

Department of Physics

**Numerical models for the simulation
of Raman amplification in plasma**

by

John Farmer

Thesis submitted for the degree of

Doctor of Philosophy in Physics

2012

This thesis is the result of the author's original research. It has been composed by the author and has not been previously submitted for examination which has led to the award of a degree.

The copyright of this thesis belongs to the author under the terms of the United Kingdom Copyright Acts as qualified by University of Strathclyde Regulation 3.50. Due acknowledgement must always be made of the use of any material contained in, or derived from, this thesis.

To the thousand other things

I could have done.

Abstract

Raman amplification in plasma is a potential method of producing ultra-short, ultra-intense laser pulses. As amplification in realistic systems involves many competing nonlinear processes, numerical simulations are often used to better understand the associated complex behaviour. Different models offer different advantages and limitations. This work presents a study of three distinct numerical models: a particle-in-cell (PIC) code, a three-wave model derived from a fluid treatment of plasma, and the envelope-PIC code aPIC.

PIC codes are known to incur large computational overheads. A new numerical instability associated with nonlinear plasma waves has been identified, which can result in a significant increase in the computational overhead required to achieve convergent results for the simulation of Raman amplification. The absence of collisional damping in available PIC codes is found to limit the parameter spaces within which simulations can be carried out.

The conventional three-wave model has been extended to improve its accuracy, and to include the effects of plasma wave har-

monics. In warm plasma, a new model is developed to estimate the fraction of trapped electrons, and their influence on the effective modulation depth of the plasma wave. Benchmarking against PIC codes shows strong agreement in regimes where the linearisation of the plasma response and Landau damping rates is valid.

aPIC, an envelope-PIC implementation suitable for the investigation of Raman amplification, has been further developed to give improved numerical accuracy, allow the use of chirped pulses, and approximate the dispersive action of the plasma on the laser pulses. Benchmarking against full PIC codes shows good agreement. The code is used to identify a new regime of broad-bandwidth Raman amplification, with nonlinearity arising through high-amplitude plasma waves. Possible future improvements to the model and numerical implementation are outlined.

The applicability of these models to the simulation of current experimental work at the University of Strathclyde is also considered.

Acknowledgements

My sincere thanks to Bernhard Ersfeld, for a great many useful and enjoyable discussions, for his encouragement, and for pointing out when I'd made a mistake or a questionable assumption. My thanks also to Dino Jaroszynski, for his support, and for giving me the freedom to follow my initiative, while reminding me of the importance of keeping my work relevant. I am also grateful to Adam Noble, who provided feedback on many of my ideas, especially during the development of the trapping model used in Chapter 4.

I am grateful to Gennady Shvets, who hosted me during my first collaboration trip, for providing not only some answers, but some interesting questions, and to Luis Silva and the group at IST, who provided the OSIRIS code and instruction on how to use it. My thanks to Alan Cairns, with whom I worked closely on the development of a model for Raman amplification with a broad-bandwidth pump; although the work was not completed for inclusion in this thesis, it should in time provide some interesting publications. I am also grateful to MinSup Hur, who generously hosted me in Korea, and provided the aPIC code and support in further developing it.

Thanks also to Gregory Vieux and Xue Yang, who were always happy to answer my questions on their experimental work, and to Silvia Cipiccia, who was always available for a chat, about physics or otherwise. I appreciate

the efforts of the entire group at Strathclyde, who provided an enjoyable atmosphere in which to conduct my PhD.

On a personal note, I would like to thank Maria Pia, Ania, Adam and Silvia, for pleasant distractions and adjournments to the pub, and Mireille, who was always very generous with her praise for my baking, as well as the friends made through collaborations, conferences and meetings. Thanks also to Mick and Sandra for some fun days out, Ian for a great many hours spent climbing, and Chau for the best nights out in Glasgow. Thanks also to Naomi and my parents for their support, and for putting up with too many conversations influenced by physics.

The role of the author

The particle-in-cell code OSIRIS, which is used in Chapter 3 to assess the suitability of PIC codes for the simulation of Raman amplification, and used in Chapters 4 and 5 for benchmarking, was developed by the ORISIS Consortium [1, 2]. The simulations and interpretation of the results were carried out by the author.

The three-wave model used in Chapter 4 was originally developed by Bernhard Ersfeld [3]. The extensions to the model discussed in this thesis were carried out by the author. Benchmarking and the case studies were carried out by the author.

The aPIC code was developed by MinSup Hur [4]. The author was responsible for the development of a new laser update, and the inclusion of chirped pulses and single-pulse ponderomotive effects. Benchmarking and the case study were carried out by the author.

Contents

Abstract	i
Acknowledgements	iii
Role of the author	v
Contents	vi
List of Figures	xii
1 Introduction	1
1.1 Lasers	2
1.1.1 Introduction	2
1.1.2 Ultraintense laser pulses	4
1.2 Raman amplification	6
1.2.1 Raman scattering	6
1.2.2 Raman amplification in plasma	9

1.3	Experimental work at Strathclyde	12
2	Analytical solutions	15
2.1	Introduction	16
2.2	Linear regime	17
2.3	Pump depletion regime	21
3	Particle-in-Cell	25
3.1	Introduction	27
3.2	PIC modelling of Raman amplification	30
3.2.1	Instabilities	30
3.2.2	Numerical stability	33
3.2.3	Model limitations	47
3.2.4	Stable parameter spaces	49
3.2.5	Computational overheads	53
3.3	Conclusion	55
4	Three-wave model	58
4.1	Introduction	61
4.2	Cold plasma	63
4.2.1	Three-wave model	63
4.2.2	Nonlinear plasma waves	64

4.2.3	Wavebreaking	66
4.3	Warm plasma	68
4.3.1	Three-wave model	68
4.3.2	Coffey wavebreaking limit	72
4.4	Particle trapping	74
4.4.1	Estimating the trapped electron density	77
4.5	Limitations and comparison to other works	80
4.5.1	Limitations of the model	80
4.5.2	Other works	82
4.6	Numerical implementation	84
4.7	Benchmarking	87
4.7.1	Cold plasma	88
4.7.2	Cold plasma: wavebreaking	94
4.7.3	Thermal plasma	97
4.8	Case studies	105
4.8.1	Collisional damping	105
4.8.2	Wavebreaking	108
4.8.3	Noise	109
4.8.4	Impact of wavebreaking and trapping	110
4.9	Conclusion	113
4.10	Future work	116

5	aPIC	117
5.1	Introduction	119
5.2	Envelope PIC	121
5.2.1	Separation of scales	121
5.2.2	Approximations required for envelope-PIC	123
5.2.3	Envelope-PIC equations	125
5.2.4	Applicability	127
5.3	Original aPIC model	128
5.3.1	Introduction	128
5.3.2	Poisson solver	130
5.3.3	aPIC equations	132
5.3.4	Laser field update	135
5.3.5	Pseudo-2D model	137
5.4	Extensions to the aPIC model	138
5.4.1	Chirps	138
5.4.2	Wake fields	141
5.4.3	Advective laser update	142
5.4.4	Dispersive laser update	143
5.5	Stability and accuracy	147
5.5.1	Numerical stability	147
5.5.2	Benchmarking	151

5.6	Model demonstration: nonlinear plasma waves	154
5.7	Conclusion	159
5.8	Future work	161
6	Conclusion	162
6.1	Conclusions	163
6.2	Future work	165
6.3	Applicability to experiments	167
	Appendix	168
A.1	Finite-difference schemes	169
A.1.1	Taylor expansions	169
A.1.2	Numerical schemes	171
A.2	Envelope equations for stimulated Raman scattering	174
A.2.1	Laser evolution	174
A.2.2	Plasma Response	180
A.3	Wave equation solvers	185
A.3.1	Advective solver	186
A.3.2	Paraxial approximation	187
A.3.3	Gordon approximation	188
A.3.4	Full wave equation solver	190

List of Figures

2.1	Analytical solutions for the linear regime Plots showing the analytical solutions for the linear regime, given in Eqs. (2.6-2.7), with $v_g = 1$, $S = 0.01$. Amplitude of the a) plasma wave, b) probe (representing the seed by a Gaussian of finite width), at time $t = 1.5, 3, 4.5$	19
2.2	Analytical solutions for the pump depletion regime Plots showing the analytical solutions for the pump depletion regime, given in Eq. (2.10), based on numerical solutions to the Sine-Gordon equation, Eq. (2.11), with $v_g = 1$. Amplitude of a) the pump, probe and plasma wave at $t = 20$, b) the probe, shown in the co-moving frame, at time $t = 20, 30, 40$	23
3.1	Flow chart for a PIC solver. Main loop for a conventional Yee-PIC code.	29
3.2	Example of pump and probe evolution. Pump and probe amplitude ($eE/mc\omega_p$) against position taken at different times. Simulation parameters as in Table (3.1).	33
3.3	Influence of plasma density on numerical stability. Pump and probe amplitude ($eE/mc\omega_p$) for simulations carried out with varying plasma density and simulation resolution. Plasma density (from left to right) of 1×10^{19} , 1.4×10^{19} , 2×10^{19} and $4 \times 10^{19} \text{ cm}^{-3}$. Resolution (from top to bottom) of 50, 100, 200 and 400 cells per pump wavelength. Other parameters as in Table (3.1).	35
3.4	Influence of pump intensity on numerical stability. Pump and probe amplitude ($eE/mc\omega_p$) for simulations carried out with varying pump intensity and simulation resolution. Pump intensity (from left to right) of 2.5×10^{13} , 5×10^{13} , 7.5×10^{13} and $1 \times 10^{14} \text{ W cm}^{-2}$. Resolution (from top to bottom) of 50, 100, 200 and 400 cells per pump wavelength. Plasma density of $1 \times 10^{19} \text{ cm}^{-3}$, other parameters as in Table (3.1).	36

3.5	Influence of simulation resolution on numerical stability. Sample macroparticle phase space and corresponding electrostatic field ($eE/mc\omega_p$) over a short region for varying simulation resolution. Resolution (from top to bottom) of 50, 100 and 200 cells per pump wavelength. Density of $1.4 \times 10^{19} \text{ cm}^{-3}$, other parameters as in Table (3.1).	39
3.6	Comparison of “numerical wavebreaking” with grid heating. Pump and probe amplitude ($eE/mc\omega_p$) and sample macroparticle phase space for different initial conditions. a) initially cold plasma with a resolution of 50 cells per pump wavelength, b) initial temperature 2.35 eV (equivalent to the saturation temperature for the grid heating instability for a resolution of 50 cells per pump wavelength) with a resolution of 400 cells per pump wavelength. Other parameters as in Table (3.1).	42
3.7	Influence of macroparticle number on numerical stability. Pump and probe amplitude ($eE/mc\omega_p$) and sample macroparticle phase space for different macroparticle densities: 5 000, 10 000 and 20 000 per pump wavelength. Grid resolution of 50 cells per pump wavelength, other parameters as in Table (3.1).	43
3.8	Influence of smoothing techniques on numerical stability. Pump and probe amplitude ($eE/mc\omega_p$) and sample macroparticle phase space for different smoothing techniques: no smoothing, smoothing of currents, and smoothing of fields. Parameters as in Table (3.1).	44
3.9	Influence of interpolation on numerical stability. Pump and probe amplitude ($eE/mc\omega_p$) and sample macroparticle phase space for different levels of particle interpolation: linear, quadratic and quartic. Parameters as in Table (3.1).	45
3.10	Influence of density on spontaneous backscatter. Pump and probe amplitude ($eE/mc\omega_p$) for simulations carried out with different plasma densities. Density of a) $1 \times 10^{19} \text{ W cm}^{-2}$, b) $1.4 \times 10^{19} \text{ W cm}^{-2}$, c) $2 \times 10^{19} \text{ W cm}^{-2}$ and d) $4 \times 10^{19} \text{ W cm}^{-2}$. Temperature of 100 eV, other parameters as in Table (3.1). . .	50
3.11	Influence of temperature on spontaneous backscatter. Pump and probe amplitude ($eE/mc\omega_p$) for simulations carried out with different plasma temperatures. Temperature of a) 10 eV, b) 20 eV, c) 30 eV, d) 40 eV, e) 50 eV and f) 60 eV. Density of $1 \times 10^{19} \text{ cm}^{-3}$, other parameters as in Table (3.1). . .	51

