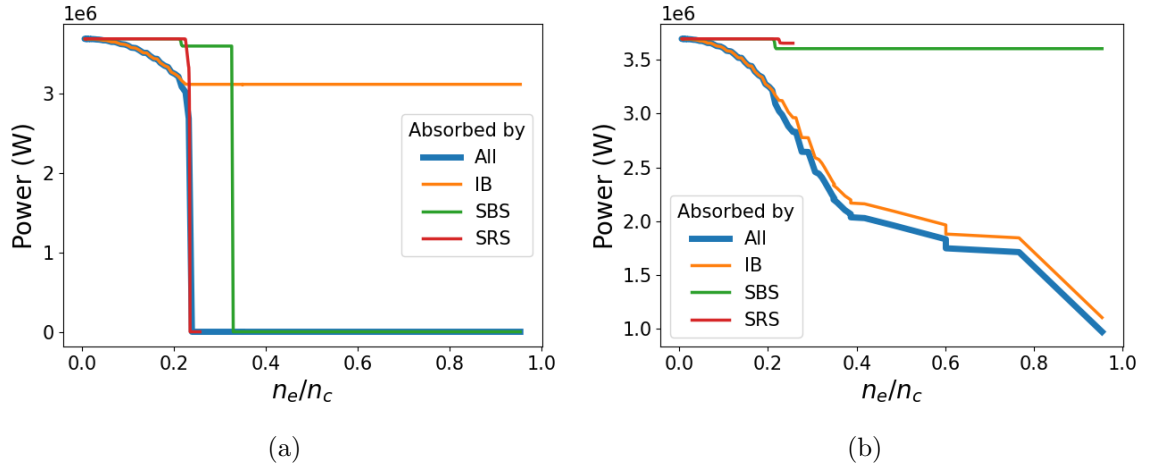


Figure 3.26: Various absorption profiles of a single ray from a simulation after (a) 9 iterations and (b) 10 iterations. This ray was one out of 900 rays making up a laser beam with intensity $\sim 3 \times 10^{15}$ W/cm². Each line represents the amount the pump is depleted by by an individual absorption mechanism, all of which are then combined to create the “total” depleted pump profile. TPD was not included in this simulation for simplicity.

pump seeing an intensity profile in the cell which has fully depleted at $0.22 n_c$ (as all of the other rays in the cell were also depleted) and so can not predict any LPI beyond this density. Instead, a small amount of SRS and SBS occur before this point and then the rest of the ray is only absorbed by IB. The iterative procedure fails to converge upon absorption profiles for each mechanism because of the extreme depletion caused by so few scattered rays.

In the above, homogeneous example where only a single ray was considered the approach was to try to force the ray towards a convergent solution and test this by rerunning the LPI calculation. When working in three dimensions with multiple rays this becomes a lot more time consuming: the profile of every ray that is in a cell that has not converged must be averaged, then the mesh based intensity that each ray sees must be calculated again, and then the LPI calculations must be restarted over all of the rays because of the fact that changing one pump ray can affect LPI growth on other pump rays near it. The method used on the single ray would add a lot of runtime to the simulations. Perhaps in the future, when this code is ported into Fortran, this approach could be tried, but currently there hasn’t been the time to implement it. Currently, no attempt is made to force convergence, instead the model take averages.

The first attempt, shown in Fig. 3.27a, took the average of the two total pump profiles (blue) to give a final result, and then also averaged each absorption profile for diagnostic




Figure 3.27: Various absorption profiles of a single ray from a simulation after (a) 9 iterations and (b) 10 iterations. This ray was one out of 900 rays making up a laser beam with intensity $\sim 3 \times 10^{15}$ W/cm². Each line represents the amount the pump is depleted by by an individual absorption mechanism, all of which are then combined to create the "total" depleted pump profile. TPD was not included in this simulation for simplicity.

purposes. The problem here is clear, the SBS pump profile is depleted by 50% but it has no effect on the total final profile and so the diagnostic fails to conserve energy as it sees >100% absorbed by the three individual instabilities but <100% of the total pump's power absorbed (before its turning point). It arises because the depletion of the SBS profile in iteration 9 has no effect on the already depleted pump.

One approach would be to say that when averaging iteration 9 and 10, the potential SBS that occurs after 0.3_c should be ignored because it does not absorb pump energy in iteration that is being averaged. On the other hand, if one assumes that in this steady-state, energy is getting past $n_c/4$ after taking an average of the SRS depletion, then it makes sense that this ray would then undergo SBS, even if the last two iterations failed to capture it. This outlines another reason why solving this problem in inhomogeneous plasmas is difficult. Once a solution has been found for the ray profiles in the cell that is failing to converge, the LPI physics along the remaining paths of these rays needs to be recalculated using the new intensities. However, recalculating LPI levels using averaged profiles is not desirable for the sake of computational efficiency.

The current approach follows the order of the absorption calculations on the ray and is shown in Fig. 3.27b. First, the model takes the average of the LPI profiles and uninterpolates these to deplete the pump ray, and then IB is applied to this pump profile. This method

solves the concerns from the previous question; by recalculating the pump (rather than taking an average) it ensures that the pump energy going into each absorption mechanism is conserved. It also suggests that SBS has an effect on the pump despite it not appearing to in iterations 9 and 10.

Another consideration when dealing with the problem of nonconverging pump profiles is that they are often caused by extreme amplification of individual scattered rays, which can be non-physical; occurring as a result of computational artefacts from dealing with the rays in the mesh as discussed in Sec. 3.3.1 or even noisy plasma profiles from the hydrodynamics simulation. Reducing the prevalence of these is one way to reduce the occurrences of nonconverging solutions and further increase the accuracy of our model.

Chapter 4

Results

4.1 Pump Depletion Benchmark

A test of Chiron was based on a benchmark of the electromagnetic code FLAME developed by Hao *et al.* [13]. They carried out separate simulations for SRS and SBS in a homogeneous plasma to look at how reflectivity varies with pump intensity. They benchmarked their code by comparing it to solutions to Tang's analytic expression for reflectivity [118]:

$$R(1 - R) = \varepsilon \left(e^{G(1-R)} - R \right), \quad (4.1)$$

where the seed amplitude is $\varepsilon = I_{seed}/I_0 = 10^{-4}$ and G is the convective gain in a homogeneous plasma. Without gradients in the plasma there is no decoupling of the waves to limit the convective gain as their wave vectors remain constant. The convective gain is

$$G = \frac{\gamma_0^2 \delta x}{v_{gs} \nu} \quad (4.2)$$

where ν is the rate of damping on the EPW or IAW and δx is the width of the plasma, i.e. the distance over which the waves amplify.

The LPI models developed in chapters 2 and 3 are only built to function in inhomogeneous plasmas with the focus of this work being application to ICF experiments. Convective saturation relies on the assumption of wave dephasing captured by the κ term or its derivative and without this the models will not work in a homogeneous plasma. Despite the limited applicability of the LPI gain in these two different scenarios, this benchmark is a useful way to test the iterative solution to pump depletion as so many other competing features of the

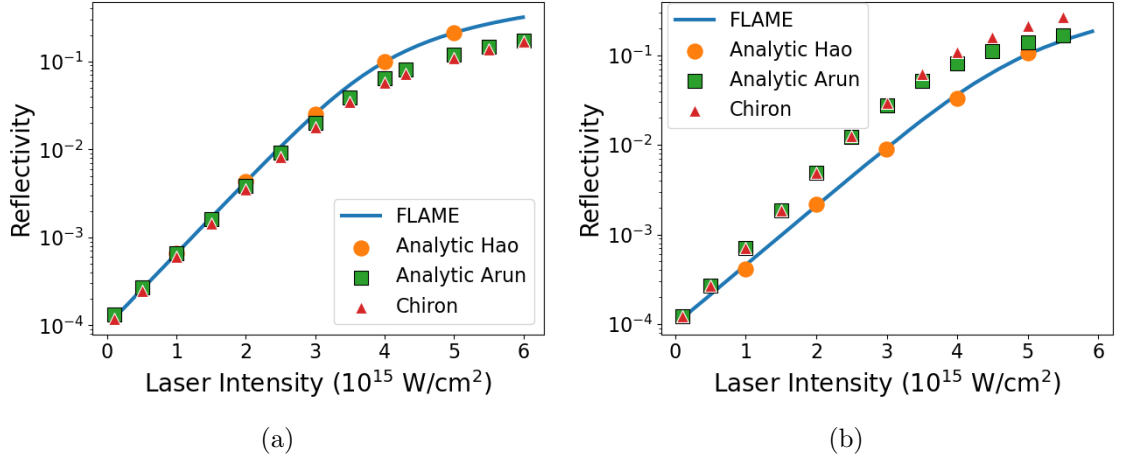


Figure 4.1: Various (a) SRS and (b) SBS reflectivities in a homogeneous hydrogen plasma with $n_e = 0.1n_c$, $T_e = 3.5 \text{ keV}$ and $T_i = 1.6 \text{ keV}$. FLAME is the EM code developed by Hao *et al.* [13] and the data series labelled analytic Hao and analytic Arun are their solutions and mine to Tang’s analytic reflectivity (eq. (4.1)).

model are not included. Only one pump and scattered ray need to be launched, the $\kappa = \kappa'x$ approximation is not necessary and there is no competition between convective and absolute saturation.

Figure 4.1 shows the results of the benchmarks for convective SRS and SBS. The data series labelled FLAME and analytic Hao were both taken from Hao’s paper, and Chiron is shows the results from my model. After finding that Hao’s analytic solutions did not solve eq. (4.1) using my expression for gain, I also solved the analytic expression using an iterative method, that starts with a guess for reflectivity, solves the equation for reflectivity using this initial guess and uses the new reflectivity solution to start the next iteration, repeating until the reflectivity solution converges.

The SRS model intially failed to converge at a pump intensity of $5 \times 10^{15} \text{ W/cm}^2$, but using the method described in sec.3.3.2 the model was able to force itself to find a stable intensity profile that is considered converged. The SBS model initially failed to converge at $5 \times 10^{15} \text{ W/cm}^2$, but similarly was able to force convergence. The analytic solutions converged across the entire intensity range.

The solution’s proided by Chiron start diverging from FLAME ad Hao’s analytic solutions at intensities above $3 \times 10^{15} \text{ W/cm}^2$. For SBS Chiron’s and FLAME’s solutions become different as low as $1 \times 10^{15} \text{ W/cm}^2$. The results provided by Chiron do however match my solution to the analytic reflectivity, apart from for the very high intensity simulations

for SBS. The fact that Chiron's solutions match my analytic solutions but not Hao's, and Hao's solution does not solve eq. (4.1) for me suggests that there is a discrepancy in our gain terms. The similarity between Chiron's and my analytic results indicate the iterative solution to pump depletion works.

4.2 Conical Target Experiment

The experiment chosen to test Chiron against was carried out by Scott *et al.* [119] at the OMEGA laser facility. Wanting to explore LPI in shock ignition conditions, they designed a conical target pictured in fig., 4.2. Twenty OMEGA beams can be used on planar targets, with the beams in three cones of constant polar angles of 23° (5 beams), 48° (5 beams) and 62° (10 beams) to the target normal. The high incidence angle of the 62° beams means they refract away from the target at low density, decreasing the coupling of laser energy into the plasma. By introducing the open cone target design, the angle between these beams and the target normal is significantly reduced, increasing coupling efficiency and allowing for the creation of long scale length plasmas relevant to shock ignition.