3.12	Influence of pump chirp rate on spontaneous backscatter. Pump and probe amplitude ($eE/mc\omega_p$) for simulations carried out with different pump chirp rates. From left to right, pump bandwidth as a fraction of pump frequency, $\Delta\omega_0/\omega_0 = 0, 0.002$ and 0.004 . Density of 1×10^{19} cm^{-3} , temperature of 50 eV, other parameters as in Table (3.1).	51
4.1	Illustration of separatrix for warm plasma Electron velocities in the wave frame, illustrating the separatrix about v_ϕ	74
4.2	Maxwellian and waterbag distributions. The model assumes that for an initially Maxwellian distribution, $\mathcal{F}_0(v)$, the distribution at Coffey's warm wavebreaking limit, $\mathcal{F}'(v)$, satisfies $\mathcal{F}'(v_\phi) = \mathcal{F}_0(\sqrt{3}v_{th})$	77
4.3	Maxwellian and modified waterbag distributions. The model assumes that for an initially Maxwellian distribution, $\mathcal{F}_0(v)$, the distribution at the modified waterbag wavebreaking limit, $\mathcal{F}'(v)$, satisfies $\mathcal{F}'(v_\phi) = \mathcal{F}_0(\sqrt{3\alpha}v_{th})$. a) $\alpha = 2$ ($f < f_{WB}$), $\mathcal{F}'(v_\phi)$ is smaller, b) $\alpha = 0.5$ ($f > f_{WB}$), $\mathcal{F}'(v_\phi)$ is larger.	79
4.4	Schematic of three-wave solver. Leapfrog update for laser field solver (LS) and electrostatic field solver (ES).	85
4.5	Illustration of high resolution required for PIC convergence. Pump and probe envelopes from PIC and three-wave model, using different PIC resolutions. From left to right, 50, 100 and 200 cells per pump wavelength. Density of 1×10^{19} cm^{-3} , pump intensity of 2.5×10^{13} W cm^{-2} , other parameters as in Table (3.1).	87
4.6	Benchmarking for varying pump intensity. Pump and probe envelopes from PIC and three-wave model for a pump intensity of: a) 2.5×10^{13} , b) 5×10^{13} , c) 7.5×10^{13} and d) 1×10^{14} W cm^{-2} . Density of 1×10^{19} cm^{-3} , other parameters as in Table (3.1).	89
4.7	Benchmarking for varying plasma density. Pump and probe envelopes from PIC and three-wave model for plasma density of: a) 1×10^{19} , b) 1.4×10^{19} , c) 2×10^{19} and d) 4×10^{19} cm^{-3} . Other parameters as in Table (3.1).	90
4.8	Benchmarking for varying detuning. Pump and probe envelopes from PIC and three-wave model pump/probe detuning of: a) $1.05 \omega_p$, b) $0.95 \omega_p$, and c) $0.9 \omega_p$. Other parameters as in Table (3.1).	91

4.9	Benchmarking for varying pump chirp. Pump and probe envelopes from PIC and three-wave model for chirped pump of bandwidth: a) 0.4%, b) 1.0% and c) 2.0%. Other parameters as in Table (3.1).	92
4.10	Influence of harmonics on Raman amplification. Pump and probe envelopes from PIC and three-wave model. In the first figure, the influence of harmonics is included in the three wave model. In the second, it is neglected. Plasma density of $2 \times 10^{19} \text{ cm}^{-3}$, other parameters as in Table (3.1).	93
4.11	Benchmarking of wavebreaking model near wavebreaking limit. Pump and probe envelopes from PIC and three-wave model in wavebreaking regimes. Both the extended three-wave model and the conventional model with no breaking (3w NB) are shown. Plasma density and pump intensity of: a) $5 \times 10^{18} \text{ cm}^{-3}$, $1 \times 10^{14} \text{ W cm}^{-2}$, b) $1 \times 10^{19} \text{ cm}^{-3}$, $2 \times 10^{14} \text{ W cm}^{-2}$, and c) $2 \times 10^{19} \text{ cm}^{-3}$, $4 \times 10^{14} \text{ W cm}^{-2}$. Other parameters as in Table (3.1).	95
4.12	Benchmarking of wavebreaking model beyond wavebreaking limit. Pump and probe envelopes from PIC and three-wave model in wavebreaking regimes. Both the extended three-wave model and the conventional model with no breaking (3w NB) are shown. Plasma density of $5 \times 10^{18} \text{ cm}^{-3}$, pump intensity: a) 1×10^{14} , b) 2×10^{14} and c) $4 \times 10^{14} \text{ W cm}^{-2}$. Other parameters as in Table (3.1).	96
4.13	Benchmarking of wavebreaking and trapping models in warm plasma. Pump and probe envelopes from PIC and three-wave model for varying temperature. The left hand plot shows the wavebreaking model using both the cold wavebreaking limit (3w CB) and Coffey's warm wavebreaking limit (3w TB). The right hand plot shows the trapping model. Initial temperature of: a) 1 eV, b) 5 eV, c) 10 eV, d) 25 eV and e) 50 eV. Plasma density of $5 \times 10^{18} \text{ cm}^{-3}$, other parameters as in Table (3.1).	99

4.14	Benchmarking of wavebreaking and trapping models in warm plasma (high intensity). Pump and probe envelopes from PIC and three-wave model for varying temperature. The left hand plot shows the wavebreaking model using both the cold wavebreaking limit (3w CB) and Coffey’s warm wavebreaking limit (3w TB). The right hand plot shows the trapping model. Initial temperature of: a) 1 eV, b) 5 eV, c) 10 eV, d) 25 eV and e) 50 eV. Plasma density of $5 \times 10^{18} \text{ cm}^{-3}$, probe intensity of $1 \times 10^{14} \text{ W cm}^{-2}$, other parameters as in Table (3.1).	102
4.15	Comparison with Bohm-Gross model Benchmarking Pump and probe envelopes from PIC and three-wave model, using both the trapping model derived in this work and that derived by Bohm and Gross (3w BG). Probe intensity a) $1 \times 10^{12} \text{ W cm}^{-2}$, b) $1 \times 10^{14} \text{ W cm}^{-2}$. Plasma density $5 \times 10^{18} \text{ cm}^{-3}$, initial temperature 5 eV, other parameters as in Table (3.1).	104
4.16	Influence of collisions. Probe envelopes, showing full thermal model and collisionless (PIC-like) model, for 5 eV and 15eV. Other parameters as in Table (3.1).	107
4.17	Influence of thermal effects with and without collisions. Parameter scans showing peak amplified probe intensity, comparing cold plasma model, full thermal model and collisionless (PIC-like) model, for varying pump/probe detuning, and varying pump chirp. Other parameters as in Table (3.1).107	107
4.18	Influence of wavebreaking. Snapshots of probe intensity during amplification in the wavebreaking regime. Pump intensity is shown for the final snapshot. Simulation carried out for cold plasma of density $5 \times 10^{18} \text{ cm}^{-3}$, initial pump of intensity $1 \times 10^{15} \text{ W cm}^{-2}$, initial probe of intensity $1 \times 10^{15} \text{ W cm}^{-2}$ and FWHM duration 100 fs. Other parameters as in Table (3.1).	108
4.19	Influence of thermal effects on Raman growth rates. Peak amplified probe intensity, comparing cold plasma model, and thermal model at 5 eV and 15 eV, using a) monochromatic, b) up chirped, and c) down chirped pump. Pump bandwidth in the chirped pump case is $0.05\omega_0$. Other parameters as in Table (3.1).	111

4.20	Influence of wavebreaking and particle trapping on Raman growth rates. Peak amplified probe intensity, comparing cold plasma model, including and excluding wavebreaking, and thermal model at 5 eV, including and excluding particle trapping. Probe intensity of a) $1 \times 10^{12} \text{ W cm}^{-2}$, b) $1 \times 10^{14} \text{ W cm}^{-2}$. Other parameters as in Table (3.1).	112
5.1	Influence of aPIC averaging length on probe laser field. Amplified probe after 2500 steps for different grid sizes. Note that due to the difference in timestep, the snapshots correspond to different times, for easier comparison.	140
5.2	Comparison of laser solvers. Absolute and real amplitudes of a laser pulse ~ 7 and ~ 35 ps after entering the plasma. Parameters as for probe in Table 3.1.	146
5.3	Influence of simulation resolution on numerical stability (constant macroparticle number). Pump and probe envelopes for aPIC simulations with varying spatial resolution for the electrostatic solver. Resolution, from top to bottom, of 50, 100, 200 and 400 cells per pump wavelength. Timestep of $0.98 \lambda_b/c$, with other parameters as in Table 3.1. Total macroparticle number is kept constant as resolution is increased.	148
5.4	Influence of simulation resolution on numerical stability (scaled macroparticle number). Probe envelope over a short range for aPIC simulations with varying spatial resolution for the electrostatic solver. Resolution, from left to right, of 100, 200 and 400 cells per pump wavelength. Timestep of $0.98 \lambda_b/c$, with other parameters as in Table 3.1, but with macroparticle number per cell kept constant as resolution is increased.	150
5.5	Influence of temporal resolution on numerical stability (constant macroparticle number). Probe envelope over a short range for aPIC simulations with varying temporal resolution. Timestep, from left to right, of 0.98, 0.49 and $0.245 \lambda_b/c$. Other parameters as in Table 3.1.	150
5.6	Benchmarking for varying detuning. Pump and probe envelopes from OSIRIS and aPIC for a detuning of a) ω_p , b) $0.95\omega_p$, and c) $1.05\omega_p$. Timestep of $0.245\lambda_b/c$, with other parameters as in Table 3.1.	152

5.7	Influence of simulation resolution. Plots of amplified probe for varying resolution. “1x” corresponds to a timestep of $0.98 \lambda_b/c$, with other parameters as in Table 3.1. Temporal and spatial resolution are increased, with macroparticle number per cell kept constant.	153
5.8	Influence of probe chirp on amplification. Plots of absolute probe gain over the interaction for a) down- and b) up-chirped probes. Initial pulse spectrum shown for comparison. Amplitudes are normalised to the peak spectral amplitude of the initial pulse.	156
5.9	Phase space evolution. Sample electron phase space, showing snapshots of the electron distribution (velocity against position) at different times and positions. Phase spaces are shown every 133 fs over $2 \mu\text{m}$ intervals, to the rear, centre and front of the interaction region, separated by $80 \mu\text{m}$	158

Chapter 1

Introduction

1.1 Lasers

1.1.1 Introduction

Since its first demonstration in 1960 [5], the laser has been incorporated into a wide range of science and technology applications. Lasers range in scale from small diode lasers used in pointers, to entire facilities built to house petawatt systems.

Their versatility arises from the range of properties they offer as a light source, such as monochromaticity, coherence, low divergence, high focusability, and ability to deliver either continuous wave beams or ultra-short, ultra-intense pulses.

Everyday applications, such as CD players and optical communication, contrast with medical applications, such as scar treatment [6] and more esoteric scientific pursuits, such as measuring the distance to the moon [7] and the creation of Bose Einstein condensates via laser cooling [8]. Current areas of research include laser wakefield acceleration [9], gravity wave detection [10] and inertial confinement schemes for nuclear fusion [11]. As the peak intensities of lasers continue to rise, they open the potential for future studies of quantum electrodynamics [12], and the ability to conduct new laboratory experiments relevant to astrophysical phenomena [13].

Amplification in a laser is limited by the peak fluence which the amplifying

medium can support. For continuous wave or pulsed lasers, this limit is due to the development of nonlinearities, which lead to intensities beyond the damage threshold of the medium. As the creation of ultra-short laser pulses typically relies on nonlinearities [14], the practical limitation is the damage threshold of optical components.

1.1.2 Ultraintense laser pulses

The limitation of peak fluence in the amplifying medium may be circumvented by using chirped pulse amplification (CPA), developed by Strickland and Mourou [15]. CPA, along with the optical parametric amplification equivalent (OPCPA) [16], are the basis of current ultra-short, ultra-intense laser pulses. The laser pulse is stretched prior to amplification by a dispersive optical system, giving rise to a chirped pulse, which is then amplified and recompressed. Pulse stretching and compression is usually achieved using gratings.

This method reduces the peak laser fluence inside the amplifying medium, avoiding optical damage and allowing much higher laser energies to be attained. The avoidance of thermal damage limits the average power, reducing the repetition rate. As the laser power increases, further stretching, and so larger gratings, are required. High powers also require that the compressed beam remain unfocussed to avoid damaging optical components, necessitating metre diameter optics in petawatt lasers, greatly increasing both the cost and size of the system.

Raman amplification in plasma has been suggested as an alternative method for the creation of ultra-intense laser pulses [17]. Using plasma as an amplifying medium avoids optical damage, as plasma has no bonded struc-

ture, and fresh plasma may be used for each shot. This technique could potentially either be used instead of CPA, with the amplified pulse directly incident on the target, removing the need for both gratings and large diameter optics, or in conjunction with CPA, as an additional amplification stage in conventional petawatt systems, analogous to the afterburner approach which has been used to increase the peak energy attainable in particle accelerators [18].

1.2 Raman amplification

1.2.1 Raman scattering

Inelastic scattering of light is characterised by a change of frequency between the incident and scattered photons. In the case of Compton scattering, first observed by Compton for X-rays [19], the difference in energy is accounted for by the acceleration of an electron. Compton scattering contributed to the evidence of the particle nature of light, as the radiation-recoil scattering of a purely electromagnetic wave would only give rise to significant differences between the incident and scattered energy at high intensity. Compton scattering, on the other hand, is a frequency dependent effect, and may be seen at low intensities, similar to the photoelectric effect.

Drawing analogy with Compton scattering, Raman and Krishnan anticipated the inelastic scattering of visible light. Their experiments for scattering in liquids and gasses showed that a small fraction of scattered light was frequency shifted relative to the source [20]. Working independently, Landsberg and Mandelstam observed the same effect in crystals [21].

In Raman scattering, the difference in energy between incident and scattered photons is deposited as a resonant excitation (or de-excitation) of the medium. Many discussions of Raman amplification describe this as the excitation of the medium to some “virtual state”, before the immediate decay to

a real state, as an analogue to fluorescence. The vibrational states in insulators allow Raman scattering to occur at much lower energies than Compton scattering, which requires the ionisation of an atom.

As the difference in energy between incident and scattered photons corresponds to a fundamental excitation of the medium, spectroscopy of Raman scattered light may be used to obtain information on the structure of the medium. The resonant excitations are typically of vibrational states, and are widely used for the identification of minerals [22], although other excitations are possible, such as collective charge-ordering [23] and magnon excitation [24].

In addition to spontaneous scattering, Raman scattering may be stimulated through the presence of a photon with a frequency similar to that of the scattered photon. This provides the basis for Raman amplification, in which an optical signal may be amplified by a pump pulse of higher frequency [25]. One application is the distributed amplification of signals in fibre optics to compensate for attenuation losses in telecommunications [26]. Fibres may be doped to increase the number of possible vibrational transitions, allowing a large amplification bandwidth. Stimulated scattering is also the basis of the Raman “laser”, in which spontaneous scattering inside a pumped cavity is amplified. While the Raman laser generates coherent radiation, it is not a true laser, as it relies on stimulated scattering rather than stimulated

emission. This removes the requirement for population inversion of the gain medium, although there exists an intensity threshold for coherent emission [27]. The broad bandwidth amplification in doped fibres allows the creation of ultra-short (< 1 ps) laser pulses [28].

1.2.2 Raman amplification in plasma

The Raman process in plasma involves the coupling of the incident and scattered photons through the resonant excitation of a plasma wave. Scattering in plasma may also be caused by Brillouin scattering [29], in which the excitation is of an ion-acoustic wave, or stimulated Compton scattering [30], in which scattering is by resonant thermal particles. Stimulated scattering has been harnessed to transfer energy between laser pulses, allowing the radiation profile in inertial confinement fusion experiments to be modified without significant changes to the laser system [31, 32].

A more physical understanding of this coupled three-wave system may be gained by considering the beat of two lasers driving a plasma wave through the ponderomotive force: the perturbation of the refractive index due to the resultant density perturbation will act as a moving Bragg grating, which reflects the pump pulse into the Doppler-frequency-shifted probe, amplifying the latter.

Early treatments were limited to the amplification of long pulses growing from noise [33–36], prompted by concerns of parasitic instabilities in inertial confinement fusion applications [37], or as a possible heating mechanism for magnetically confined plasmas [38].

The narrow resonance of the plasma wave might be expected to make

Raman amplification in plasma unsuitable for the amplification of short pulses due to gain narrowing, and indeed, in linear, homogeneous systems this is the case. However, Shvets *et al* [17] realised that at high intensities, the ponderomotive force of the laser beat exceeds the electrostatic force associated with the plasma perturbation. In this case, the excitation is charge bunching at the beat frequency, rather than the resonant plasma frequency. In the limit where electrostatic forces may be neglected, this nonlinear coupling gives rise to superradiant amplification of a short pulse, with amplitude scaling proportionally, and duration varying inversely, to propagation distance [39].

This is often referred to as the Compton regime, in analogy to the equivalent behaviour in the free-electron laser, where space-charge effects are negligible, leading to superradiant growth [40, 41].

The ability to provide broad-bandwidth amplification, and the absence of a damage threshold and the need for pulse stretching prior to amplification, makes plasma a strong candidate for the amplification of ultra-short, ultra-intense laser pulses. Further works showed that the amplification of short pulses can be achieved in the linear regime through the exploitation of nonlinearities: Malkin *et al* showed that pump depletion acts to suppress amplification to the rear of the pulse [42], while Ersfeld and Jaroszynski showed that a chirped pump allowed broad-bandwidth amplification despite the nar-

row resonance of the plasma wave [43]. Both methods are characterised by superradiant scaling.

However, the use of plasma as an amplification medium is not without its limitations. While plasma has no damage threshold, there exist many instabilities that can act to limit amplification or reduce the quality of the amplified pulse. As small signals can be amplified exponentially, noise arising from the spontaneous scattering of the pump can be significantly amplified [44]. This parasitic amplification can prematurely deplete the pump, reducing amplification of the probe, and lead to a large pedestal or prepulse. Wavebreaking in the linear regime can act to significantly damp the excited plasma wave, limiting amplification. Heating results in a shift of the plasma resonance, lowering the growth rate. Filamentation and self focussing can lower growth and degrade the quality of the amplified beam [45].

Experimental work has shown bandwidth-broadening of an amplified pulse [46, 47] and pulse compression [47]. Energy gains of a factor of 350 have been observed in a double-pass amplifier, with an energy transfer efficiency over two passes of 6.4 % [47].

1.3 Experimental work at Strathclyde

The Alpha-X project [48] aims to create a compact radiation source operating in the X-ray and gamma ray regimes [49]. A laser wakefield accelerator [9] is used to create a highly relativistic short electron bunch, which is then passed through an external undulator to produce high energy photons. Alternatively, interaction of the electron bunch with the plasma electrostatic field and the electromagnetic field of the laser driver can lead to betatron oscillations of the accelerated electrons inside the plasma, again producing high energy synchrotron-like radiation [50].

A plasma based Raman amplifier would be a synergistic choice for the laser system, as, in addition to the cost and space savings, there is a significant overlap with the laser wakefield accelerator, both in terms of technology and required expertise. For this reason the University of Strathclyde has been working on Raman amplification in plasma as part of the Alpha-X collaboration [3, 43, 51, 52].

Current experiments are based around a plasma capillary [53]. An electrical discharge creates a 4 cm plasma channel for amplification, and provides a guiding mechanism for the laser pulses [53], which should allow an improved interaction through matching the spot size of the two counterpropagating pulses. The length of a channel allows the use of a long pump pulse, decreas-

ing the intensity, avoiding instabilities such as spontaneous backscatter and filamentation. Plasma densities are typically in the region of $2 \times 10^{18} \text{ cm}^{-3}$, with initial temperatures around 5 eV [54]. The current experimental work is aimed at characterising the interaction in the “linear” chirped pump regime [51], with a typical intensity for pump and probe pulses in the region of $4 \times 10^{13} \text{ W cm}^{-2}$.

Many competing effects are possible during amplification, which include heating, guiding, self-focussing, particle trapping, Landau damping, ponderomotive nonlinearity and pump depletion. For this reason, a purely analytical treatment of the experimental results is not sufficient, and numerical simulations of the interaction are necessary to identify the dominant processes.

The aim of this work is to identify and develop suitable methods for the simulation of the ongoing experimental work at Strathclyde. Simulations must capture the physical processes involved, but must also be viable with the computational resources available.

This thesis addresses three candidates for these simulations: Particle in Cell (PIC) codes, which are widely used to simulate laser-plasma interactions [45], but are typically computationally intensive [55] and are subject to numerical instabilities [56]; three-wave models derived from a fluid treatment of plasma, which have been used to investigate Raman amplification [42, 43], but have limitations due to the absence of particle effects; and the aPIC

code, developed by Min Sup Hur [4, 57], which reduces the computational overhead of conventional PIC simulations through the use of approximations, but has not been extensively tested, and requires further development if it is to simulate the parameter regime of interest.

Chapter 2

Analytical solutions

2.1 Introduction

The evolution of the pump, probe and plasma wave involved in the Raman backscatter interaction may be described using a three-wave theory. The full equations, which are derived in Appendix A.2 and are the basis for the three-wave model discussed in Chapter 4, may also be used to obtain analytical results in certain idealised situations. We start with Eqs. (A.21, A.31), and assume optimal detuning and real envelopes for the pump and probe. As we are interested only in the form of the analytical solutions, we additionally neglect the scaling constants, to give:

$$\begin{aligned}\left(\frac{\partial}{\partial t} - v_g \frac{\partial}{\partial z}\right) a_0 &= -f a_1, \\ \left(\frac{\partial}{\partial t} + v_g \frac{\partial}{\partial z}\right) a_1 &= f a_0, \\ \frac{\partial}{\partial t} f &= a_0 a_1.\end{aligned}\tag{2.1}$$

We here derive an analytical solution for the linear regime, and a semi-analytical solution for the pump depletion regime [42]. An analytical solution is also possible for the case of a chirped pump [43], while an analytical treatment of the Compton regime [17] allows the scaling properties of the amplified probe to be calculated.

2.2 Linear regime

In the linear regime, pump depletion is negligible, and so a_0 may be taken as constant. For convenience we set $a_0 = 1$, equivalent to measuring time in units of inverse growth-rate. The two remaining coupled equations may be simplified by transforming into new variables [58], $\eta = z/v_g$, $\sigma = t - z/v_g$:

$$\begin{aligned}\frac{\partial}{\partial \eta} a_1 &= f, \\ \frac{\partial}{\partial \sigma} f &= a_1.\end{aligned}\tag{2.2}$$

We consider the case of a short seed pulse, \hat{a}_1 , entering an unperturbed plasma. The Raman interaction with the pump gives rise to a scattered wave. In this linear regime, we may write the total field as a linear superposition of the seed and the scattered wave, \tilde{a}_1 :

$$a_1 = \hat{a}_1 + \tilde{a}_1\tag{2.3}$$

We represent the seed pulse as a scaled delta-function located at $\sigma = 0$ with integral $S \ll 1$, i.e. $\hat{a}_1 = S\delta(\sigma)$. Formally, this corresponds to the Green's function for a seed of arbitrary shape, but for a short seed we may use the delta-function to represent the entire seed. The scattered wave will initially have zero amplitude, while the interaction of the pump and seed will excite a

plasma wave of amplitude S . We therefore eliminate a_1 to obtain an equation for f :

$$\frac{\partial^2}{\partial \eta \partial \sigma} f = f. \quad (2.4)$$

Such equations have a self-similar solution, i.e. dependent on a single variable.