One of the principal concerns with shock ignition designs for laser fusion, which this experiment was designed to investigate, is the potential for severe losses to LPI due to spike




Figure 4.2: The open cone target design with one beam shown for each cone of 5 beams at polar angles of 23° , 48° and 62° out of the low intensity beams shown. The 48° and 62° were turned on initially at lower intensity to generate the long scale length plasma and then the high intensity 23° beams were used to investigate laser-plasma interactions. Two of the 23° beams had FABS detectors behind them, one to collect light generated by SRS and TPD and the other for SBS and CBET generated light.

in laser profile that will lead to intensities up to 10^{16} W/cm². These laser intensities are beyond the upper limit set by the fluid theory that Chiron is based upon and so one of the lower intensity experiments which were within the design-region of the code was chosen. This chapter will examine a lower intensity shot labelled 88980 that did not use all of the twenty laser beams. The shot used the five 48° and five of the 62° beams initially to create the long scale length plasma and then used only the five 23° beams after these to generate a high intensity pulse for LPI studies. The overlapped intensity of the plasma creation beams was predicted to be 0.9×10^{15} W/cm² and the higher intensity beams were 2.5×10^{15} W/cm². Other shots used the remaining five 62° beams to drive the second laser pulse to intensities up to 8.5×10^{15} W/cm².

Figure 4.3 shows scattered light spectra recorded by the FABS diagnostic from this experiment. The two lobes at either side of 700 nm that occur at around 5 ns are the $\omega_0/2$ light generated by interactions involving the TPD EPWs as described in sec.2.2.6. The continuous spectrum that occurs after 1 ns is scattered light from SRS. The horizontal band at ~ 600 nm and diagonal line pointing towards the bottom right corner which register decreased scattered light are a result of damage on the detector.

The beginning of the Raman scattered light is aligned with the turning on of the higher intensity beams. This could be because the initial lower intensity beams were below the threshold for SRS to grow. Alternatively, it might be because the FABS diagnostic was positioned in one of the 23° beam ports so it was able to pick up backscatter. The low intensity beams were at large angles away from the FABS detector so any backscatter they produced would be less likely to be recorded. The other notable feature in the SRS spectrum is at ~ 2 ns there is a maximum in scattered light at 550 nm and the previously present scattered light above 600 nm has disappeared. The lack of light above 600 nm is likely due to the pump being more depleted because of the increase in SRS at lower densities (shorter wavelengths). It is also true that at later times as the plasma continues to expand, collisional absorption becomes on the scattered rays becomes more significant, however this is unlikely to account for such a large change in the spectrum. There is also no scattered light at 700 nm which is sometimes seen in laser fusion experiments, like the FABS data shown in fig. 2.7 for example.

The scattered light from two-plasmon decay is only detected while the lower intensity beams were on. Light scattering off of the TPD plasmons occurs at non-zero angles to the

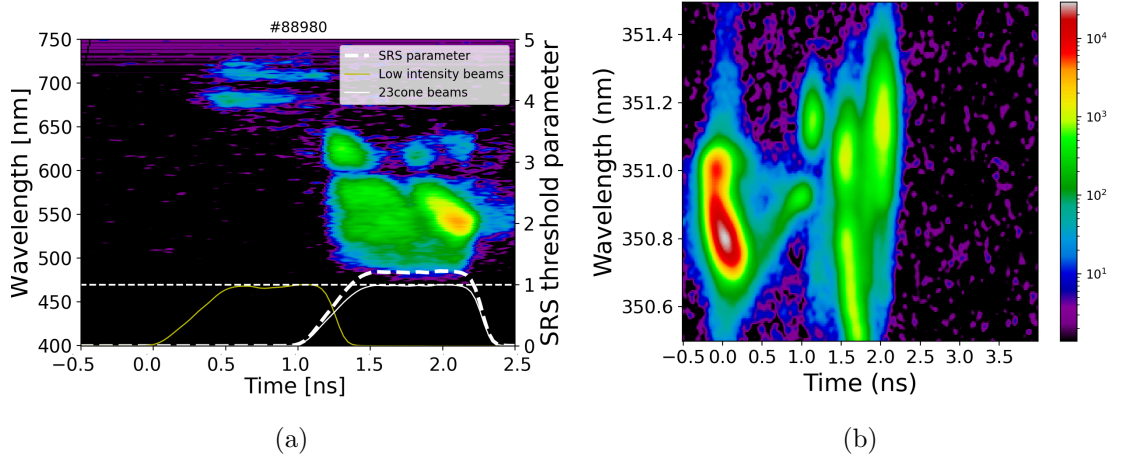


Figure 4.3: Backscattered light captured by the FABS diagnostic, which was positioned behind one of the 23° beams. White lines show laser pulse. The detector shown in (a) was sensitive to scattered light generated by SRS and TPD, while (b) was sensitive to SBS scattered light.

laser beam so this light generated by the higher angle TPD may have been better oriented to hit the FABS detector. Pump depletion could also be a factor in why TPD is seen during the low beams but not the high beams. More convective SRS means less laser energy reaching the quarter critical surface such that the laser intensity may be below the TPD threshold. Figure 4.3a only shows SRS during the high intensity beams, which would appear to support the pump depletion argument, however it is possible that SRS is occurring during the low intensity beams (or TPD during the high intensity beams) and that the scattered light does not reach the FABS detector. The higher intensity beams will be more prone to driving LPI as a result of their higher intensity and the plasma scale length being longer later in time does mean that more SRS is expected to occur then and so less energy would reach the quarter critical surface.

Figure 4.3b is the scattered light signal from the FABS detector designed to measure light near the same wavelength as the laser, i.e. Brillouin scattered light. There is a strong signal when the low intensity beams turn on, which is commonly seen in scattered light data at the laser wavelength [85,120]. This could be SBS, but it has also been attributed to laser light that is reflected off of the outer surface of the expanding plasma [121]. Its Doppler shift away from the laser wavelength gives information on the motion of the surface [122]. Reflection of laser light off of the plasma is not a feature included in Chiron so this will not be recreated in the simulation. Brillouin scattered light is seen during the peak of the high

intensity beams.

4.3 The Chiron Simulation

To model this experiment using Chiron a simulation was carried out using H2D, the two-dimensional version of the radiation-hydrodynamics code Hyades [105], with 3D laser ray-tracing and SESAME equations of state [123]. As Chiron is not yet integrated with a hydrocode, the LPI absorption physics are modelled separately and without feedback into the h2d simulation. Thus, because it was run with no LPI modelling, the h2d simulation likely overestimated the inverse bremsstrahlung absorption of laser energy, leading to different evolution of the plasma dynamics.

Two approaches to the laser beams were explored when setting up the simulations. The first simulated only three laser beams with each beam at a different polar angle to represent an entire cone of five beams at constant polar and varying azimuthal angle. Each of the 3 simulated beams had the power of the summed 5 beams in that cone of beams. The other simulation modelled all 15 beams at their unique polar and azimuthal angles. This should have lower intensities driving LPI growth as the intensity of each crossing beam is summed following the description in sec. 3.2.2. Figure 4.4 shows the fraction of laser energy absorbed by each mechanism (or not absorbed) using the hydrodynamic outputs from the h2d simulation every 0.2 ns apart from between 1 and 1.4 ns where an error occurred and no h2d outputs were saved.

Figure 4.5 shows the fraction of laser energy absorbed by IB and all of the LPI in the two simulations with different beam set ups. The simulation that models all 15 beams had lower levels of LPI and higher IB absorption until 1 ns, which was the anticipated outcome. After 1.4 ns the amount of IB in both simulations dips and LPI increases. This is plausible at this time in the simulation as the high intensity laser is at its maximum intensity for the first time so LPI growth rates will be at their highest until that point. The fact that the all beam simulation sees higher levels of LPI to the three beam simulation between 1.4 and 1.8 ns was not expected. The rays in the all beam simulation will have lower cell-based intensities and thus smaller LPI growth rates than the rays in the three beam simulation.

Figure 4.6 shows a close up of fig. 4.4 to better examine the interplay between the different instabilities. The all beam simulation having larger total LPI levels to the three beam

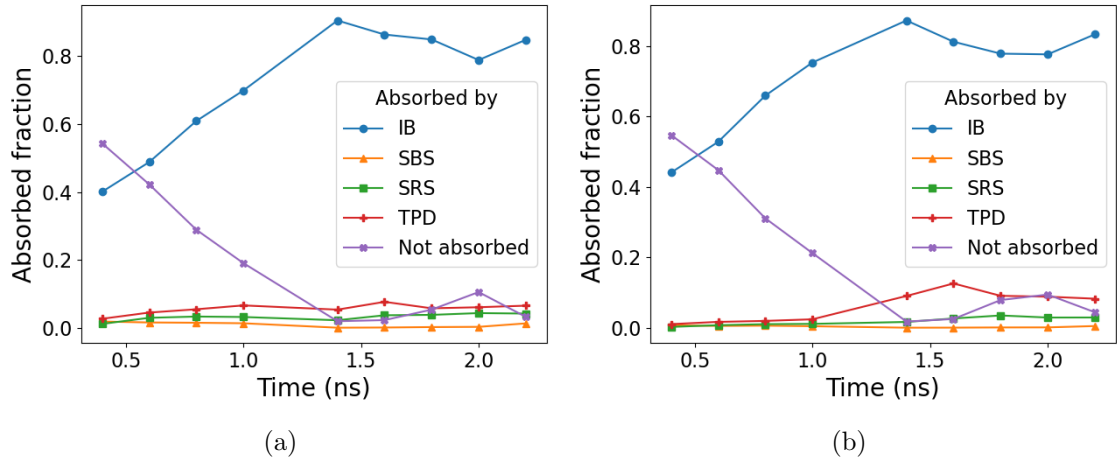


Figure 4.4: Fraction of laser energy absorbed by IB, each LPI or not absorbed after applying Chiron to the hydrodynamic outputs from h2d simulations of experiment 88980. Two different simulation set ups were tested: (a) used 3 laser beams with one beam per cone of beams at constant polar angle, (b) simulated the propagation of all 15 beams. The total power used in each simulation was the same.




Figure 4.5: Fraction of laser power absorbed by IB and the summed total of all three LPI for the three beam simulation (blue) and the all beam simulation (orange).




Figure 4.6: The same data presented in fig.4.4 but with a reduced limit on the y axis to allow a more detailed view of the LPI interplay. (a) is from the three beam simulation, (b) is from the all beam simulation.

simulation after 1.4 ns is due to a large increase in TPD during the high intensity beams in the all beam simulation. The level of TPD in the three beam simulation stays at approximately the same level between the low and high intensity beams.

The increase in laser energy that is not absorbed at 1.8 and 2 ns is also unanticipated and raises further questions about the simulations at these times. The plasma should get longer with increasing time so the laser path lengths within it will increase and thus increase the chance of absorption via IB or LPI. This is clear when looking at the not absorbed fraction decreasing (and IB in turn increasing) between 0.4 and 1.4 ns.

Figure 4.7 shows the density profiles of all of the rays launched in the 3 beam simulation by the high intensity beam at 1.6, 1.8 and 2.0 ns. The unabsorbed laser energy at these three times in this simulation is seen in fig.4.6a. Out of these three times, at 1.6 ns the maximum laser energy is absorbed via IB and there is a minimum in the laser energy not absorbed. This unabsorbed energy then increases at 1.8 ns and increases even further at 2.0 ns. Most rays depicted in fig.4.7 have an approximately symmetrical density profile that increases up to the maximum at the ray's turning point and decreases as the ray propagates out of the plasma. At 1.8 ns, fig.4.7b shows a number of rays that exit the plasma directly after reaching their turning point, visible as the vertical lines at $2500\mu\text{m}$. These rays are likely to be significantly less absorbed in the plasma as they propagate half of the distance that the other rays do. These rays that exit the plasma are even more numerous at 2.0 ns and non-existent at 1.6 ns, which fits as an explanation for the increasing amount of laser energy




Figure 4.7: Density profiles of all 1600 rays launched by the high intensity, 23° beam in the 3 beam simulation at (a) 1.6 ns, (b) 1.8 ns and (c) 2.0 ns.

not being absorbed across these three times. It is likely an error in the hydrodynamic output or the ray trace that has caused this and can be rectified in the future. It is unlikely to affect the LPI physics much as most rays have significantly been absorbed by their turning point such that they are below the LPI intensity thresholds and mostly undergo IB absorption on their way out of the plasma.

4.3.1 Stimulated Raman Scattering

4.3.1.1 SRS in the high intensity beams

The SRS scattered light data presented in fig. 4.3a only shows Raman scattered light during the high intensity, 23° beams. Figure 4.8a is a synthetic FABS diagnostic using scattered light from only the high intensity beams in the three beam simulation. As the model does not include ray tracing of the scattered light rays this assumes that all of the Raman scattered light is collected by the FABS detector. The simulation includes collisional absorption on the scattered light rays and also the group velocity correction to power described in sec. 3.1.4.

The experiments maximum in scattered light intensity at 2 ns and 540 nm has been almost exactly reproduced. The precursor to this peak in the experimental scattered light spectrum is an amplitude increase that starts slightly earlier 550 and 600 nm before moving to the maximum at 540 nm. The scattered light from the simulation also has an increase in scattered light between 550 and 600 nm at 1.8 ns, preceding the simulation at 2 ns that contained the spectrum's maximum.

Another feature of the experimental spectrum is the lack of scattered light between 600

and 650 nm after 1.6 ns. The simulated spectrum does show a small decrease in the amount of scattered light in this wavelength range at later time, but does not replicate the experimental spectrum which has a much larger difference.

Below 500 nm there is no scattered light recorded in the experiment. Although the simulated scattered light at this wavelength shows a decrease of approximately an order of magnitude in the scattered power, this is not as significant as the cut off in the experiment. At the low densities the fluid modelling of the plasma becomes less appropriate as kinetic effects become more important. Seaton [92] uses a fully kinetic expression for Landau damping as he says the expression derived from fluid theory is not appropriate when $k\lambda_D > 0.3$. This may also be comparable to TPD being above threshold at very low densities discussed in sec. 2.2.6.1 and fig. 2.15a. Similar to the approach taken there, a hard cut off could be placed on the SRS model, not allowing anything to grow below ~ 500 nm.

Figure 4.8b uses the same colour scale to show the scattered light from the high intensity beam in the same simulation except without applying any inverse bremsstrahlung absorption to the scattered rays. The difference in scattered power measured by this diagnostic and in fig. 4.8a is larger at longer wavelengths as these rays are generated at higher densities further into the plasma and so undergo more absorption. The synthetic diagnostic for scattered light that included IB shares many more similarities with the actual experimental data. Figure 4.8b does not reproduce the same drop off in scattered power at wavelengths over 650 nm and also has a much larger and differently shaped region of maximum scattered power. The difference in scattered power shows the importance of considering the effects of IB on the scattered light in these simulations. It also indicates that analysis of scattered light data from experiments should always bear in mind that if only shorter wavelength Raman scattered light is being measured, it is likely higher wavelength Raman is still occurring but is being absorbed. The comparison between these two figures also supports the argument that the Raman gap is caused by collisional absorption, as the biggest difference between the two diagnostics is above 650 nm — the region the Raman gap is seen. Typically the Raman gap is used as a descriptor when there is visible light scattered at 700 nm, which isn't present in the experimental or simulation data here. In differing conditions, the scattered light at 700 nm can grow absolutely to much higher amplitudes so despite the significant absorption of the emitted light above 600 nm this will still be present.

In addition to collisional absorption, refraction may play a role in measurements showing




Figure 4.8: Synthetic FABS image created by collecting all of the Raman scattered light generated by the high intensity beam from the three beam simulation. The correction to power due to varying group velocity described in sec. 3.1.4 is applied to the scattered rays. (a) collisionally absorbs the scattered rays as they propagate out of the plasma, whilst (b) does not.

less longer wavelength scattered light. Chiron does not trace the scattered rays and assumes they are all collected by the FABS diagnostic, but in reality these rays are also subject to refraction. The rays that have longer wavelengths are further away from the laser wavelength of 351 nm so are more likely to undergo differing paths due to refraction. They are also generated further into the plasma so travel a larger distance and thus have more chance to be refracted differently to the laser's path. Therefore, refraction could be an additional factor in reducing the amount of higher wavelength scattered light measured in experiment that Chiron is currently unable to capture.

The scattered light data presented so far in this section came from the three beam simulation. The synthetic FABS in fig. 4.9a shows the scattered light from one of the five high intensity beams from the simulation that modelled all 15 laser beams separately. The scattered light data from each of the five high intensity beams from this simulation looks almost identical. This simulation has done a worse job of recreating the FABS image than the simulation that only used 3 beams. The maximum here is too early in time and is at higher wavelengths. Above 600 nm this simulation does not see the same decrease in scattered light either.

Figure 4.9b is the scattered light from all five of the high intensity beams from the all beam simulation using the same colour scale as in fig. 4.8. Given these five beams all had different azimuthal angles their scattered light could not be collected by a single FABS




Figure 4.9: Simulated Raman scattered light diagnostics created from the high intensity, 23° beams in the all beams simulation including collisional absorption and the group velocity correction on the scattered rays. (a) For the beam that had the FABS in its port in the experiment (b) The simulated Raman light summed from all of the five high intensity beams and using the same colour scale as the three beam simulation FABS images in fig. 4.8.

diagnostic but it is useful to compare the magnitudes to the results from the three beam simulations. As all five high intensity beams had virtually identical spectra, this has the same shape as fig. 4.9b but with a five times increase in magnitude. The reflectivity from this simulation is lower than the three beam simulation, as anticipated based on the fact the SRS levels are lower as shown in fig. 4.6. This simulation also generated a larger fraction of its light at higher wavelengths, so has a larger fraction of this light being collisionally absorbed.

It is unexpected that the simulation modelling all of the beams has performed worse than the simulation using one beam per cone given the fact that the setup is closer to the experiment. The method used for adding the intensity of crossing rays is likely to be incorrect, but that does not directly explain why the shape of the spectrum is better matched by the three beam simulation, when the main difference is an increase in intensity in that simulation. One might expect the spectra to be similar in shape, but of different amplitudes. One factor is that the higher intensities in the three beam simulation drove more SRS at lower density (shorter wavelengths) and thus drove less SRS at higher wavelengths as the pump had decreased intensity at higher plasma densities. Whereas the lower laser intensities in the all beam simulation did not drive as much low density SRS, so did not deplete as much and in turn generated proportionally more of its scattered light at higher wavelengths.

4.3.1.2 SRS in the low intensity beams

During the low intensity beams there is no experimental SRS signal in fig. 4.3. The three beam simulation shows significant SRS levels absorbing almost 4% of the laser energy during the low intensity beams which at a comparable level to the amount of SRS during the high intensity beams in this simulation. The simulation of all of the laser beams had very low levels of SRS, taking less than 1% of laser energy, and importantly significantly less than during the high intensity beams in this simulation, implying that the SRS does “switch on” here like the FABS shows. The FABS diagnostic is in one of the ports of the high intensity beams so is much more likely to detect backscattered light from that beam. It is highly likely that any backscatter from the low intensity would not hit the FABS detector as the beams have such high angular separation.

On fig. 4.3a the thick white dashed line represents the SRS threshold parameter calculated by the group that carried out and analysed the experiment [119]. They used the threshold provided by Montgomery [8], which was initially derived by Liu [10] for substantial amplification based on Rosenbluth’s convective gain expression for backscatter:

$$\frac{\gamma_0^2}{v_{g1}v_{g2}k'} > 1. \quad (4.3)$$

Liu shows that this is equivalent to $(\frac{v_{osc}}{c})^2 k_0 L_n > 1$, whereas Montgomery gives it as $(\frac{v_{osc}}{c})^2 k_0 L_n > 2$ and provides it in experimental units as

$$I_0[\text{W}/\text{cm}^2] > \frac{4 \times 10^{17}}{L_n[\mu\text{m}]\lambda_0[\mu\text{m}]}, \quad (4.4)$$

which is the expression used in Scott *et al.*’s analysis. It is thus likely that this threshold, if based off of Liu’s expression should be a factor of 2 smaller. In the work by Scott, they only plot the SRS threshold parameter during the high intensity beams. The threshold parameter is significantly below 1 by their calculations, but making some changes to the expression shows the vagueness of the parameter and that it can actually be surpassed in these conditions with the low intensity beams. Their radiation-hydrodynamics simulations returned an increasing density scale length during the heater beams which rises to a maximum of 500 μm by 1.5 ns when they turn off. The distribution of density scale lengths at the quarter critical surface seen by individual rays in the simulation reveals a wide spread of




Figure 4.10: SRS threshold parameter as formulated by Scott *et al.* for the two sets of beams and the scale lengths they considered using eq. (4.4). The alternative threshold I provide for the low intensity beams uses Liu’s expression which is a factor of two smaller than eq. (4.4) and a wider range of scale lengths based on the distribution of values seen in the ray tracer. The intensities used are the maximum overlapped intensities of 0.9 and $2.5 \times 10^{15} \text{ W/cm}^2$ for the low and high beams respectively.

values. The plasma protrudes from the target in an elliptical shape, with the target normal along the major axis of the ellipse. The two lower intensity cones of beams have a shorter average scale length than the high intensity beams as they are further from the target normal. However, the low intensity beams are at high angles and see a much wider distribution of scale lengths with 13% of the rays in each beam seeing a scale length larger than $500 \mu\text{m}$. LPI grow unstably from feedback loops, so while average parameters are important, it is at these extrema of parameter distributions where LPI growth can be pushed over threshold and grow rapidly between a fraction of the rays. This is similar to how the average intensity of a laser beam may be below threshold, but the inhomogeneity in its intensity profile leads to speckles of high intensity that seed growth.

Figure 4.10 shows an estimate at recreating the threshold plotted by Scott *et al.* on fig. 4.3a and their sub-threshold parameter found during the low beams. These use the maximum overlapped intensities of the low and high beams but are limited in the range of scale lengths considered. An alternative threshold for the low beams is plotted in green, this time using Liu’s threshold which is half the value of Montgomery’s in eq. 4.4 and considering a wider range of scale lengths to reflect the broad distribution seen by rays in the high angle beams. For rays with scale length above 600 nm the threshold for convective backscatter is exceeded at $9 \times 10^{14} \text{ W/cm}^2$ so according to this interpretation it is plausible that Raman backscatter could grow during the low beams.

Above it is shown that there is an arbitrary factor of two between the thresholds given by Liu and Montgomery. The definition given by Liu of “substantial amplification” leaves space for interpretation as to which values of the convective gain exponent constitute this [10]. Experimental work by Rosenberg *et al.* [73] investigated thresholds for the onset of Raman back and sidescattering driven by low and high angle beams. They find intensity thresholds always below $6 \times 10^{14} \text{ W/cm}^2$ in every scenario, further suggesting that SRS could occur during the low intensity beams.

4.3.2 Two-Plasmon Decay

Light scattered from the two-plasmon decay EPWs is seen in fig. 4.3a only during times when the low intensity beams are turned on. Notably, it peaks in amplitude at 0.5 ns, when the low beams reach their maximum, but then decays despite the low intensity beams remaining at their maximum intensity. The angle of each EPW relative to the pump is found by Chiron giving a cone of possible wavevectors. The specific orientation of the EPW is in the same plane of polarisation as the pump wave. Chiron currently does not include the laser polarisation, so, in combination with the absence of tracing of the scattered rays, it becomes difficult to predict where scattered light from TPD is expected to be emitted and detected. Therefore, despite no TPD signal being seen experimentally during the high intensity beams, the possibility of it occurring can not be dismissed. Similarly, during the low intensity beams, the model can not predict which set of beams (48° or 62°) is generating the TPD scattered light.

The main takeaway from the simulation regarding TPD is that Chiron predicts TPD to be occurring at all times in both the three beam and all beam simulation, despite only limited signal on the FABS detector, implying that much of the TPD generated light is going undetected.