We introduce a new variable $\Phi = \eta\sigma$, to give:

$$\left(\Phi \frac{d^2}{d\Phi^2} + \frac{d}{d\Phi} \right) f = f, \quad (2.5)$$

which (for positive η) has the solution:

$$f = \begin{cases} S I_0(\sqrt{\Phi}), & \Phi \geq 0, \\ 0, & \Phi < 0, \end{cases} \quad (2.6)$$

where I_n is the n th-order modified Bessel function of the first kind. The solution for the scattered wave is then:

$$\tilde{a}_1 = \begin{cases} \frac{S I_1(\sqrt{\Phi})}{2\sqrt{\Phi}}, & \Phi \geq 0, \\ 0, & \Phi < 0. \end{cases} \quad (2.7)$$

The results, illustrated in Figure 2.1, show that \tilde{a}_1 and f increase mono-

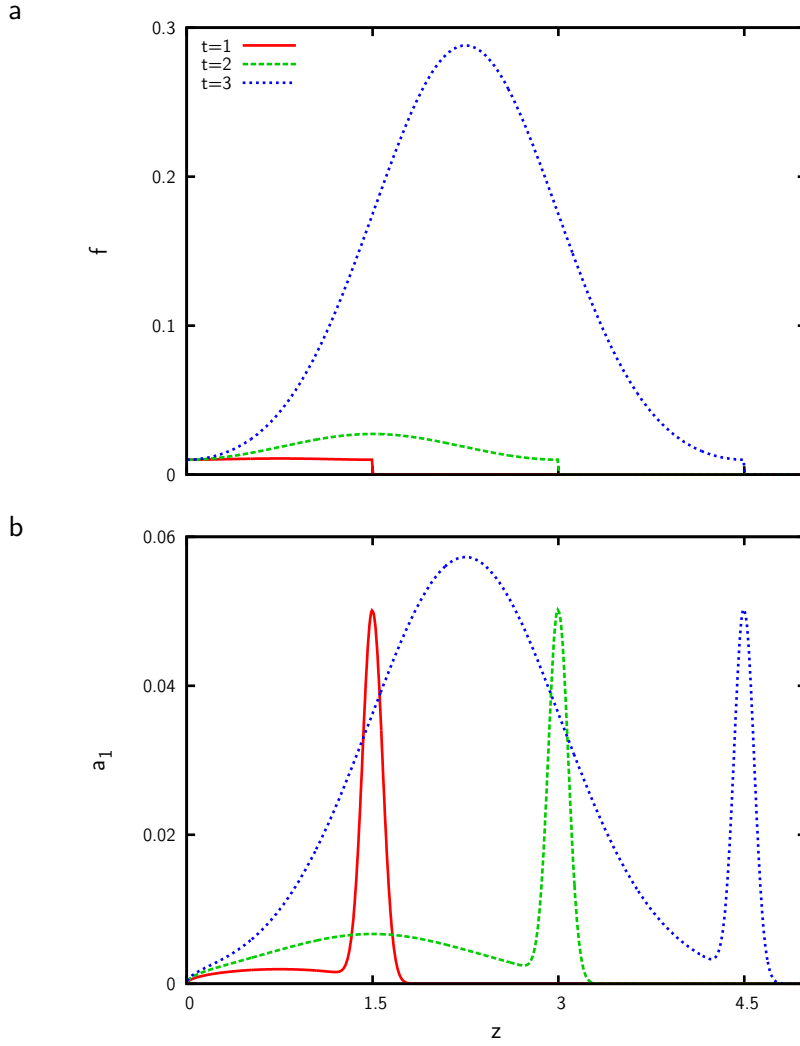


Figure 2.1: **Analytical solutions for the linear regime**

Plots showing the analytical solutions for the linear regime, given in Eqs. (2.6-2.7), with $v_g = 1$, $S = 0.01$. Amplitude of the a) plasma wave, b) probe (representing the seed by a Gaussian of finite width), at time $t = 1.5, 3, 4.5$.

tonically with Φ , giving a maximum at $z = v_g t/2$. The solution may be understood as follows: the short seed pulse excites a plasma wave, which continues to scatter the pump after the seed has passed. The scattered wave further excites the plasma wave, leading to more scattering, giving exponential-like growth.

σ and η represent the time over which the plasma and scattered waves, respectively, have evolved. At the plasma edge, σ is a maximum but η , and so \tilde{a}_1 , is zero, leading to zero growth. The roles are reversed at the position of the seed: η is a maximum but σ , and so f , is zero, again giving zero growth. The point $z = v_g t/2$ represents the maximum product of η and σ , i.e. the maximum product of growth rate and evolution time, giving the largest amplitude for both the plasma and scattered waves.

It is evident that this regime is therefore unsuitable for the amplification of short pulses: while the scattered wave grows exponentially, its width increases proportionally to the interaction length.

2.3 Pump depletion regime

It is clear that the assumption of negligible pump depletion is only valid for a small scattered wave. For larger probes, it becomes necessary to model the evolution of the pump pulse. We transform into the co-moving frame of the probe, $\zeta = z - v_g t$, $\varsigma = t$, to give:

$$\begin{aligned} \left(\frac{\partial}{\partial \varsigma} - 2 \frac{\partial}{\partial \zeta} \right) a_0 &= -f a_1, \\ \frac{\partial}{\partial \varsigma} a_1 &= f a_0, \\ \left(\frac{\partial}{\partial \varsigma} - \frac{\partial}{\partial \zeta} \right) f &= a_0 a_1. \end{aligned} \tag{2.8}$$

Again we seek self-similar solutions, assuming the evolution of the three waves depends on the variable $\Theta = \zeta \varsigma$. The governing equations then become:

$$\begin{aligned} (\zeta - 2\varsigma) \frac{d}{d\Theta} a_0 &= -f a_1, \\ \zeta \frac{d}{d\Theta} a_1 &= f a_0, \\ (\zeta - \varsigma) \frac{d}{d\Theta} f &= a_0 a_1. \end{aligned} \tag{2.9}$$

From the evolution equations for a_0 and f , a set of solutions may be found in the limit where $|\zeta| \ll |\varsigma|$, i.e. close to the leading edge of the probe:

$$\begin{aligned}
 a_0 &= \cos(u/2), \\
 a_1 &= -\frac{\varsigma}{\sqrt{2}} \frac{du}{d\Theta}, \\
 f &= \sqrt{2} \sin(u/2).
 \end{aligned}
 \tag{2.10}$$

The governing equation for u may be found by substituting these solutions into the equation for a_1 :

$$\begin{aligned}
 \zeta \frac{d}{d\Theta} \left(\frac{\varsigma}{\sqrt{2}} \frac{du}{d\Theta} \right) &= -\frac{1}{\sqrt{2}} \sin(u) \\
 \Theta \frac{d^2 u}{d\Theta^2} + \frac{du}{d\Theta} &= -\sin(u)
 \end{aligned}
 \tag{2.11}$$

This is the Sine-Gordon equation[59]. Results based on numerical solutions for u are shown in Figure 2.2. These equations describe superradiant [39] growth of the probe, with an amplitude increasing proportionally to the propagation distance, while duration scales inversely to the propagation distance.

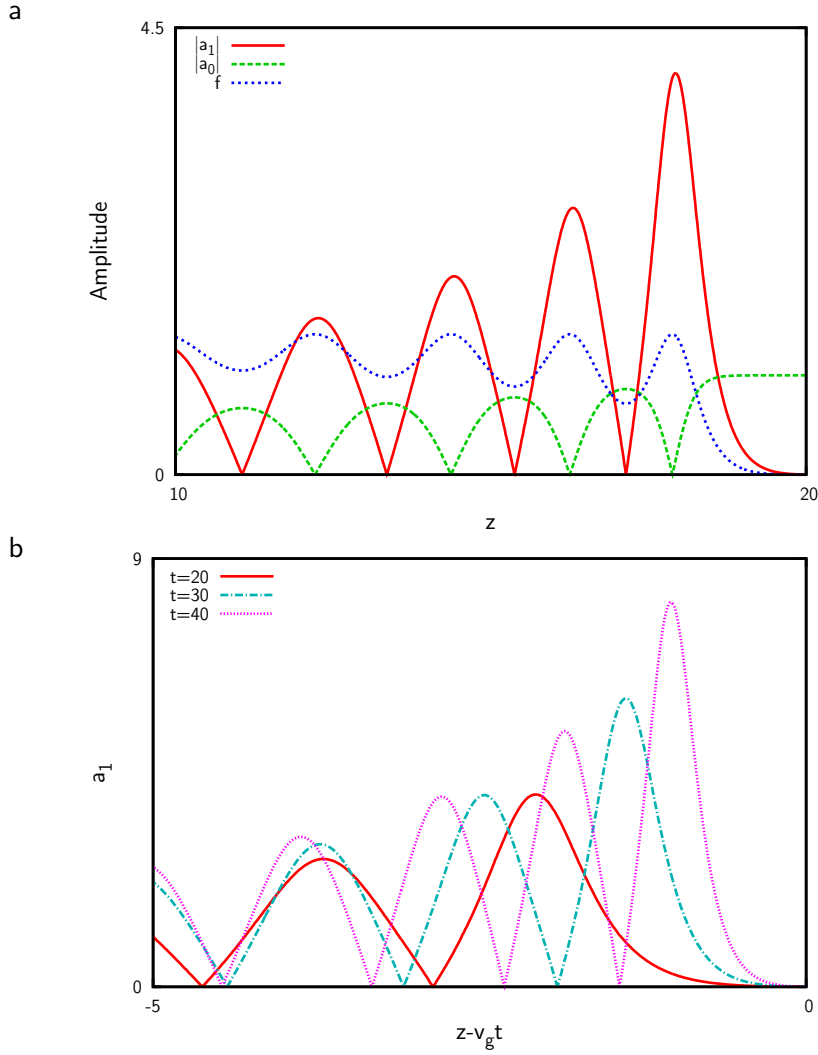


Figure 2.2: Analytical solutions for the pump depletion regime

Plots showing the analytical solutions for the pump depletion regime, given in Eq. (2.10), based on numerical solutions to the Sine-Gordon equation, Eq. (2.11), with $v_g = 1$. Amplitude of a) the pump, probe and plasma wave at $t = 20$, b) the probe, shown in the co-moving frame, at time $t = 20, 30, 40$.

The combination of superradiant amplification and pulse compression, which is also observed for the case of a chirped pump or in the Compton regime, makes Raman amplification in plasma an attractive method for the creation of ultra-intense, ultrashort laser pulses.

Chapter 3

Limitations of particle in cell simulations of Raman amplification

Chapter Abstract

The limitations of PIC codes relevant to the simulation of Raman amplification are investigated. Numerical stability and computational overhead are both issues, each of which can be ameliorated at the expense of the other. We identify a new numerical instability, which we call “numerical wavebreaking”, in which failure to resolve the harmonics of high amplitude plasma waves results in particles escaping the ordered oscillations of the plasma wave, leading to damping. We further show that both smoothing and higher-order interpolation increase the effect of this instability.

The absence of collisional damping, which can act to reduce instability growth in physical systems, will not only reduce the accuracy of the model, but limit the parameter space in which simulations may be carried out.

We conclude that PIC codes are not currently suitable for the simulation of several parameter ranges of interest, including the experimental work at Strathclyde. Parameter spaces in which instabilities can be managed are identified, to allow benchmarking of the other simulation tools considered in this thesis.

3.1 Introduction

Of the computational models for the study of laser-plasma interactions, Particle-in-Cell (PIC) codes are one the most widely used, due to their ability to model both particle and (with sufficient computational resources) collective effects [60]. They are perhaps most extensively used in the simulation of plasma-wakefield acceleration, which requires both collective effects, for the simulation of the wake, and particle effects, to model the injection and acceleration processes.

Particles, typically macroparticles that represent many physical particles, are modelled by their position and velocity; the forces on the particles are calculated every timestep, and used to update these values. For a large number of particles, calculating the forces that arise from all particle-particle interactions becomes computationally prohibitive. PIC codes avoid this problem by depositing the particle fields onto a grid; particles then interact with the grid field, reducing the number of calculations required (overheads scaling with the number of particles, N , compared with N^2 when particle-particle interactions are included explicitly).