The amplitude of the TPD scattered light signal is also not something that is currently modelled by Chiron. Section 3.1.8 explains how the amplitude of the EPW is used to estimate the amplitude of scattered light generated by TPD. The relation between the EPW and light amplitudes is 1:1, i.e. the scattered light waves are given identical power to the EPW. The intent here was not to discuss the absolute magnitude of the scattered light, but the relative magnitude of scattered light generated by different beams.

The following analysis of scattered light from TPD in the two simulations is thus ten-




Figure 4.11: Power of scattered light generated by TPD from (a) 62° , (b) 48° and (c) the high intensity 23° degree beams in the three beam simulation. Collisional absorption was applied to the light rays. The initial powers of the scattered rays were assumed to have powers equal to the EPWs they scatter off, as is explained in Sec. 3.1.8.

uous because it is expected that the relation between EPW amplitude and scattered light amplitude is complicated by the multitude of mechanisms that could be generating the light. Given the relation between EPW amplitude and scattered light amplitude has been assumed as linear here, the following figures showing light spectra and analysis of them could instead be used to just discuss the spectra of EPWs that TPD has generated, removing some of the issues with the discussion.

4.3.2.1 Three beam simulation

Figure 4.6a showed that in the simulation that used one beam to represent each cone of beams at constant polar angle, the fraction of laser power being absorbed by TPD started at 2% at 0.4 ns, before climbing to approximately 6% absorption from 0.8 ns to the end apart from a spike at 1.6 ns.

Figure 4.11 shows the scattered light spectra from the three simulated beams. For the 48° beam in fig. 4.11b, the lower wavelength lobe of scattered light has its maximum at 0.6 ns and 0.8 ns before a decrease in amplitude at 1.0 ns. The lower wavelength lobe of scattered light on the experimental FABS data bears some similarity to this as it decays in amplitude, however its maximum is earlier at 5 ns and it has already started decaying by 0.8 ns. The lower wavelength lobe from the 62° beam in fig. 4.11a does however show a maximum at 0.4-0.6 ns that decays later in time. This similarity in shape suggests that the scattered light detected by the FABS detector may be generated by the higher angle laser beam, although more in depth modelling with tracing of the scattered rays would be required to evaluate this conclusively.

The magnitude of spectrally integrated scattered light from the 62° beam in fig. 4.11a is significantly lower than from the 48° beam fig. 4.11b despite having the same laser profile. Plotting the energy absorption by TPD in each beam shows that the 48° beam does have a higher fraction of energy absorbed by TPD by a few percent — not enough to create the massive difference in scattered power seen in the simulation. Plotting the two spectra without the effects of collisional absorption in figs. 4.13a and 4.13b shows that the total magnitude of scattered light from each beam is of the same order of magnitude, with the small difference in absorption fraction accounted for. The high angle rays are propagating longer distances and being more heavily collisionally absorbed.

The 23° beam has far less scattered power than the 48° beam in as is demonstrated in




Figure 4.12: Wavelength integrated spectra for TPD scattered light from each laser beam from the spectra presented in fig. 4.11. The solid and dashed lines are from the shorter and longer wavelength scattered light lobes respectively.

figs. 4.11 and 4.12. When collisional absorption is ignored in fig. 4.13, it still has less power in TPD scattered light than both of the low intensity beams. However, in fig. 4.6a, the laser absorption fraction by TPD after 1 ns is always equal to or higher than the fraction absorbed by TPD before 1 ns. This is because the high intensity beam has half of the power of the each of the low intensity beams (its higher intensity is by virtue of different phase plates creating a smaller focal spot). So the lower scattered power during the 23° beam is a result of lower power in the pump; proportionally more laser energy is being absorbed by TPD when the 23° is on.

4.3.2.2 All beam simulation

The simulation using all of the beams had a large difference in the amount of laser energy being absorbed by TPD between the low and high intensity beams, as depicted in fig. 4.6b. During the low beams, less than 2% of the laser power is absorbed by TPD, whereas during the high beams, over 8% of laser power is lost to TPD.

The three plots in fig. 4.14 show the summed TPD spectra from each of the five beams that constitutes each cone of constant polar angle. In this simulation, the 23° beams do produce more scattered light as a result of a much larger amount of TPD occurring during these beams than the low intensity beams.

Unlike with the SRS example above, the individual beams in each cone do not create near identical spectra. Figure 4.15 shows the integrated spectrum of the lower wavelength




Figure 4.13: Power of scattered light generated by TPD from (a) 62° , (b) 48° and (c) the high intensity 23° degree beams in the three beam simulation. This is from the same simulation as in fig. 4.11 except it ignores collisional absorption on the scattered light.




Figure 4.14: Total TPD scattered light from the all beam simulation created by summing each of the five beams in (a) the 62° cone of beams, (b) the 48° cone and (c) the 23° cone.

scattered light lobe for all of the low intensity beams. None of the ten low intensity beams generates a spectrum that peaks at 0.4 and 0.6 ns as the FABS data does, while the 62° beam does in the 3 beam simulation. All of the low intensity beams here have a large increase in scattered light amplitude between 0.4 and 0.6 ns. Four out of the five beams in the 62° cone do also decay in amplitude after this peak at 0.6 ns, as is seen in the FABS data. The 48° degree beams all maintain their amplitude instead. This evolution of magnitude of the scattered light amplitude is similar to that seen in the three beam simulation. The trend of the amplitude of the TPD generated by the 62° beams appears closer to the the experimental TPD scattered light in fig. 4.3a.




Figure 4.15: Wavelength integrated spectra for TPD scattered light from each laser beam from the spectra presented in fig. 4.14. This light is only from the lower wavelength lobe of scattered light (<700 nm).




Figure 4.16: TPD scattered light from one of the five (a) 62° beams and (b) 48° beams from the all beam simulation.

4.3.3 Excessive TPD Depletion

Figure 4.16 shows the spectra of two beams which have steep maxima that are separated in wavelength and time. In fig. 4.16a, the maxima in scattered light at 0.8 and 1.0 ns are an order of magnitude larger than the scattered light amplitude at the neighbouring wavelengths. Figure 4.16b has three local maxima at 0.8 ns and two at 1.0 ns. This feature seems unlikely to be physically realistic as the plasma parameters which determine LPI growth are not expected to vary so drastically over such short distances. Both of these examples indicate the TPD model may be being dominated by fluctuations and noise in the hydrodynamic mesh parameters that are causing spikes in LPI gain.

Examining the growth of scattered rays and depletion of pump rays in the simulations reveals that it is fluctuations in plasma and ray parameters that are causing spikes in gain. Each simulation presented in this results section (the 9 timesteps in two different beam setups) failed to have the power mesh completely converge as described at the end of sec. 3.2.2. Typically, when comparing the power in each cell between two iterations most cells do converge after a few iterations, however there are always a few cells in the mesh which fail to converge. In these cells there is normally one or multiple pump rays being entirely depleted by high LPI growth, which the method fails to resolve.

Focusing in one of the cells which failed to converge in the 15 beam simulation, figure 4.17 presents parameters along a single pump ray that prevented that cell from converging. Figure 4.17a shows the pump profile after the 9th and 10th iteration over the LPI calculation and the final averaged pump profile that was settled on. At densities below $0.22 n_c$ the pump is depleted gradually by a combination of IB, SRS and TPD and the profile has clearly converged as it is identical after the 9th and 10th iteration. However, just after $0.22 n_c$, the profile after the ninth iteration (and any odd iteration) is entirely depleted instantaneously. The model can not handle such a steep depletion and so fails to converge, alternating between the two profiles seen in iterations 9 and 10.

This depletion is caused by one pair of TPD daughter waves growing to a massive amplitude, much bigger than the TPD EPWs that have preceded them along the same pump ray. The EPWs that spiked in power were saturated via LDI, so figure 4.17b plots various wave amplitudes to explain how the spike was created. Equation (3.14) gives the expression for finding the amplitude at which LDI saturates an EPW. The electric field amplitude that LDI saturates the forward pointing EPWs along the pump ray is the blue line on the plot. Notably, it is fairly smooth around $0.22 n_c$ where the pump depletes. This LDI saturated amplitude is first converted into the intensity of the other EPW by conserving action to give I_{epw1} . Then its power is found using $P = \frac{IV}{dl}$, where dl is the length of the ray in a cell with volume V . The intensity profile (orange) also shows relatively gradual changes with density, but the power of the EPW saturated by LDI (green) has an increase of many orders of magnitude after $0.21 n_c$. This problem arises from the conversion between power and intensity using the parameters of cell volume and ray path length.

The red line shows the EPW power after considering both convective and LDI saturation. At $0.21 n_c$ it is not affected by the spike in the LDI saturation as the ray saturates convectively

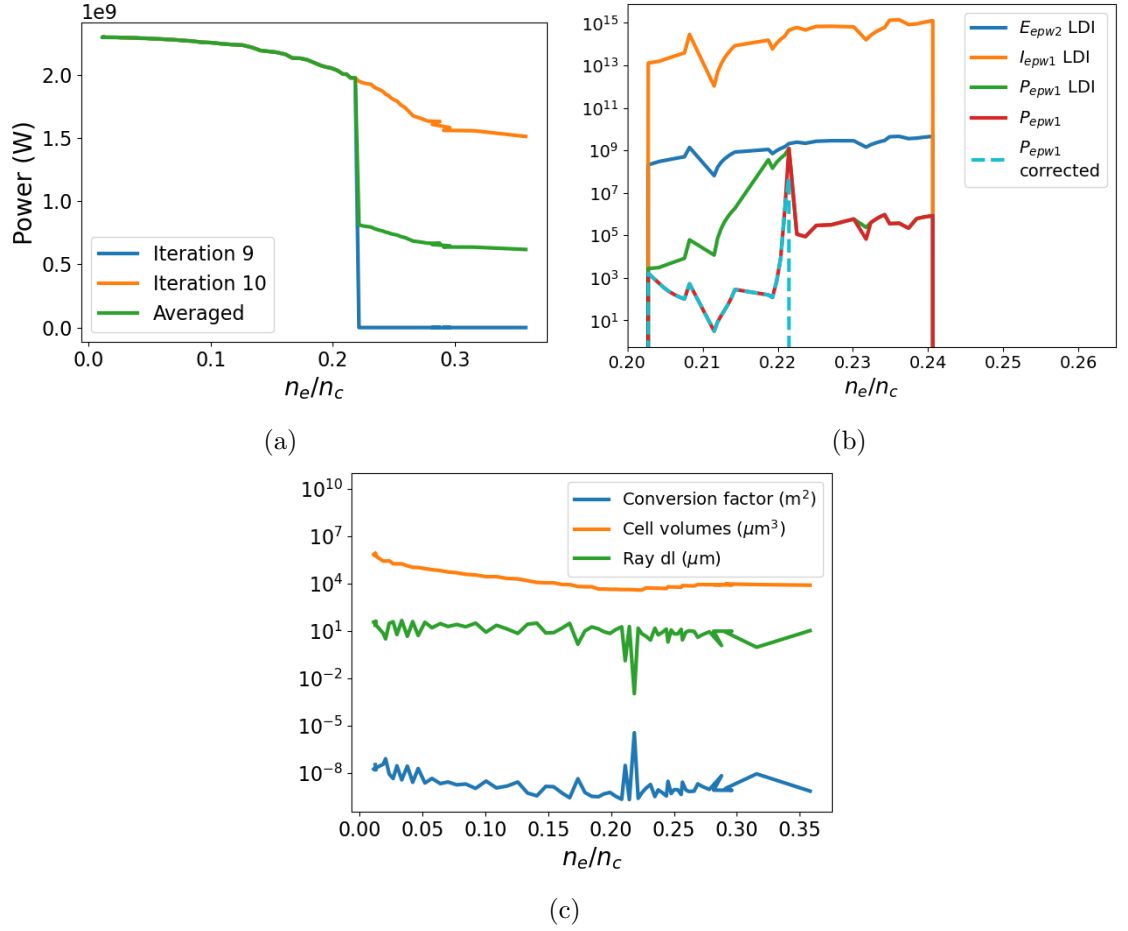


Figure 4.17: (a) Daughter wave amplitudes for TPD occurring along the single pump ray that contributes most to the largest spike at 0.8 ns in the beam presented in fig. 4.16b. LDI indicates these amplitudes are only considering saturation via LDI, whereas those without it have also considered convective saturation. (b) The conversion factor and the two contributing variables as defined in the text to move between intensity and power for the pump ray considered in (a).

here. When convective saturation fails at $0.22 n_c$, the EPW power increases drastically as LDI saturation takes over. The corrected power in cyan shows the power after considering the fact the pump is fully depleted. Note that the maximum amplitude of the corrected power is lower than the red line as energy is conserved and the EPW can only grow this large when fully depleting the pump.