One consequence of this method is that direct particle-particle interactions, which give rise to processes such as inverse-bremsstrahlung (collisional) heating, are not modelled. It is possible to include an additional calculation

for binary collisions for spatially localised particles, with more distant particles interacting only with the grid field. While this inevitably involves a significant increase in computational overheads, it remains far less onerous than using a full particle-particle model. However, such a treatment is made more complicated by the use of macroparticles, as collisions between groups of particles will differ from Coulomb collisions between individual particles [61, 62]. At the time this work was undertaken, neither OSIRIS nor VORPAL, two of the most commonly used PIC codes [63], supported collisions.

The most common field solver used for PIC codes is the finite-difference time domain implementation of Faraday’s and Ampere’s laws derived by Yee [64]. The full Yee algorithm is an extension of the leap-frog method, detailed in Appendix A.1, to three-dimensions. In the one-dimensional case, used in this work, the Yee method reduces to a simple leap-frog method, with the electric and magnetic fields calculated at half-cell offsets. A typical Yee-PIC update is illustrated in Figure 3.1.

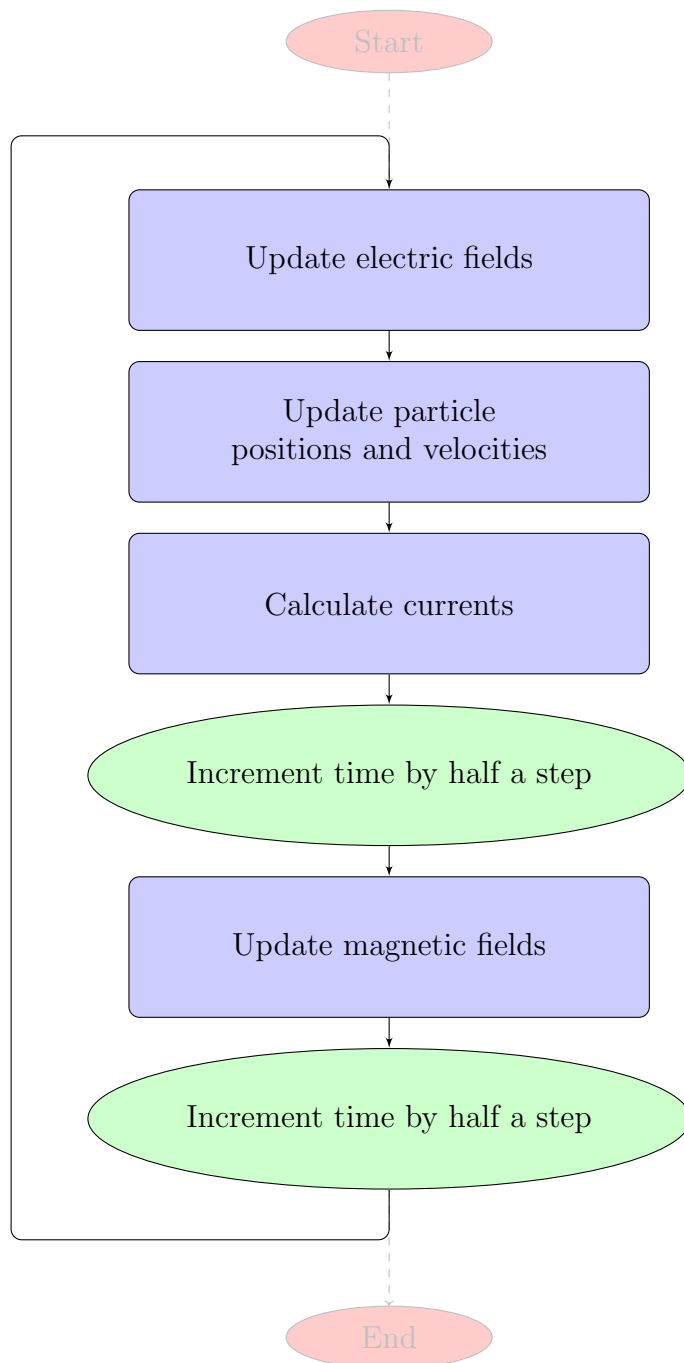


Figure 3.1: **Flow chart for a PIC solver.**
Main loop for a conventional Yee-PIC code.

3.2 PIC modelling of Raman amplification

3.2.1 Instabilities

Instabilities can broadly be defined as any processes that involve positive feedback. Raman scattering in plasma, which we seek to harness as an amplifying mechanism, is one example. The pump and probe fields beat to drive a plasma wave, which leads to the amplification of the probe. In regimes where the pump amplitude is much larger than that of the probe, the increase in the probe amplitude leads to a larger beat, driving the plasma wave more strongly, further amplifying the probe. In the linear regime, probe growth is proportional to the plasma wave amplitude, which in turn grows proportionally to the probe amplitude (c.f. Appendix A.2), which leads to exponential growth.

In damped systems, the growth rate of the instability must exceed the rate at which it is damped for net growth to occur. This introduces an instability threshold, below which the instability does not grow. In the case of Raman scattering in warm plasma, the probe will be subject to collisional damping, while the plasma wave will be subject to both collisional and Landau damping. The pump amplitude, which determines the Raman growth rate, must therefore be sufficiently high to exceed the influence of these damping mechanisms before the probe can be amplified.

Some instabilities saturate when they reach a certain amplitude. For Raman scattering, when the pump and probe have equal amplitudes, further amplifying the probe will deplete the pump, lowering the beat amplitude. In this pump-depletion regime [42], the probe amplitude grows linearly, rather than exponentially.

In addition to the real instabilities present in physical, i.e. real world, systems that we seek to model, simulations may be subject to numerical instabilities. One well known example arises in explicit finite-difference methods applied to partial differential equations, such as the Yee field solver. In order for the solutions to converge, the Courant condition [65] must be satisfied, which requires that the step size in the domain be below a certain threshold (for the Yee solver in 1D, the condition is $\Delta t < \Delta z/c$, with Δt and Δz the time and spatial step size, and c the vacuum speed of light). Failure to satisfy this condition will lead to nonphysical solutions that grow exponentially.

While this instability can lead to nonphysical behaviour, the threshold for stability is well-known, and it is typically clear that the results are in error, e.g. exponential growth of a single laser pulse propagating through a medium. Perhaps more dangerous are numerical instabilities that go unnoticed, either changing the simulation conditions, or being mistaken for physical processes.

Grid heating, also known as Debye heating, provides a good example of such a process [60]. In PIC simulations, the position and velocity of particles

are calculated to machine precision. However, the introduction of a grid to calculate the forces acting upon them results in a simulation resolution determined by the grid size. The grid must be sufficiently fine to resolve the physics of the processes being modelled. Grid heating occurs when the Debye length of the plasma is not resolved, manifesting as exponential heating. This numerical instability saturates when the plasma temperature becomes sufficiently high that the Debye length is resolved by the grid.

The nonphysical increase in temperature associated with grid heating can have an impact on many of the real processes we seek to model, for example the resonant frequency of the plasma (discussed in more detail in Chapter 4). The increase in temperature can also act to artificially increase the Landau damping rate, leading to further heating. Landau damping is a real process, but in this case it is enhanced by the presence of numerical instabilities. We refer to such processes as pseudo-physical, processes that exist in real systems but are artificially increased in simulations due to nonphysical effects.




Figure 3.2: **Example of pump and probe evolution.**

Pump and probe amplitude ($eE/mc\omega_p$) against position taken at different times. Simulation parameters as in Table (3.1).

3.2.2 Numerical stability

In laser-plasma interactions for cold plasma, the smallest spatial scale is that of the laser phase. For example, in the case of wakefield acceleration, a resolution of 20 cells per laser wavelength is typically sufficient to model the process [63]. The smallest scale in the Raman interaction is that of the excited plasma wave. As the wavelength of the plasma wave is approximately half that of the laser pulses, it might be expected that simulations require a resolution approximately 40 cells per laser wavelength.

Simulations have been carried out using the PIC code OSIRIS [2]. An example showing the envelopes of the pump and probe at different stages during their interaction, separated by post-processing techniques, is shown in Fig. (3.2). A second pulse is seen to develop behind the head of the probe, consistent with the semi-analytical results obtained in Section 2.3. Circularly polarised pulses are used, with amplitudes normalised to $e|E|/mc\omega_p$, where

Pump	Intensity Duration Wavelength	$1 \times 10^{14} \text{ W cm}^{-2}$ 5.46 ps 800 nm
Probe	Intensity Duration Wavelength	$1 \times 10^{12} \text{ W cm}^{-2}$ 500 fs* Chosen for optimal detuning
Plasma	Density Length Initial temperature	$2 \times 10^{19} \text{ cm}^{-3}$ 1 mm Cold plasma
Simulation	Resolution Macroparticle density Interpolation	100 cells per pump wavelength 5 000 per pump wavelength Linear

*The pulse uses the default OSIRIS pulse shape, a fifth order polynomial. The FWHM amplitude is 500 fs, with a full-width of 1 ps.

Table 3.1: Typical parameters used for simulations

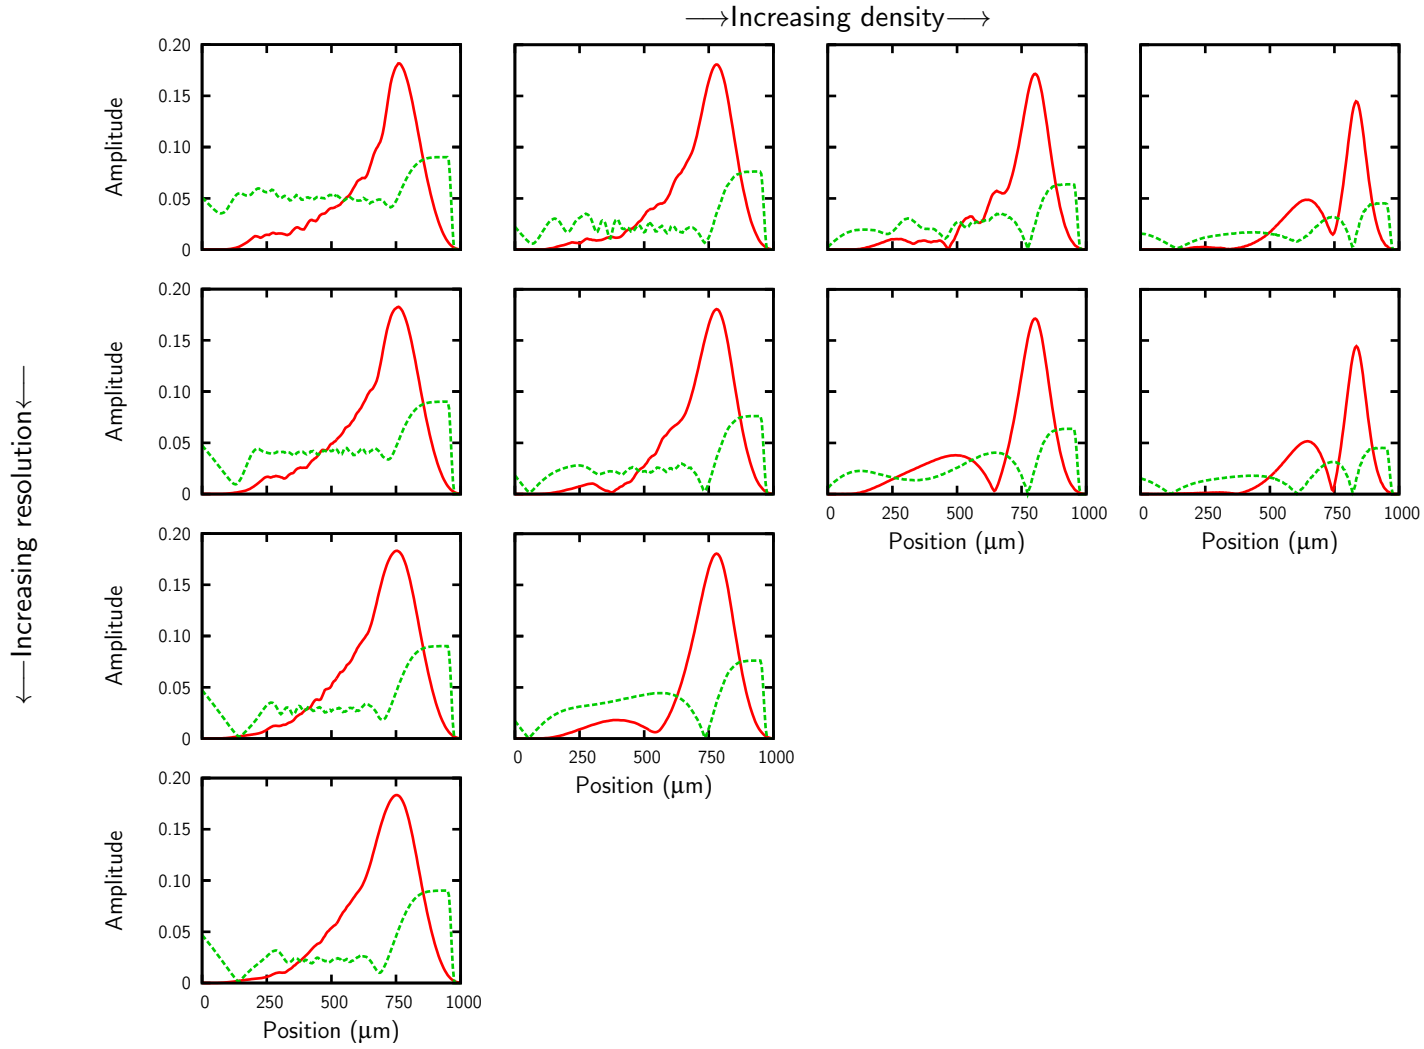


Figure 3.3: **Influence of plasma density on numerical stability.**

Pump and probe amplitude ($eE/mc\omega_p$) for simulations carried out with varying plasma density and simulation resolution. Plasma density (from left to right) of 1×10^{19} , 1.4×10^{19} , 2×10^{19} and 4×10^{19} cm⁻³. Resolution (from top to bottom) of 50, 100, 200 and 400 cells per pump wavelength. Other parameters as in Table (3.1).