To further understand why the conversion between intensity and power creates the problem here, fig. 4.17c plots the cell volumes, ray path lengths and the conversion factor of $\frac{V}{dl}$. While the volumes decrease gradually as expected as the ray propagates into the mesh, the problem is clearly created by the length of the ray in each cell. When the ray spends a very short distance in one cell it causes the conversion factor to spike. This is similar to the problem that was encountered in sec. 3.3.1 and depicted in fig. 3.20b. The ray path length has a large trough at $0.21 n_c$ and another at $0.22 n_c$, which cause the increase in the LDI saturated EPW power (green) in fig. 4.17b.

Converting between power and intensity is a repeated feature of the model as it works with ray powers most of the time but needs intensities for LPI growth rates and calculates LDI saturation in intensity too. Using powers was first implemented to get around issues created by the mesh in the intensity calculations, however the conversions between power and intensity still clearly cause problems as they rely on parameters from the mesh. In the future, moving to the neighbouring ray method (discussed in Sec.3.2.2) to calculate ray intensities should reduce the noise that is causing the problems here.

4.3.4 SBS

The FABS image in fig. 4.3b shows the scattered light in the wavelength range relevant to SBS. During the low intensity beams the spectrum is dominated by reflected laser light so that if SBS was happening here it would be indistinguishable from the reflected light. However, as a primarily backscattering instability it is likely that any SBS happening during the low beams would not be detected by the FABS detector which is in the port of one of the high intensity beams.

Comparing the fraction of laser power being absorbed in the three beam and the all beam simulation in fig. 4.18, reveals qualitative similarities between the simulations. Both see the largest amount of absorption by SBS earlier in time during when the low intensity beams are on. The absorbed fraction decreases throughout the duration of the low intensity beams




Figure 4.18: Fraction of laser energy absorbed by SBS in the two simulations

and then is very low throughout the high intensity beams apart from at 2.2 ns where it rises again. The total magnitude of SBS in both simulations is low, in particular the all beam simulation never sees SBS absorb more than 1% of laser energy. The three beam simulation sees 1-1.5% of laser energy absorbed, except for the dip during the high intensity beams. The three beam simulation had higher laser intensities so higher SBS levels are to be expected.

Both see a significant increase in absorption by SBS after 1.0 ns, i.e. when the high intensity beams turn on. Between 1.4 and 1.8 ns both simulations see big increases in the amount of SBS absorption. This continues in the three beam simulation after 1.8 ns, but in the all beam simulation, SBS declines. SBS is exclusively saturated convectively in the simulation, so the amplitude of the scattered waves will vary exponentially with changes in pump intensity. Hence the three beam simulation, which has larger pump intensities, has consistently and significantly higher SBS levels than the all beam simulation.

The decrease in SBS levels that both simulations see during the high intensity beams may be caused by the same issue that is led to the increase in laser energy that was not absorbed at this time. The error in the simulation that caused some rays to immediately exit the plasma as soon as they reached their turning point rather than propagating out through the plasma. This region near the turning point of the rays of the high intensity 23% beam will likely have been at densities where SBS may have taken place.

Figure 4.19a shows the wavelength integrated spectrum from the FABS data alongside the pulse shapes. During the high intensity beams there is a small local maximum as the beams ramp up in power. Then there are two larger local maxima while the beam is at peak intensity.




Figure 4.19: (a) Spectrally integrated signal from the experimental FABS data in fig. 4.3b. (b) Spectrally integrated scattered light from SBS in the three beam simulation.

Figure 4.19b is the spectrum from the three beam simulation. The temporal evolution of the total Brillouin scattered light from the low intensity 48° and 62° beams does not align particularly well with the FABS signal at early times in Fig. 4.19a. It also seems unlikely that Brillouin backscatter would from these beams would be measured by the FABS that was behind one of the 23° beams. The backscatter has near identical frequency to the pump so should follow roughly the same path as the pump ray, so most of it should end up around the beam port. It is possible this early time emission is instead from laser light that has propagated through the plasma without being absorbed, which Fig. 4.4a indicates is prevalent at early times. This early time emission from the experiment also seems temporally and spectrally similar to what Froula [121] describes as direct reflection of the wings of the laser off of the expanding plasma. This seems the most likely explanation for it.

The simulated SBS emission during the high intensity beam shows a spike at 2.2 ns, which does match the final maxima of the experimental SBS data in Fig. 4.19a. Before that however, there is minimal SBS occurring in the simulation, unlike the experiment which has other peaks in emission during the high intensity beam. The lack of SBS before 2.2 ns during the high intensity beam needs to be further investigated as it is likely an issue as explained above with the hydrodynamic simulation or the ray tracer. The plasma parameters and other LPI behaviour, does not change significantly between this period and at 2.2 ns and so it seems likely that the reason SBS in the high intensity beam suddenly turns on at 2.2 ns is due to a numerical error.

Chapter 5

Conclusions

This project began with the aim to develop an LPI model based on analytic theory that is computationally efficient so that it can be combined with with radiation-hydrodynamics codes and further ICF simulation capabilities.

Chapter 2 described the analytic theory used to create the model. While Chiron is mostly based off of existing theory, a derivation of Rosenbluth's [5] expression for convective gain was adapted to create an alternate expression, titled here as the "bounded" gain. This expression limited the amplification of the shortest wavelength Raman scattered rays, which was a common problem faced when using the Rosenbluth gain. It also changed the integral to begin at zero (the LPI seed position) rather than $-\infty$. The correct choice of integral lower bound will be decided once there is a self-consistent, physics-based model for seed amplitude.

A common problem in the development of steady-state models that assume LPI saturation is how to handle absolute instability growth. An expression for instability saturation via secondary wave decays has been derived for this model by balancing the growth rates of the laser-plasma instability and the plasma wave decay. In Chiron it is used to saturate both SRS and TPD by considering the decay of EPWs via the Langmuir decay instability. In reality there are many other potentially more dominant, competing effects which can nonlinearly saturate absolute instabilities, such as mode coupling and density profile modification. These effects need to be considered too for the model to be more complete. The LDI saturated amplitude is linearly proportional to pump intensity, whereas the convectively saturated amplitude is exponentially dependent on pump intensity. This pair of relations is similar to what was observed in experimental work by Rosenberg *et al.* [73], although it should be noted that LDI may not have been saturating their absolute SRS.

Chapter 3 outlines the method that takes the theory of chapter 2 and turns it into a functioning computational model capable of simulating LPI behaviour in different conditions. A model was developed for instability growth and saturation along a single one dimensional pump ray that could later be applied to the many rays of a hydrodynamic simulation with three dimensional paths, using curvilinear abscissa.

An additional one dimensional model was developed based around eq. (2.71) too. This left out the approximation of linear gradients used by Rosenbluth to make analytic integration of an expression for convective gain possible. Instead, this expression had to be numerically integrated, considering the evolving plasma parameters seen by the scattered ray. This should have given a more accurate prediction of the saturated amplitude of a convective instability, particularly with how it captured the physics of wave dephasing, so it was compared to the model based on the analytically integrated expression in sec.3.1.5. This was a test of the applicability of the assumptions made in the analytically integrated model. The regimes where it succeeded in providing results similar to the numerically integrated gain were the intensities most appropriate for the fluid theory used to derive both models. At higher intensities, both models are expected to fail to describe the LPI behaviour which can become highly kinetic. Thus, the approximations in the analytically integrated model allow it to run at significantly reduced computational expense while still providing similar results to the more physically complete numerically integrated model.

By using an iterative approach to the LPI levels, a solution to the nonlinear problem posed by pump depletion has been reached and verified as being accurate in sec.4.1. This allowed the model to use theory that treats the pump amplitude as a constant and significantly simplify the derivations for saturated wave amplitudes.

The model combines the simultaneous growth of the three instabilities and inverse bremsstrahlung, capturing the pump depletion by one mechanism suppressing other instabilities further along a ray. The framework developed will make it simple in the future to add further absorption mechanisms if so desired. New models can be developed in 1D along a single ray and then simply added to run in tandem with the existing models.

The importance of having an LDI model to limit SRS was also shown. The model prevents unbounded amplification by SRS at the quarter critical surface when convective saturation fails. This also decreases the runtime of Chiron by requiring fewer scattered rays to be launched to accurately sample plasma parameters along a pump ray. The quarter critical

surface is an important region for LPI physics so having the beginnings of a physics model to account for SRS and TPD in this region is an important first step and can be built upon with more saturation models.

Chiron was compared to the results from a planar target experiment on OMEGA, by applying it to hydrodynamic outputs from an H2D simulation in chapter 4. The presence of LPI in the experiment was diagnosed with two FABS detectors. The data provided from the FABS detectors was not absolutely calibrated so the absolute magnitude of scattered light is unknown. Comparisons can only be made by considering relative levels of scattered light across the time and wavelength axes of each detector. It also means that the relative prevalence of SRS and SBS can not be commented upon because their scattered light is recorded by different detectors.

Two separate LPI simulations were carried at each timestep comparing two methods of setting up the laser beams. The experiment had three cones consisting of five beams at constant polar angle and varying azimuthal angle. The two simulation set ups were called: three beam, for a simulation with only three beams simulated, one to represent each cone; and all beam, where all fifteen beams were simulated. The main difference would be the intensities of rays within the plasma as the three beam simulation sums the intensity of rays of all the beams in each cone, whereas the all beam simulation simulates them separately and instead combines the intensity of individual rays using the method devised in sec. 3.2.2.

The main success from this project is the Raman scattered light signal from the high intensity beam in the three beam simulation. This scattered light showed a remarkable similarity to the experimental data, with the shape of the maxima region in wavelength and time almost identical. It also showed a drop off at longer wavelengths when collisional absorption was applied to the scattered light rays. The all beam simulation did not produce a Raman scattered light spectra that was so similar to the experiment's. This indicates that the way in which ray intensities are being combined is not quite working as anticipated. Including a way to account for different beam angles seems very important. For example, in the three beam simulation an entire cone of beams has its intensity summed to be one beam, so two beams which have an azimuthal angle of 144° between them are assumed to be driving the same instability waves, which is highly unlikely unless the polar angle is very small.

The simulations both indicate SRS is present earlier in time when the low intensity beams

are on despite there being no FABS signal then. With the FABS diagnostic being in a port behind a high intensity beam, it is very plausible that Raman backscattered light from the other beams would not have been detected.

TPD is difficult to compare to the experimental data because of the spread of angles that TPD wavevectors encompass. It is also detected by a secondary scattering process which Chiron has no way to model, so currently just assumes the intensity of scattered light is proportional to the amplitude of the EPWs. Both simulations showed TPD to be the instability the dominant instability throughout the low and high intensity beams, despite its minimal presence in experimental scattered light data. Given the difficulty with diagnosing TPD through scattered light alone, hard x-ray signal — bremsstrahlung emission generated by hot electrons stopping in the fuel — would be ideal as this is assumed to be emitted isotropically. Distinguishing between SRS and TPD generated hard x-rays could still prove a problem however.

The SBS modelling indicated that the early time emission near the laser wavelength is not generated by SBS and instead may be reflected or refracted laser light. During the high intensity beams it appears that the SBS physics is not fully captured because of an error in the ray tracer which prevented some rays from propagating past their turning point.