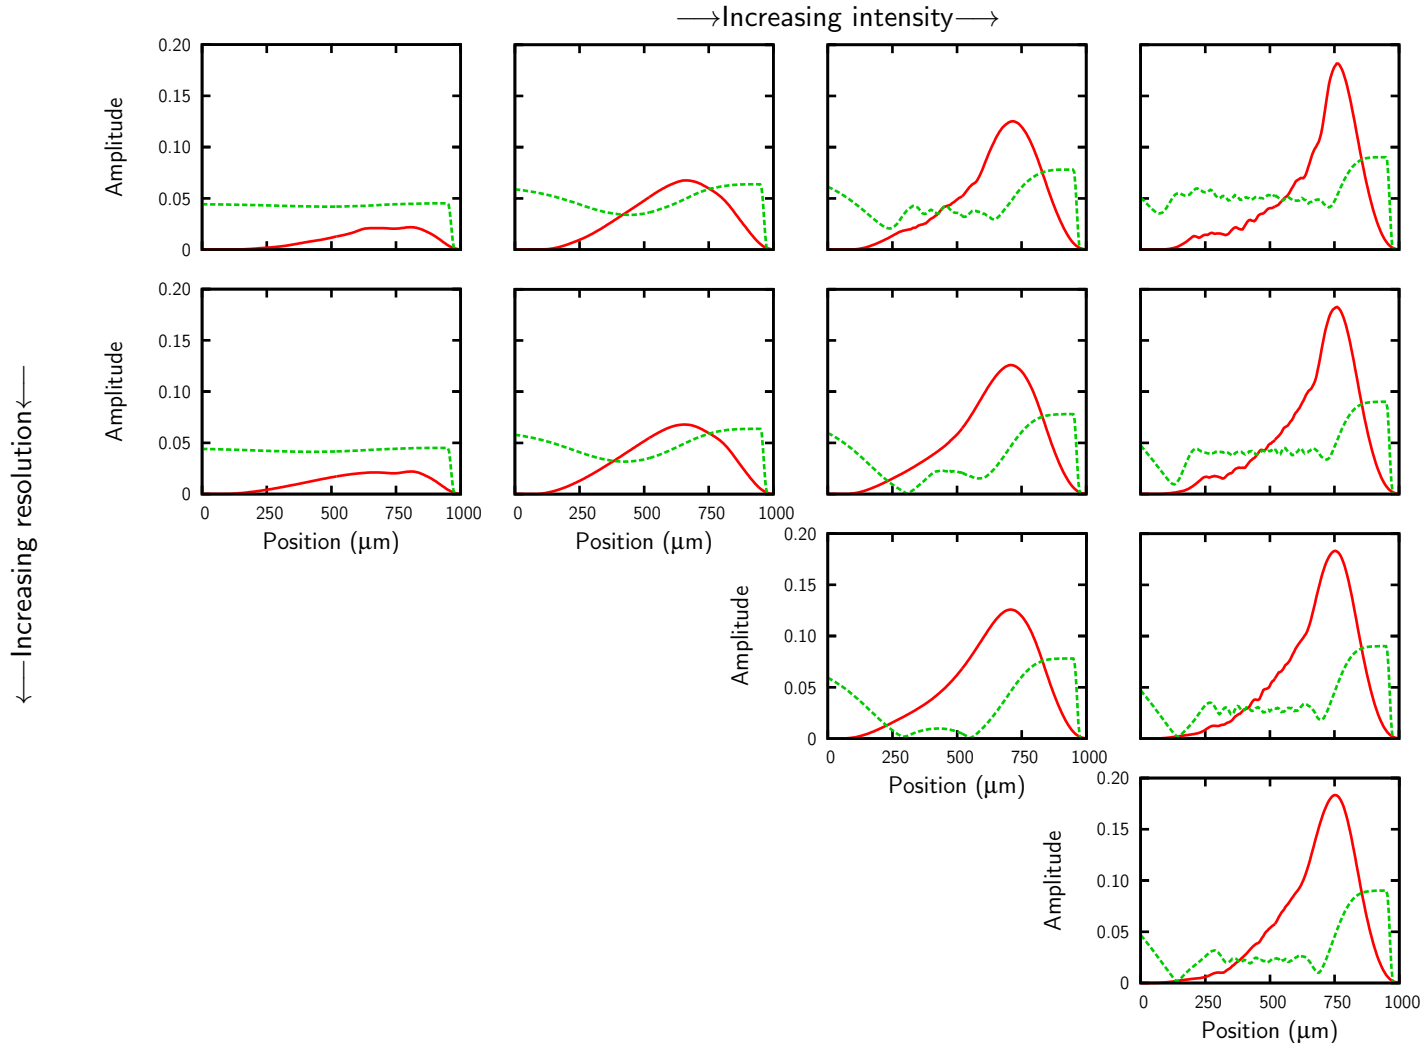


Figure 3.4: **Influence of pump intensity on numerical stability.**

Pump and probe amplitude ($eE/mc\omega_p$) for simulations carried out with varying pump intensity and simulation resolution. Pump intensity (from left to right) of 2.5×10^{13} , 5×10^{13} , 7.5×10^{13} and 1×10^{14} W cm⁻². Resolution (from top to bottom) of 50, 100, 200 and 400 cells per pump wavelength. Plasma density of 1×10^{19} cm⁻³, other parameters as in Table (3.1).

$-e$ and m are the electron charge and mass, E the laser electric field strength, and ω_p the plasma frequency. Typical simulation parameters are shown in Table 3.1. Parameters have been chosen to be similar to those used in the experiment by Ren *et al* [47], which, as shown in Figure 3.2, allow significant growth to be observed over a relatively short interaction length. The length of the pump was chosen to allow the probe to interact with the full length of the pump.

Fig. (3.3) shows the envelopes of the pump and amplified probe for different plasma densities using different grid resolutions. To provide a direct comparison, the total number of macroparticles is the same for all cases. It can be seen that at lower densities, a higher resolution is required for the simulation results to converge. At the lowest density shown, $1 \times 10^{19} \text{ cm}^{-3}$, a resolution of 400 cells per wavelength is not sufficient to reach a convergent solution. Fig. (3.4) shows the effect of varying the pump amplitude on the resolution stability. In this case, a higher resolution is required to achieve convergent results at high intensity.

A common feature in both these parameter scans is that higher resolution is required for plasma waves with a larger fractional density modulation. Simulations of cold plasma are not susceptible to grid heating [60], and so this behaviour is attributed to the development of spatial harmonics of the plasma wave as it reaches high amplitudes [66], which require a finer grid to

be resolved. For the lowest plasma density/highest pump amplitude, simulations using a fluid model (described in Chapter 4) suggest the plasma wave reaches 90% of the wavebreaking amplitude, leading to a strongly anharmonic structure.

The phase-space distribution for a selection of electrons over a short sample is shown in Fig. (3.5). As can be seen, at low resolution, there is an increase in the plasma temperature, with some macroparticles escaping the ordered oscillations of the plasma altogether. In analytical treatments of cold plasma, particles can escape the ordered oscillations of the wave due to wavebreaking, which corresponds to particle crossing [66]. We therefore label this nonphysical instability as “numerical wavebreaking”.

Plots of the longitudinal electric field corresponding to the phase space show a significant decrease in the amplitude of the plasma wave. This decrease in amplitude may be due to pseudo-physical effects caused by the increase in temperature, such as Landau damping of the plasma wave, or the Bohm-Gross shift in the plasma resonance. These processes are discussed in the case of physical systems in Chapter 4. The decrease in amplitude may also be directly related to the “numerical wavebreaking” instability, as particles escaping from the plasma wave will remove energy from it, as occurs due to wavebreaking in physical systems.

The presence of pseudo-physical processes make direct study of the nu-




Figure 3.5: **Influence of simulation resolution on numerical stability.** Sample macroparticle phase space and corresponding electrostatic field ($eE/mc\omega_p$) over a short region for varying simulation resolution. Resolution (from top to bottom) of 50, 100 and 200 cells per pump wavelength. Density of $1.4 \times 10^{19} \text{ cm}^{-3}$, other parameters as in Table (3.1).

numerical instability difficult. It is also possible that different numerical instabilities may couple together. For example, simulations of cold plasma are not susceptible to grid heating, but it may occur as the result of heating by another numerical instability. To determine if grid heating plays a significant role, a simulation has been carried out for cold plasma at low resolution, in which non-convergent behaviour may be seen, and another at high resolution, for which a convergent solution is achieved, but with an initial temperature chosen to match the saturation temperature for grid heating in the low-resolution case. Spontaneous backscatter ahead of the probe, seen for the case of initially warm plasma, is discussed in Section 3.2.4. grouped into bins of 100. The temperature over a region was calculated by taking the average of the standard deviation of all bins in the region. This gave a temperature over the 500-505 μm region of 194 eV for the low resolution simulation, compared to 165 eV for the high resolution case. The high resolution case gives a lower final temperature than the low resolution case, which shows the “numerical wavebreaking” observed here can exceed the maximum temperatures attainable through grid heating.

This numerical instability due to anharmonic plasma waves is similar to that observed by Cormier-Michel *et al* in the case of simulations of wake-field acceleration [56], which led to significant heating of cold plasma. Both instabilities arose in the context of high amplitude plasma waves. Here, how-

ever, the instability is potentially more problematic, as even small changes in temperature can result in pseudo-physical processes, which can strongly impact on the Raman process, such as Landau damping and the Bohm-Gross shift in plasma resonance. In wakefield acceleration, a much larger shift in temperature is required before pseudo-physical effects have a strong impact, as electron injection into the wake has a much smaller temperature dependence. However, unlike the result of Cormier-Michel *et al*, the macroparticle number does not have any discernible impact on the instability, as shown in Fig. (3.7). This may be due to the much larger macroparticle density used in this work.

As illustrated in Figs. 3.3,3.4, numerical wavebreaking can lead to a significant error in simulations of Raman amplification, modifying both the pump and probe pulses. As can be seen, this error is typically only significant some distance behind the leading edge of the probe, and, for the parameters used here, the peak amplitude of the amplified probe is correctly predicted at low resolution. At low densities/high intensities, however, the duration of the leading peak is underestimated at low resolution. While correctly simulating the intensity and duration of the leading pulse of the amplified probe is sufficient for many applications, in some cases the structure of the full pulse is important, for example in shock-ignition inertial confinement fusion [67]. Similarly, resolution of the depleted pump is necessary, for example, to gauge




Figure 3.6: **Comparison of “numerical wavebreaking” with grid heating.** Pump and probe amplitude ($eE/mc\omega_p$) and sample macroparticle phase space for different initial conditions. a) initially cold plasma with a resolution of 50 cells per pump wavelength, b) initial temperature 2.35 eV (equivalent to the saturation temperature for the grid heating instability for a resolution of 50 cells per pump wavelength) with a resolution of 400 cells per pump wavelength. Other parameters as in Table (3.1).




Figure 3.7: **Influence of macroparticle number on numerical stability.** Pump and probe amplitude ($eE/mc\omega_p$) and sample macroparticle phase space for different macroparticle densities: 5 000, 10 000 and 20 000 per pump wavelength. Grid resolution of 50 cells per pump wavelength, other parameters as in Table (3.1).




Figure 3.8: **Influence of smoothing techniques on numerical stability.** Pump and probe amplitude ($eE/mc\omega_p$) and sample macroparticle phase space for different smoothing techniques: no smoothing, smoothing of currents, and smoothing of fields. Parameters as in Table (3.1).

the viability of multi-pass amplification systems [47].

Smoothing techniques have been implemented for some applications, for example to reduce pseudo-physical spontaneous Raman backscatter, discussed in Sections 3.2.3, 3.2.4. Smoothing may be applied to currents or fields every timestep, to reduce the influence of noise. However, as shown in Fig. (3.8), the use of smoothing, in either case, leads to an increase in




Figure 3.9: **Influence of interpolation on numerical stability.**

Pump and probe amplitude ($eE/mc\omega_p$) and sample macroparticle phase space for different levels of particle interpolation: linear, quadratic and quartic. Parameters as in Table (3.1).

numerical heating, and a decrease in simulation stability. This is likely due to the fact that the smoothed fields provide a less accurate representation of the anharmonic wave, which is sharply peaked, and so a higher resolution is required to attain the same level of accuracy.

To deposit particle currents onto the grid, and apply grid fields to individual particles, interpolation methods must be used. Higher order interpolation is known to suppress grid heating [56]. However, as shown in Fig. (3.9), the use of higher order interpolation acts to increase the numerical wavebreaking instability. This is attributed to higher-order interpolation leading to a “smearing-out” of individual macroparticles, reducing the accuracy with which the anharmonic plasma wave is modelled, just as smoothing does. While the influence is weaker than that due to smoothing, it leads to a significant increase in numerical effort while reducing the accuracy of the simulation. This increase in nonphysical heating is of interest, as it contrasts with grid heating, which is reduced by the use of higher order interpolation.

3.2.3 Model limitations

There are several instabilities that may limit the practicality of Raman amplification in plasma. One example is spontaneous Raman scattering, where the Raman process is seeded by a source of noise, e.g. a density perturbation arising due to the finite temperature of the plasma [68]. The spontaneously scattered signal then acts as a seed for further scattering, allowing it to be amplified exponentially. This can deplete the pump before it interacts with the probe pulse, reducing amplification, or impair the quality of the amplified probe through the introduction of a pre-pulse or pedestal.

As discussed in Section 3.2.1, numerical instabilities can lead to pseudo-physical processes, which exist in real systems but are artificially enhanced. Such effects can be avoided by suppressing the numerical instability that causes them. However, pseudo-physical processes can also be caused by the limitations of the simulation model.

Historically, PIC codes have been considered more vulnerable to instabilities that grow from noise, such as spontaneous Raman scattering, due to the relatively small number of macroparticles used [56, 60]. This increases the granularity of the simulated plasma, artificially increasing the size of density perturbations that may act to seed these processes. However, with modern computational resources, it is now feasible to simulate the physical number

of particles in lower dimensionality simulations, such as those carried out in this work.

Pseudo-physical instabilities, however, remain a concern, as PIC codes typically do not include collisional damping, which will result in an artificially high growth rate for Raman amplification, increasing the influence of spontaneous Raman scattering. Heating of the plasma by collisional damping, which has been suggested as a method to suppress spontaneous backscatter [44], is also not simulated.

In addition to the pseudo-physical instabilities that arise, the lack of collisional damping prevents PIC codes from accurately simulating Raman amplification in regimes where collisional processes are important, e.g. as shown in Section 4.8.

3.2.4 Stable parameter spaces

While PIC codes are not appropriate simulation tools for some parameter spaces due to the absence of collisional effects, they provide a useful tool for benchmarking other models in regimes where they give accurate results. We therefore investigate how pseudo-physical backscatter may be suppressed, to provide a simulation space in which comparisons may be carried out. As the enhancement of spontaneous backscatter arises due to a model limitation, rather than a numerical instability, it is here suppressed using the same methods as could be applied in physical systems.

Fig. (3.10) shows the pump and amplified spontaneous Raman backscatter for different plasma densities. It can be seen that backscatter decreases with decreasing density. This is attributed to both a decrease in amplification and an increase in Landau damping at lower densities. The latter arises due to the decrease in phase velocity of the resonantly excited plasma wave at low density, resulting in more particles with velocities close to the phase velocity of the plasma wave. This role of Landau damping in suppressing noise, which has been suggested as a stabilising factor in physical systems [44], is also observed in Fig. (3.11), as less backscatter is observed for higher temperatures, where damping is stronger. Landau damping is discussed in more detail in Chapter 4.

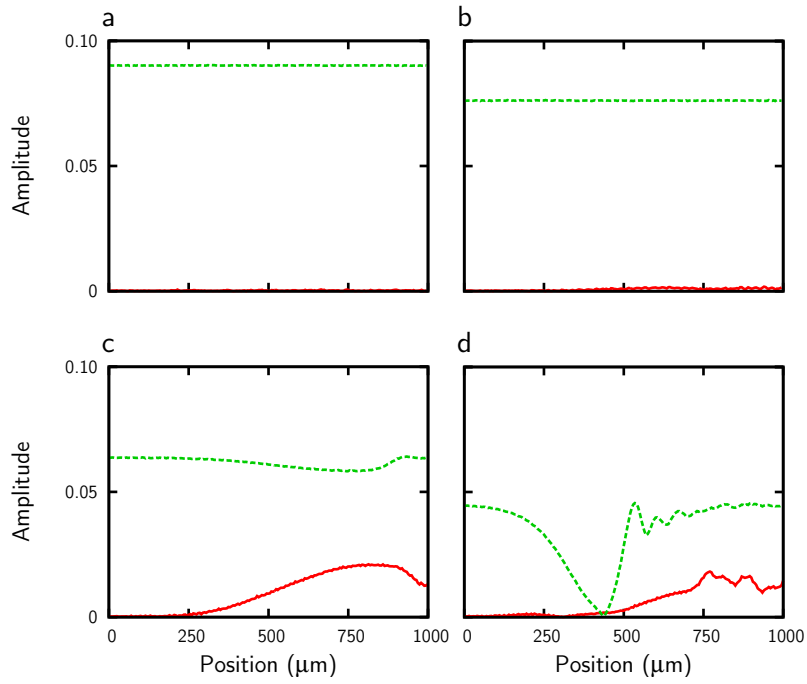


Figure 3.10: **Influence of density on spontaneous backscatter.**

Pump and probe amplitude ($eE/mc\omega_p$) for simulations carried out with different plasma densities. Density of a) $1 \times 10^{19} \text{ W cm}^{-2}$, b) $1.4 \times 10^{19} \text{ W cm}^{-2}$, c) $2 \times 10^{19} \text{ W cm}^{-2}$ and d) $4 \times 10^{19} \text{ W cm}^{-2}$. Temperature of 100 eV, other parameters as in Table (3.1).

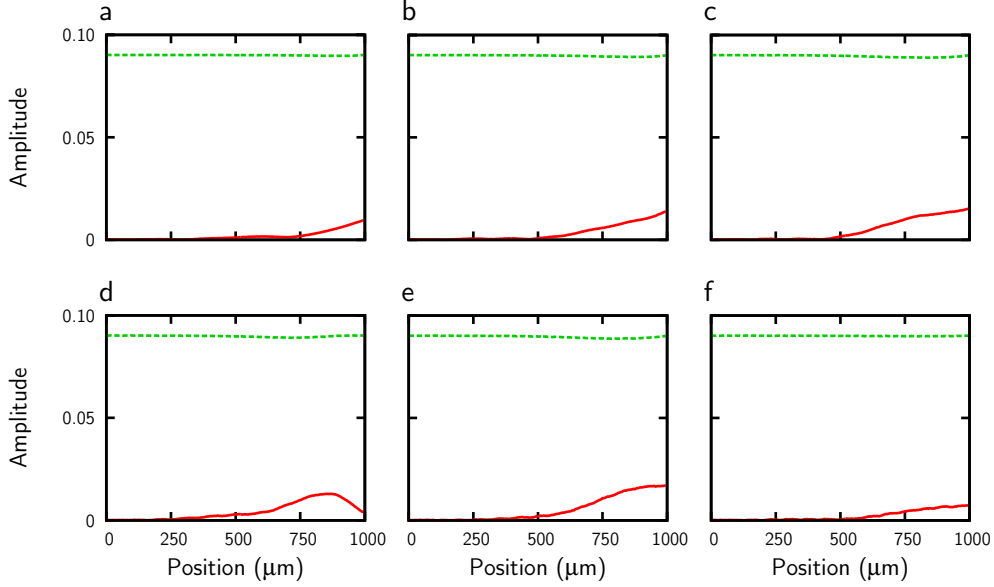


Figure 3.11: **Influence of temperature on spontaneous backscatter.** Pump and probe amplitude ($eE/mc\omega_p$) for simulations carried out with different plasma temperatures. Temperature of a) 10 eV, b) 20 eV, c) 30 eV, d) 40 eV, e) 50 eV and f) 60 eV. Density of $1 \times 10^{19} \text{ cm}^{-3}$, other parameters as in Table (3.1).

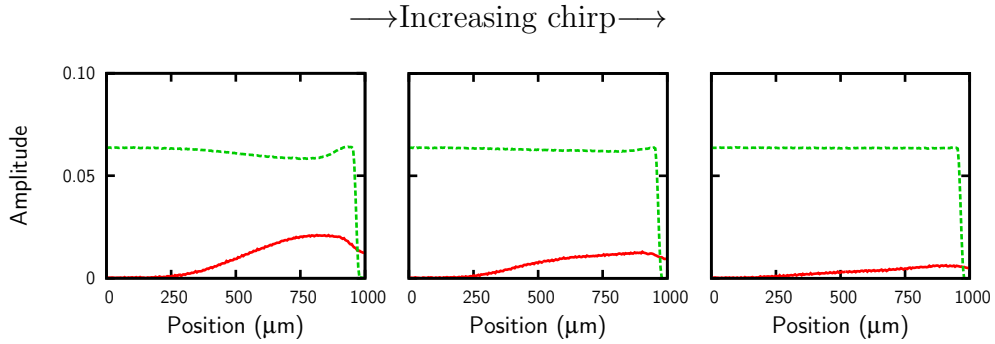


Figure 3.12: **Influence of pump chirp rate on spontaneous backscatter.** Pump and probe amplitude ($eE/mc\omega_p$) for simulations carried out with different pump chirp rates. From left to right, pump bandwidth as a fraction of pump frequency, $\Delta\omega_0/\omega_0 = 0, 0.002$ and 0.004 . Density of $1 \times 10^{19} \text{ cm}^{-3}$, temperature of 50 eV, other parameters as in Table (3.1).

The use of a chirped pump pulse has also been suggested as a method to overcome noise [68], as the growth rate becomes linear for small signals, rather than exponential [43]. Fig. (3.12) shows decreasing backscatter with increasing chirp-rate.

3.2.5 Computational overheads

In Raman amplification, two counterpropagating laser pulses are made to collide in plasma. For the case of linearly polarised pulses, the beatwave will have two components: one with high frequency, and one with low frequency (for circularly polarised light, only one of these components will exist, determined by the relative polarisation of the laser pulses). For the correct detuning, the difference frequency matches the resonant frequency of the plasma electrons, allowing the laser beat to resonantly drive a plasma wave.