5.1 Future Work

An overriding feature of the comparison with experimental data is shortcomings due to limited data to compare to. The experimental data only recorded scattered light from SRS and SBS at one position in the chamber, meaning backscattered light from the low intensity beams would likely have not been detected and thus no comparison can be made to Chiron's outputs here. Having an experiment that diagnosed the hard x-ray signal would also support the analysis of TPD, and SRS to a lesser extent. So, thoroughly diagnosed experiments would be helpful for future applications of the model, however part of the problem was that many of the well diagnosed LPI experiments were not useful to use because of Chiron's shortcomings. The main problem is the lack of a CBET model when CBET is such a pervasive instability in direct drive experiments with overlapping beams. For this project, the choice of experiments to use for comparison was drastically limited by searching only for planar experiments where it could be assumed that CBET was less significant.

The other key addition to Chiron is a physics based model for the magnitude of LPI seeds. Without this instability amplitudes are dependent on seed levels, which are estimates based on other simulation work and attempts to fit simulation results to experiment. A self-consistent seed model would also help to resolve where in space the convective amplification integral should begin, currently models have the lower bound at 0 or $-\infty$.

Other physics model additions include a two-ion decay model to saturate SBS and a stimulated Raman sidescatter model. The sidescatter model could be separate to backscatter with both models assuming the exact angle at the point a scattered ray is launched. Or the method for launching rays could be rethought and a range of angles could be sampled over for an instability and the angles where rays grow fastest should reveal when back or sidescatter are prevalent. Struggles modelling SBS showed the need to use the more complete kinetic descriptions of convective amplification. It is likely a kinetic model for SRS will also be developed to use instead of the fluid model.

Including ray tracing of the scattered light rays from all instabilities is another important aim for the near future. Currently it is assumed that all Raman and Brillouin backscatter propagates directly back along the pump ray's path. It would also help to understand TPD scattered light which is emitted at a wide range of angles.

Early work carried out in moving the one dimensional models to three dimensions and using results from hydrodynamics failed to appreciate the impact of noise in the mesh and so results have been sensitive to it. Smoothing of profiles has been applied with success, but sensitivity to noise is still present and causes nonphysical spikes in gain in the TPD and SBS models most notably. Further work to address this is needed to help make results more accurate and decrease the number of iterations needed for results to converge, thus reducing computational expense. This should also remove the need for the code to switch between power and intensity for the pump and scattered waves, which has brought its own problems.

The ultimate goal of this project is to have the model Chiron coupled to the radiation-hydrodynamics code Odin. Then, LPI calculations will occur during hydro simulations of ICF implosions and, importantly, will feedback information of laser absorption to the hydro simulation, thus affecting plasma dynamics.

Despite it being part of one of the initial aims, the computational cost of the model has not been discussed in this thesis. The computational efficiency is not yet realised, because the model is written in Python and the long term goal has always been to rewrite it in

Fortran when it is incorporated into Odin. Huge speed ups are inevitable, most obviously in parallelising the LPI calculations for each individual ray which currently happen one at a time as each ray is iterated over.

These improvements will be addressed over the next two years as development of Chiron is intended to continue as part of the UK effort towards inertial confinement fusion.

Appendix A

Ion-acoustic Waves in Multi-ion Plasmas

This section presents modelling of ion-acoustic wave (IAW) parameters in a two ion plasma mostly based on theory provided by Williams *et al.* [14] with comparisons made to some theory by Vu *et al.* [55].

For a two ion plasma, both authors describe the two ion modes as being fast and slow based on which has the larger phase velocity. Both of them focus on the situation where the two ion modes have phase velocities such that

$$v_{\phi,F} > v_{th,l} > v_{\phi,S} > v_{th,h}, \quad (\text{A.1})$$

where the F and S subscripts refer to the fast and slow ion acoustic modes. The two thermal velocities are for each ion species where one species is termed h for heavy and one is subscripted l for the lighter ions. The kinetic theory first derived by Williams and that derived by Vu both require this inequality to remain true.

Williams also derives a fluid expression that depends on adiabatic indices for each ion in the plasma, which they point out can be different for each mode. In their first example, depicted in fig. A.1, they use a 99%H 1%Xe plasma where the Xe is partly ionised such that it has $Z^* = 40$. They say for the fast mode the adiabatic index is $\gamma_{Xe} = \gamma_H = 3$, but for the slow mode $\gamma_H = 1$, while $\gamma_{Xe} = 3$. The phase velocity they derive from fluid theory, eq.(16) in their paper, also has an ambiguous \pm sign. Fig. A.1 plots the results for the phase velocity for each mode using both the plus and the minus to see which better matches the




Figure A.1: Phase velocity (v_ϕ) of the fast (red) and slow (blue) ion-acoustic mode from a 99%H 1%Xe plasma. The solid lines are the results presented by Williams *et al.* [14], and the pluses and minuses are for our results using their expression with either the plus or minus sign. Velocities are normalised to $v_r = \sqrt{k_B T_e / m_p}$.

plot in Williams' paper. This shows quite clearly that the plus should be used for the fast mode and the minus for the slow mode. This is a small point to make but nice to clarify for modelling and any future work using this theory. Throughout this section velocities are normalised to $v_r = \sqrt{k_B T_e / m_p}$, where m_p is the mass of a proton.

Figure A.2a shows the two IAW phase velocities and ion thermal velocities for the XeH simulation for a fixed $k\lambda_D = 0.1$ and varying ion-electron temperature ratio. This scenario clearly obeys the condition in eq. (A.1) for the kinetic theory in both papers to be valid. Using Williams' theory, the results in his paper are well recreated. However, using Vu's theory the phase velocity of the slow mode is significantly lower than the results in Williams' paper.

Figure A.2b compares the values of damping rate presented in Williams' paper, to my results using the kinetic damping rates Williams and Vu provide for both modes. The damping rates for both modes are not well reproduced by models based on either of the theories, but their general shapes show similar trends. For the fast mode both sets of data show a steep upwards trend around $0.2T_i/T_e$ similar to the results in the paper. However from here their damping rate decreases at higher temperature ratios, unlike the results in the paper. At very low temperature ratio, Williams' fast mode theory recreates the shape in his paper, whereas Vu's theory decreases towards zero rapidly. For the slow mode both set of theory stay with a factor of ~ 3 of the results in the paper.

In cases with two ions that have much closer masses, the thermal and phase velocities

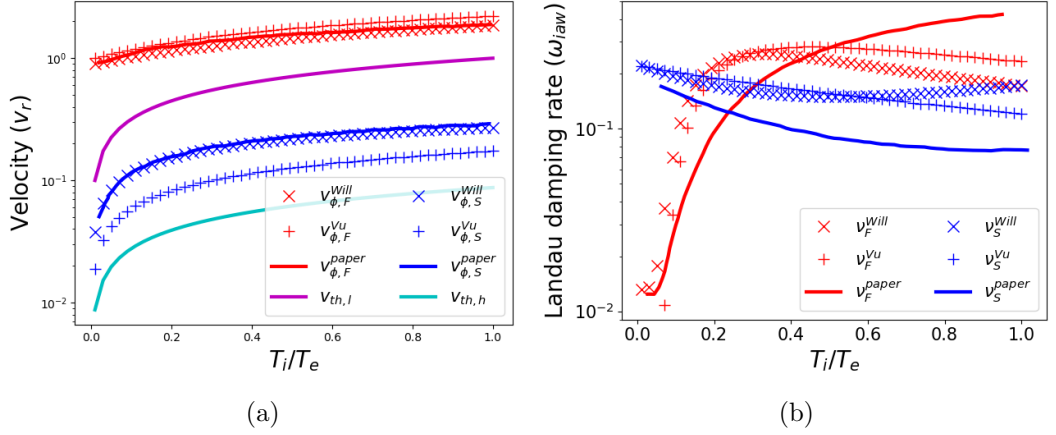


Figure A.2: (a) Phase velocity (v_ϕ) of the fast (red) and slow (blue) ion acoustic modes as well as the thermal velocities (v_{th}) of the light (magenta) and heavy (cyan) ion species. The solid lines are results taken from Williams' paper, while the crosses and pluses represent my results using the kinetic expressions of Williams and Vu respectively. Velocities are normalised to $v_r = \sqrt{k_B T_e / m_p}$. (b) Landau damping rate (ν) of the two ion modes again comparing results with Williams with my model using kinetic theory from Williams and Vu using the same legend as in (a). Damping rate is normalised to the wave's frequency.




Figure A.3: (a) The phase velocities (v_ϕ) of the two IAWs and the thermal velocities in a 50% C 50% H plasma. The phase velocity data was taken directly from Williams *et al.* [14]. (b) The Landau damping rates (ν) for the fast IAW (red) and slow IAW (blue) in a CH plasma. The data from Williams paper are the solid lines and our solutions using Williams' and Vu's theory are the crosses and pluses respectively.




Figure A.4: Phase velocities (v_ϕ) for the (a) fast and (b) slow ion mode in a CH plasma. The solid lines are the results presented in Williams paper from his new kinetic (blue) and fluid (orange) theory. Our results using Williams' and Vu's kinetic theory based on the unfulfilled condition in eq. (A.1) are the blue crosses and pluses respectively. The orange pluses are our results using Williams' fluid theory.

often do not meet the condition in eq. (A.1). Most ICF ablation plasmas will be CH, and fig. A.3a, shows the relevant thermal velocities and IAW phase velocities. Clearly the slow mode often has a phase velocity above the light ion's thermal velocity. Williams presents a new expression for the frequency and damping of ion modes based on kinetic theory using the plasma dielectric function without using the approximations they had previously used. This has not yet been solved numerically as part of this project, so instead the solutions Williams provides are compared with my solutions using the original, approximated kinetic theory.

Figure A.3b compares the damping rates between Williams' and Vu's kinetic theory that was derived in the regime not applicable to the CH plasma and the results that Williams presents from the new kinetic theory. For the fast mode, the general shape of the damping as a function of ion temperature is followed by all three sets of data, however magnitude is significantly different, especially for Williams' old theory which is a factor ~ 7 smaller than their new theory. For the slow mode Vu's theory is once again larger than Williams' approximated theory. For this mode Williams' approximated theory is closer in magnitude to the new theory, never with more than a factor 2 between them. For both modes Vu's theory is a factor ~ 4 apart from Williams' new theory for a lot of the temperature range.

Figures A.4a and A.4b compare various methods of calculating the phase velocity for the fast and slow IAW modes respectively. The fluid and kinetic models are included as Williams

does the same and there is no indication that fluid models are invalid in this regime. Williams finds fairly good agreement with their fluid model model and the new kinetic theory (which they view as the correct result) for both modes, but particularly the fast mode. For both modes, my model using Williams' old kinetic theory remains significantly closer to Williams' new theory than my model using Vu's theory. It is unclear why my model using Williams' fluid theory no longer agrees with his fluid results. This is particularly puzzling given the good agreement between fluid models in fig. A.1 for the XeH example. Based on these results this project use Williams' kinetic theory that he derives under the assumption that eq. (A.1) is true. Although it is not quite appropriate for CH plasmas, it does a better job than the fluid model. One advantage of kinetic theory is that it also provides a Landau damping rate, whereas fluid theory will not, although currently that rate is not in agreement with the paper, especially for the fast mode.

List of References

- [1] J. Myatt, R. Follett, J. Shaw, D. Edgell, D. Froula, I. Igumenshchev, and V. Goncharov. A wave-based model for cross-beam energy transfer in direct-drive inertial confinement fusion. *Physics of Plasmas*, 24(5), 2017.
- [2] D. Strozzi, E. Williams, D. Hinkel, D. Froula, R. London, and D. Callahan. Ray-based calculations of backscatter in laser fusion targets. *Physics of Plasmas*, 15(10), 2008.
- [3] A. C. Angus. *Probabilistic modelling of stimulated raman back-scatter in laser direct-drive fusion plasmas at ignition scale*. Ph.D. thesis, University of Warwick, 2023.
- [4] D. Bénisti, O. Morice, C. Rousseaux, A. Debayle, P. Masson-Laborde, and P. Loiseau. Nonlinear kinetic estimate of raman reflectivity in a hohlraum plasma with a radiation hydrodynamics code. *Physical Review E*, 109(4):L043201, 2024.
- [5] M. N. Rosenbluth. Parametric instabilities in inhomogeneous media. *Physical Review Letters*, 29(9):565, 1972.
- [6] A. Bose. *Direct drive inertial confinement fusion: analysis of the implosion core*. Ph.D. thesis, University of Rochester, 2017.
- [7] R. Betti and O. Hurricane. Inertial-confinement fusion with lasers. *Nature Physics*, 12(5):435, 2016.
- [8] D. S. Montgomery. Two decades of progress in understanding and control of laser plasma instabilities in indirect drive inertial fusion. *Physics of Plasmas*, 23(5), 2016.
- [9] A. Solodov, et al. Laser-plasma interaction experiments at direct-drive ignition-relevant plasma conditions at the national ignition facility. *Bulletin of the American Physical Society*, 62, 2017.

- [10] C. Liu, M. N. Rosenbluth, and R. B. White. Raman and Brillouin scattering of electromagnetic waves in inhomogeneous plasmas. *The Physics of Fluids*, 17(6):1211, 1974.
- [11] D. Forslund, J. Kindel, and E. Lindman. Theory of stimulated scattering processes in laser-irradiated plasmas. *The Physics of Fluids*, 18(8):1002, 1975.
- [12] P. Michel. *Introduction to laser-plasma interactions*. Springer, 2023.
- [13] L. Hao, R. Yan, J. Li, W. Liu, and C. Ren. Nonlinear fluid simulation study of stimulated raman and brillouin scatterings in shock ignition. *Physics of Plasmas*, 24(6), 2017.
- [14] E. Williams, R. Berger, R. Drake, A. Rubenchik, B. Bauer, D. Meyerhofer, A. Gaeris, and T. Johnston. The frequency and damping of ion acoustic waves in hydrocarbon (ch) and two-ion-species plasmas. *Physics of Plasmas*, 2(1):129, 1995.
- [15] IEA. Electricity 2024. <https://www.iea.org/reports/electricity-2024>, 2024. Accessed: 2024-08-18.
- [16] Ember energy institute. Yearly electricity data. <https://ember-climate.org/data-catalogue/yearly-electricity-data/>, 2024. Accessed: 2024-08-18.
- [17] T. Letcher. *Storing Energy: with Special Reference to Renewable Energy Sources*. Elsevier, 2022.
- [18] IEA. The role of critical minerals in clean energy transitions. <https://www.iea.org/reports/the-role-of-critical-minerals-in-clean-energy-transitions>, 2021. Accessed: 2024-08-18.
- [19] H. Abu-Shawareb, et al. Achievement of target gain larger than unity in an inertial fusion experiment. *Physical Review Letters*, 132(6):065102, 2024.
- [20] T. C. Sangster, et al. Cryogenic dt and d2 targets for inertial confinement fusion. *Physics of Plasmas*, 14(5), 2007.
- [21] S. E. Bodner. Rayleigh-taylor instability and laser-pellet fusion. *Physical Review Letters*, 33(13):761, 1974.
- [22] R. Craxton, et al. Direct-drive inertial confinement fusion: A review. *Physics of Plasmas*, 22(11), 2015.

- [23] J. Lindl. Development of the indirect-drive approach to inertial confinement fusion and the target physics basis for ignition and gain. *Physics of plasmas*, 2(11):3933, 1995.
- [24] H. Abu-Shawareb et al. Lawson criterion for ignition exceeded in an inertial fusion experiment. *Physical review letters*, 129(7):075001, 2022.
- [25] M. Tabak, D. Hinkel, S. Atzeni, E. M. Campbell, and K. Tanaka. Fast ignition: Overview and background. *Fusion science and technology*, 49(3):254, 2006.
- [26] R. Betti, C. Zhou, K. Anderson, L. Perkins, W. Theobald, and A. Solodov. Shock ignition of thermonuclear fuel with high areal density. *Physical review letters*, 98(15):155001, 2007.
- [27] N. Meezan, et al. Indirect drive ignition at the national ignition facility. *Plasma Physics and Controlled Fusion*, 59(1):014021, 2016.
- [28] P. Michel, et al. Symmetry tuning via controlled crossed-beam energy transfer on the national ignition facility. *Physics of Plasmas*, 17(5), 2010.
- [29] S. Atzeni and J. Meyer-ter Vehn. *The physics of inertial fusion: beam plasma interaction, hydrodynamics, hot dense matter*, volume 125. OUP Oxford, 2004.
- [30] Y. Kato, K. Mima, N. Miyanaga, S. Arinaga, Y. Kitagawa, M. Nakatsuka, and C. Yamanaka. Random phasing of high-power lasers for uniform target acceleration and plasma-instability suppression. *Physical Review Letters*, 53(11):1057, 1984.
- [31] W. Seka, R. Bahr, R. Short, A. Simon, R. Craxton, D. Montgomery, and A. Rubenchik. Nonlinear laser–matter interaction processes in long-scale-length plasmas. *Physics of Fluids B: Plasma Physics*, 4(7):2232, 1992.
- [32] S. Skupsky, R. Short, T. Kessler, R. Craxton, S. Letzring, and J. Soures. Improved laser-beam uniformity using the angular dispersion of frequency-modulated light. *Journal of Applied Physics*, 66(8):3456, 1989.
- [33] J. E. Rothenberg. Two-dimensional beam smoothing by spectral dispersion for direct-drive inertial confinement fusion. In *Solid State Lasers for Application to Inertial Confinement Fusion (ICF)*, volume 2633, pages 634–644. SPIE, 1995.

- [34] S. P. Regan, et al. Experimental investigation of smoothing by spectral dispersion. *JOSA B*, 17(9):1483, 2000.
- [35] D. Montgomery, J. Moody, H. Baldis, B. Afeyan, R. Berger, K. Estabrook, B. Lasinski, E. Williams, and C. Labaune. Effects of laser beam smoothing on stimulated raman scattering in exploding foil plasmas. *Physics of Plasmas*, 3(5):1728, 1996.
- [36] J. Moody, B. MacGowan, J. Rothenberg, R. Berger, L. Divol, S. Glenzer, R. Kirkwood, E. Williams, and P. Young. Backscatter reduction using combined spatial, temporal, and polarization beam smoothing in a long-scale-length laser plasma. *Physical review letters*, 86(13):2810, 2001.
- [37] S. Skupsky and R. Craxton. Irradiation uniformity for high-compression laser-fusion experiments. *Physics of Plasmas*, 6(5):2157, 1999.
- [38] R. Craxton and S. Skupsky. 2d ssd and polarization wedges for omega and the nif. *Bull. Am. Phys. Soc.*, 40:1826, 1995.
- [39] T. Arber, et al. Contemporary particle-in-cell approach to laser-plasma modelling. *Plasma Physics and Controlled Fusion*, 57(11):113001, 2015.
- [40] S. Spencer, A. G. Seaton, T. Goffrey, and T. D. Arber. Inflationary stimulated raman scattering in shock-ignition plasmas. *Physics of Plasmas*, 27(12), 2020.
- [41] J. F. Myatt, J. G. Shaw, R. K. Follett, D. H. Edgell, D. H. Froula, J. P. Palastro, and V. N. Goncharov. Lpse: A 3-d wave-based model of cross-beam energy transfer in laser-irradiated plasmas. *Journal of Computational Physics*, 399:108916, 2019.
- [42] LLE. Laser plasma simulation environment. <https://www.lle.rochester.edu/education/research-areas/plasma-ultrafast-science-engineering/laser-plasma-simulation-environment/>. Accessed: 2024-08-21.
- [43] R. Follett, J. Myatt, J. Shaw, D. Michel, A. Solodov, D. Edgell, B. Yaakobi, and D. Froula. Simulations and measurements of hot-electron generation driven by the multibeam two-plasmon-decay instability. *Physics of Plasmas*, 24(10), 2017.
- [44] V. N. Goncharov, et al. Improving the hot-spot pressure and demonstrating ignition hydrodynamic equivalence in cryogenic deuterium–tritium implosions on omega. *Physics of Plasmas*, 21(5), 2014.

- [45] P. Radha, et al. Direct drive: Simulations and results from the national ignition facility. *Physics of Plasmas*, 23(5), 2016.
- [46] R. Berger, C. Still, E. Williams, and A. Langdon. On the dominant and subdominant behavior of stimulated raman and brillouin scattering driven by nonuniform laser beams. *Physics of Plasmas*, 5(12):4337, 1998.
- [47] R. Town. Overview of laser-plasma interaction codes used for icf research at llnl. Technical report, Lawrence Livermore National Lab.(LLNL), Livermore, CA (United States), 2018.
- [48] A. Debayle, C. Ruyer, O. Morice, P.-E. Masson-Laborde, P. Loiseau, and D. Benisti. A unified modeling of wave mixing processes with the ray tracing method. *Physics of Plasmas*, 26(9), 2019.
- [49] C. K. Williams and C. E. Rasmussen. *Gaussian processes for machine learning*, volume 2. MIT press Cambridge, MA, 2006.
- [50] D. Bénisti, O. Morice, C. Rouseaux, A. Debayle, P. Masson-Laborde, and P. Loiseau. A predictive inline model for nonlinear stimulated raman scattering in a hohlraum plasma. *arXiv preprint arXiv:2304.11589*, 2023.
- [51] T. J. M. Boyd and J. J. Sanderson. *The physics of plasmas*. Cambridge university press, 2003.
- [52] W. Kruer. *The physics of laser plasma interactions*. crc Press, 2003.
- [53] A. Richardson. *2019 NRL plasma formulary*. Naval Research Laboratory Washington, DC, 2019.
- [54] A. N. Simakov. Electron transport in a collisional plasma with multiple ion species in the presence of a magnetic field. *Physics of Plasmas*, 29(2), 2022.
- [55] H. Vu, J. Wallace, and B. Bezzerides. An analytical and numerical investigation of ion acoustic waves in a two-ion plasma. *Physics of Plasmas*, 1(11):3542, 1994.
- [56] P. Bertrand, M. Albrecht-Marc, T. Reveille, and A. Ghizzo. Vlasov models for laser-plasma interaction. *Transport Theory and Statistical Physics*, 34(1-2):103, 2005.

- [57] T. Johnston and J. Dawson. Correct values for high-frequency power absorption by inverse bremsstrahlung in plasmas. *Physics of Fluids*, 16(5):722, 1973.
- [58] A. B. Langdon. Nonlinear inverse bremsstrahlung and heated-electron distributions. *Physical Review Letters*, 44(9):575, 1980.
- [59] D. Turnbull, et al. Inverse bremsstrahlung absorption. *Physical Review Letters*, 130(14):145103, 2023.
- [60] D. Turnbull, J. Katz, M. Sherlock, A. Milder, M. Cho, L. Divol, N. Shaffer, D. Strozzi, P. Michel, and D. Froula. Reconciling calculations and measurements of inverse bremsstrahlung absorption. *Physics of Plasmas*, 31(6), 2024.
- [61] Y. Zhao, S.-M. Weng, H.-H. Ma, X.-J. Bai, and Z.-M. Sheng. Mitigation of laser plasma parametric instabilities with broadband lasers. *Reviews of Modern Plasma Physics*, 7(1):1, 2022.
- [62] J. W. Bates, et al. Suppressing cross-beam energy transfer with broadband lasers. *High Energy Density Physics*, 36:100772, 2020.
- [63] P. Mounaix, D. Pesme, W. Rozmus, and M. Casanova. Space and time behavior of parametric instabilities for a finite pump wave duration in a bounded plasma. *Physics of Fluids B: Plasma Physics*, 5(9):3304, 1993.
- [64] M. Rosenbluth, R. White, and C. Liu. Temporal evolution of a three-wave parametric instability. *Physical Review Letters*, 31(19):1190, 1973.
- [65] K. Nishikawa. Parametric excitation of coupled waves i. general formulation. *Journal of the physical society of Japan*, 24(4):916, 1968.
- [66] R. J. Briggs et al. *Electron-stream interaction with plasmas*, volume 187. MIT press Cambridge, MA, 1964.
- [67] R. Yan, A. Maximov, and C. Ren. The linear regime of the two-plasmon decay instability in inhomogeneous plasmas. *Physics of Plasmas*, 17(5), 2010.
- [68] E. A. Williams. Convective growth of parametrically unstable modes in inhomogeneous media. *Physics of Fluids B: Plasma Physics*, 3(6):1504, 1991.

- [69] K. Estabrook and W. Kruer. Theory and simulation of one-dimensional raman backward and forward scattering. *The Physics of fluids*, 26(7):1892, 1983.
- [70] K. Estabrook, W. Kruer, and B. Lasinski. Heating by raman backscatter and forward scatter. *Physical Review Letters*, 45(17):1399, 1980.
- [71] P. Michel, L. Divol, E. Dewald, J. Milovich, M. Hohenberger, O. Jones, L. B. Hopkins, R. Berger, W. Kruer, and J. Moody. Multibeam stimulated raman scattering in inertial confinement fusion conditions. *Physical review letters*, 115(5):055003, 2015.
- [72] P. Michel, et al. Theory and measurements of convective raman side scatter in inertial confinement fusion experiments. *Physical Review E*, 99(3):033203, 2019.
- [73] M. Rosenberg, et al. Stimulated raman scattering mechanisms and scaling behavior in planar direct-drive experiments at the national ignition facility. *Physics of Plasmas*, 27(4), 2020.
- [74] S. Hironaka, J. Sivajeyan, J. Wang, M. J. Rosenberg, A. Solodov, T. Filkins, C. Xiao, Q. Wang, W. Seka, and J. F. Myatt. Identification of stimulated raman side scattering in near-spherical coronal plasmas on omega ep. *Physics of Plasmas*, 30(2), 2023.
- [75] J. Drake and Y. Lee. Temporally growing raman backscattering instabilities in an inhomogeneous plasma. *Physical Review Letters*, 31(19):1197, 1973.
- [76] B. Afeyan and E. Williams. Stimulated raman sidescattering with the effects of oblique incidence. *The Physics of fluids*, 28(11):3397, 1985.
- [77] B. B. Afeyan and E. A. Williams. A variational approach to parametric instabilities in inhomogeneous plasmas i: Two model problems. *Physics of Plasmas*, 4(11):3788, 1997.
- [78] H. Baldis, D. Villeneuve, C. Labaune, D. Pesme, W. Rozmus, W. Kruer, and P. Young. Coexistence of stimulated raman and brillouin scattering in laser-produced plasmas. *Physics of Fluids B: Plasma Physics*, 3(8):2341, 1991.
- [79] R. Drake and S. Batha. The influence of subsidiary langmuir decay on the spectrum of stimulated raman scattering. *Physics of Fluids B: Plasma Physics*, 3(11):2936, 1991.

- [80] I. Barth and P. Michel. The raman gap and collisional absorption. *Physics of Plasmas*, 31(1), 2024.
- [81] K. Tanaka and L. M. Goldman. Observations of brillouin sidescatter in laser-produced plasmas. *Physical Review Letters*, 45(19):1558, 1980.
- [82] D. Turnbull, P. Michel, J. Ralph, L. Divol, J. Ross, L. B. Hopkins, A. Kritcher, D. Hinkel, and J. Moody. Multibeam seeded brillouin sidescatter in inertial confinement fusion experiments. *Physical Review Letters*, 114(12):125001, 2015.
- [83] B. D. Fried and S. D. Conte. *The plasma dispersion function: the Hilbert transform of the Gaussian*. Academic press, 1961.
- [84] W. L. Kruer, S. C. Wilks, B. B. Afeyan, and R. K. Kirkwood. Energy transfer between crossing laser beams. *Physics of Plasmas*, 3(1):382, 1996.
- [85] I. Igumenshchev, et al. Crossed-beam energy transfer in direct-drive implosions. *Physics of Plasmas*, 19(5), 2012.
- [86] W. Seka, B. Afeyan, R. Boni, L. Goldman, R. Short, K. Tanaka, and T. Johnston. Diagnostic value of odd-integer half-harmonic emission from laser-produced plasmas. *The Physics of fluids*, 28(8):2570, 1985.
- [87] H. Baldis and C. Walsh. Growth and saturation of the two-plasmon decay instability. *The Physics of Fluids*, 26(5):1364, 1983.
- [88] N. Ebrahim, H. Baldis, C. Joshi, and R. Benesch. Hot electron generation by the two-plasmon decay instability in the laser-plasma interaction at 10.6 μm . *Physical Review Letters*, 45(14):1179, 1980.
- [89] W. Seka, D. Edgell, J. Myatt, A. Maximov, R. Short, V. Goncharov, and H. Baldis. Two-plasmon-decay instability in direct-drive inertial confinement fusion experiments. *Physics of Plasmas*, 16(5), 2009.
- [90] R. K. Follett. *The multiple-beam two-plasmon-decay instability*. University of Rochester, 2016.

- [91] W. Seka, J. Myatt, R. Short, D. Froula, J. Katz, V. Goncharov, and I. Igumenshchev. Nonuniformly driven two-plasmon-decay instability in direct-drive implosions. *Physical review letters*, 112(14):145001, 2014.
- [92] A. G. Seaton. *Particle-in-cell simulations of laser-plasma instabilities in shock-ignition*. Ph.D. thesis, University of Warwick, 2019.
- [93] C. Liu and M. N. Rosenbluth. Parametric decay of electromagnetic waves into two plasmons and its consequences. *The Physics of Fluids*, 19(7):967, 1976.
- [94] R. Yan, A. Maximov, C. Ren, and F. Tsung. Growth and saturation of convective modes of the two-plasmon decay instability in inertial confinement fusion. *Physical review letters*, 103(17):175002, 2009.
- [95] C. Liu, V. Tripathi, and B. Eliasson. *High-power laser-plasma interaction*. Cambridge University Press, 2019.
- [96] J. Zhang, J. Myatt, R. Short, A. Maximov, H. Vu, D. DuBois, and D. Russell. Multiple beam two-plasmon decay: linear threshold to nonlinear saturation in three dimensions. *Physical Review Letters*, 113(10):105001, 2014.
- [97] D. Russell and D. DuBois. $3/2 \omega$ radiation from the laser-driven two-plasmon decay instability in an inhomogeneous plasma. *Physical review letters*, 86(3):428, 2001.
- [98] D. DuBois, D. Russell, and H. A. Rose. Saturation spectra of the two-plasmon decay instability. *Physical review letters*, 74(20):3983, 1995.
- [99] H. Vu, D. DuBois, J. Myatt, and D. Russell. Hot-electron production and suprathermal heat flux scaling with laser intensity from the two-plasmon-decay instability. *Physics of Plasmas*, 19(10), 2012.
- [100] J. Meyer. Mode coupling of the two-plasmon decay instability to ion-acoustic waves and the effect on $(3/2)$ -harmonic emission. *Physics of Fluids B: Plasma Physics*, 4(9):2934, 1992.
- [101] C. Riconda and S. Weber. Raman-brillouin interplay for inertial confinement fusion relevant laser-plasma interaction. *High Power Laser Science and Engineering*, 4:e23, 2016.

- [102] H. Vu, D. DuBois, and B. Bezzerides. Kinetic inflation of stimulated raman backscatter in regimes of high linear landau damping. *Physics of Plasmas*, 9(5):1745, 2002.
- [103] R. Scott, D. Barlow, W. Trickey, A. Ruocco, K. Glize, L. Antonelli, M. Khan, and N. Woolsey. Shock-augmented ignition approach to laser inertial fusion. *Physical Review Letters*, 129(19):195001, 2022.
- [104] D. Strozzi. *Vlasov simulations of kinetic enhancement of Raman backscatter in laser fusion plasmas*. Ph.D. thesis, MIT Plasma Science and Fusion Center, 2005.
- [105] J. T. Larsen and S. M. Lane. Hyades—a plasma hydrodynamics code for dense plasma studies. *Journal of Quantitative Spectroscopy and Radiative Transfer*, 51(1-2):179, 1994.
- [106] J. Colvin and J. Larsen. *Extreme physics: properties and behavior of matter at extreme conditions*, chapter 4. Cambridge University Press, 2013.
- [107] J. I. Castor. *Radiation hydrodynamics*, chapter 3. Cambridge University Press, 2004.
- [108] J. Colvin and J. Larsen. *Extreme physics: properties and behavior of matter at extreme conditions*, chapter 11. Cambridge University Press, 2013.
- [109] J. I. Castor. *Radiation hydrodynamics*, chapter 2. Cambridge University Press, 2004.
- [110] C. W. Hirt, A. A. Amsden, and J. Cook. An arbitrary lagrangian-eulerian computing method for all flow speeds. *Journal of computational physics*, 14(3):227, 1974.
- [111] D. Edgell, R. Follett, I. Igumenshchev, J. Myatt, J. Shaw, and D. Froula. Mitigation of cross-beam energy transfer in symmetric implosions on omega using wavelength detuning. *Physics of Plasmas*, 24(6), 2017.
- [112] D. DuBois, B. Bezzerides, and H. A. Rose. Collective parametric instabilities of many overlapping laser beams with finite bandwidth. *Physics of Fluids B: Plasma Physics*, 4(1):241, 1992.
- [113] R. Follett, J. Shaw, J. Myatt, D. Froula, and J. Palastro. Multibeam absolute stimulated raman scattering and two-plasmon decay. *Physical Review E*, 101(4):043214, 2020.

- [114] G. Cristoforetti, et al. Time evolution of stimulated raman scattering and two-plasmon decay at laser intensities relevant for shock ignition in a hot plasma. *High Power Laser Science and Engineering*, 7:e51, 2019.
- [115] G. Cristoforetti, et al. Observation and modelling of stimulated raman scattering driven by an optically smoothed laser beam in experimental conditions relevant for shock ignition. *High Power Laser Science and Engineering*, 9:e60, 2021.
- [116] M. M. Škorić, M. S. Jovanović, and M. Rajković. Transition to turbulence via spatiotemporal intermittency in stimulated raman backscattering. *Physical Review E*, 53(4):4056, 1996.
- [117] L. Yin, B. J. Albright, H. Rose, K. Bowers, B. Bergen, D. Montgomery, J. L. Kline, and J. C. Fernández. Onset and saturation of backward stimulated raman scattering of laser in trapping regime in three spatial dimensions. *Physics of Plasmas*, 16(11), 2009.
- [118] C. Tang. Saturation and spectral characteristics of the stokes emission in the stimulated brillouin process. *Journal of Applied Physics*, 37(8):2945, 1966.
- [119] R. H. Scott, et al. Shock ignition laser-plasma interactions in ignition-scale plasmas. *Physical review letters*, 127(6):065001, 2021.
- [120] S. Depierreux, et al. Experimental investigation of the collective stimulated brillouin and raman scattering of multiple laser beams in inertial confinement fusion experiments. *Plasma Physics and Controlled Fusion*, 62(1):014024, 2019.
- [121] D. Froula, et al. Laser-plasma interactions in direct-drive ignition plasmas. *Plasma Physics and Controlled Fusion*, 54(12):124016, 2012.
- [122] T. Dewandre, J. R. Albritton, and E. Williams. Doppler shift of laser light reflected from expanding plasmas. *The Physics of Fluids*, 24(3):528, 1981.
- [123] B. Bennett, J. Johnson, G. Kerley, and G. Rood. Recent developments in the sesame equation-of-state library. Technical report, Los Alamos National Lab.(LANL), Los Alamos, NM (United States), 1978.