However, due to the counterpropagating geometry, this low frequency component will have the difference of the laser frequencies and the sum of the laser wavenumbers, i.e. it varies rapidly in space, but slowly in time. The Yee solver, or indeed any explicit solution to a set of partial differential equations, may be considered an inefficient choice in this case, as satisfying the Courant condition requires that the same high resolution be used for both the temporal and spatial response of the plasma wave, leading to a large computational overhead.

The presence of numerical instabilities, such as the “numerical wavebreaking” instability identified in Section 3.2.2, can act to further increase the resolution required to achieve convergent solutions. This can lead to significant increases in the computational overhead, as doubling the spatial resolution

to better resolve the harmonic structure of the plasma wave requires the number of timesteps to be doubled, increasing computational overhead by a factor of four.

Moving window geometries are widely used for PIC simulations of wakefield acceleration in order to reduce the computational overhead of simulations [63]. For a stationary window, simulation time increases with the square of the laser propagation distance, while for a moving window the dependence is linear. A similar model could be applied to Raman amplification, as the amplified pulse is typically short. A moving window would also act to reduce the influence of noise, by reducing the time that parasitic instabilities of the counterpropagating pump can develop. To allow a counterpropagating pulse, without the introduction of nonphysical effects, requires that the boundary condition at the leading edge of the window be modified to give not only the correct electromagnetic field, but also the correct electron response. At the time this work was undertaken, neither OSIRIS nor Vorpal supported moving windows that allow the use of a counterpropagating laser pulse.

As discussed in Section 3.2.4, the inclusion of collisional effects would remove the pseudo-physical increase in Raman growth rates. While not currently supported, the inclusion of such effects would lead to a significant increase in the computational effort required to complete a simulation.

3.3 Conclusion

Modelling Raman backscatter using PIC codes requires high resolution, as the short-wavelength plasma wave must be resolved. This work shows that as the plasma wave amplitude increases, higher simulation resolution is required due to the development of harmonics of the wave. Failure to use a sufficiently high resolution leads to a numerical instability, characterised by nonphysical heating, with some particles escaping the collective motion of the plasma wave, and nonphysical damping of the wave. We name this instability “numerical wavebreaking”, in analogy with wavebreaking in physical systems, which exhibits the same behaviour of electrons escaping the collective motion of the wave, leading to damping.

The use of numerical smoothing for either the electromagnetic fields or currents is shown to increase the effect of this numerical instability. This is consistent with the instability arising due to anharmonic plasma waves, which become sharply peaked at high amplitude, and so are less well modelled by smoothed fields. Higher-order interpolation is shown to have a similar, although weaker, effect, attributed to the smoothing which higher order interpolation introduces.

In cases where collisions are not modelled, the applicability of simulations is limited to cases where collisional effects do not play a significant role. The

absence of collisional damping also acts to artificially increase the growth rate of instabilities such as spontaneous Raman backscatter, which acts to prematurely deplete the pump pulse. Smoothing cannot be used to reduce the level of backscatter due to its role increasing numerical instability. Neither OSIRIS nor Vorpal currently support collisions.

The high resolution required to model Raman amplification in plasma, further increased by the presence of numerical instabilities, leads to high computational overheads. Neither OSIRIS nor Vorpal currently support moving-window geometries compatible with Raman amplification, and so the whole plasma must be modelled over the simulation, leading to prohibitive computational overheads for long interaction lengths. Should collisions be supported in future, they would lead to a significant further increase in computational overhead.

The ongoing experiments at Strathclyde use plasma capillaries that are several centimetres in length, in parameter spaces where collisional damping and heating are expected to play a significant role. The high resolution required to model Raman backscatter would result in computational overheads of tens-of-thousands of CPU hours, while failing to simulate collisional effects. Although conventional PIC codes remain an important simulation tool for the study of laser-plasma interactions, we find that in their current state of development, with current computational resources, they are unsuitable for

the simulation of the Strathclyde experiments.

In other sections of this work, PIC codes will be used as a benchmark for the alternative simulation tools investigated. Parameter spaces are identified in which conventional PIC simulations are stable. Numerical wavebreaking can be avoided through the use of high resolution, and by avoiding parameters which result in high amplitude plasma waves. Pseudo-physical Raman backscatter can be controlled by using low plasma densities, or a chirped pump.

Chapter 4

Three-wave models:

wavebreaking and harmonics

Chapter Abstract

The three-wave model, which is widely used in the simulation of Raman amplification in plasma, is extended to include a phenomenological treatment of harmonics, wavebreaking, and spectral broadening due to particle trapping.

The harmonic content and wavebreaking limit of the plasma wave in cold plasma are readily calculated from analytical results in the literature. For warm plasma, a model based on modified waterbag distributions is developed to estimate the fraction of trapped electrons, without the need to model the evolution of the entire electron velocity distribution. A simple model is developed to estimate the impact of these trapped electrons on the wavenumber spectrum of the plasma wave.

Benchmarking is carried out against a particle-in-cell (PIC) code across a range of regimes, which affirm the model. The removal of common numerical approximations allows much better agreement for the case of cold plasma than seen in previous work. It is shown that the inclusion of spectral effects greatly improves the agreement between the three-wave model and PIC simulations in warm plasma. Agreement is limited by the use of

linearised Landau damping rates, which become invalid at high temperatures and/or low densities.

The low computational overhead allows large simulations to be carried out which would otherwise have been infeasible. This allows either long propagation distances or large parameter scans to be carried out. The model has the additional advantage of including collisional effects, which are typically not included in PIC models, and are shown to have a significant impact on Raman amplification in some situations.

As the current experimental work at Strathclyde is in a regime in which nonlinear Landau damping occurs, further development is required before the model may be used for these simulations. The model is, however, applicable to experimental work carried out elsewhere, in which higher plasma densities have been used.

4.1 Introduction

Three-wave models are typically derived from a fluid description of plasma and have long been used in the study of Raman amplification [29, 42, 43, 69]. They offer the advantage that they have a very low computational overhead, with a typical simulation in this work taking tens of seconds on a single processor, rather than the tens or hundreds of processor hours that a PIC simulation would require. The simplicity of the model also offers other benefits, as, in some cases, it allows analytical solutions to be found [42, 43].

As they are derived from a fluid model, the evolution equations cannot explicitly model kinetic processes, such as wavebreaking and Landau damping. However, the simplicity of the model allows such processes to be included phenomenologically [3, 69], while maintaining the advantages of simplicity and low computational overhead. Indeed, this offers the additional advantage that the contribution of a single physical process may be easily evaluated by excluding it from the model, which may not be possible in a PIC simulation. Furthermore, the simplicity of the three-wave model suggests that it may be more readily extendable than a PIC model, for example by the inclusion of collisional effects.

Due to the relatively low plasma densities used in the capillary experiments at Strathclyde, we are interested in evaluating the impact of thermal

effects, high amplitude plasma waves, particle trapping and wavebreaking. The influence of damping and the shift in the plasma resonance due to thermal effects have been studied in previous works [3, 69, 70]. This work extends the general model used in these works through the phenomenological inclusion of harmonic content and wavebreaking of the plasma wave for case of cold plasma, and the impact of spectral broadening due to particle trapping for the warm plasma case.

This work is developed as follows: the three-wave model for cold plasma is extended to include harmonics and wavebreaking; the model for warm plasma is then extended to calculate the fraction of trapped electrons, and model the resulting impact on the wavenumber spectrum of the plasma wave; the limitations of the models are discussed, and benchmarking carried out against PIC simulations; simulations are carried out to illustrate the applicability of the model, conclusions drawn, and possible future work outlined.

4.2 Cold plasma

4.2.1 Three-wave model

The evolution of the pump, probe and Langmuir wave involved in the Raman interaction can be described using a set of slowly-varying envelope equations.

For the case of cold, nonrelativistic plasma¹, these are:

$$\begin{aligned} \left(\frac{\partial}{\partial t} - v_g \frac{\partial}{\partial z} \right) a_0 &= -\omega_p \frac{\omega_0 + \omega_1}{4\omega_0} f^* a_1, \\ \left(\frac{\partial}{\partial t} + v_g \frac{\partial}{\partial z} \right) a_1 &= \omega_p \frac{\omega_0 + \omega_1}{4\omega_1} f a_0, \\ \left(\frac{\partial}{\partial t} + i(\omega_0 - \omega_1 - \omega_p) \right) f &= \frac{\omega_p(\omega_0 + \omega_1)}{2(\omega_0 - \omega_1 + \omega_p)} a_0^* a_1. \end{aligned} \quad (4.1)$$

Here, $|a_{0,1}| = (e/mc\omega_{0,1})\sqrt{2I_{0,1}/\varepsilon_0 c}$ are the reduced vector potentials of the pump and probe, respectively, with $I_{0,1}$ and $\omega_{0,1}$ the associated intensities and frequencies. $f = eE_z/m\omega_p c$ is the normalised plasma wave electric field amplitude, E_z , with $\omega_p = \sqrt{ne^2/\varepsilon_0 m}$ the plasma frequency. v_g is the laser group velocity, n the plasma electron density, $-e$ and m the electron charge and mass, ε_0 the permittivity of free space and c the vacuum speed of light.

We expect that harmonics and wavebreaking of the plasma wave, discussed in the following sections, will reduce the Raman growth rate. We

¹The electrons are assumed to be in the steady-state with respect to the instantaneous laser field (quasi-static approximation), the ions are assumed to be stationary, and dispersive effects are neglected. Full derivations are shown in Appendix A.2.

modify the pump and probe evolution equations to include a factor μ , which denotes the fractional amplitude of the plasma wave that contributes to Raman amplification:

$$\begin{aligned} \left(\frac{\partial}{\partial t} - v_g \frac{\partial}{\partial z} \right) a_0 &= -\omega_p \frac{\omega_0 + \omega_1}{4\omega_0} \mu f^* a_1, \\ \left(\frac{\partial}{\partial t} + v_g \frac{\partial}{\partial z} \right) a_1 &= \omega_p \frac{\omega_0 + \omega_1}{4\omega_1} \mu f a_0. \end{aligned} \quad (4.2)$$

We develop models to calculate the contributions to μ due to harmonics and wavebreaking in Sections 4.2.2 and 4.2.3, respectively. Benchmarking of the extended model is carried out in Section 4.7.1.

4.2.2 Nonlinear plasma waves

At high amplitudes, the plasma wave develops spatial harmonics, while the resonant frequency remains constant [66]. This arises because the amplitude of an electron oscillation depends on its equilibrium position, while the restoring electrostatic force depends on its displacement, allowing the density to become sharply peaked.

The Raman process requires that both frequency and wavenumber matching be satisfied for efficient amplification to occur: $\omega_0 - \omega_1 = \omega_p$, $k_0 + k_1 = k_p$. These relations are a consequence of conservation of energy and momentum. They are readily understood by considering the plasma electrons as inde-

pendent radiators driven by the pump field: if they are not bunched on the correct scale, the radiation emitted by different electrons will interfere destructively; if the bunching does not have the correct phase velocity, the emitted radiation will interfere destructively with the probe.

We therefore expect that the generation of harmonics will reduce the amplification rate, as harmonics arise in k_p but not ω_p . An analytical treatment of the relative amplitude of these harmonics has been carried out by Khachatryan and Elbakian [71]. The amplitude of the n th harmonic of f , f_n , is given by:

$$f_n = 2 \frac{J_n(n|f|c/v_\phi)}{n^2}, \quad (4.3)$$

where $v_\phi = \omega_p/k_p$ is the phase velocity of the plasma wave, and J_n is the n th order Bessel function of the first kind.

Neglecting harmonics as nonresonant, we may set $\mu = \mu_h$, the relative amplitude of the plasma wave fundamental:

$$\mu_h = 2 \frac{J_1(|f|c/v_\phi)}{|f|c/v_\phi}. \quad (4.4)$$

4.2.3 Wavebreaking

The above equations are incomplete because they place no limit on the amplitude to which the plasma wave may grow. For cold plasma, wavebreaking occurs when the oscillation velocity of electrons in the plasma wave exceeds the phase velocity of the wave, i.e. $fc > v_\phi$ [66]. This is a consequence of electrons gaining sufficient momentum to escape the collective behaviour of the plasma wave. Envelope models are typically unsuitable for the study of wavebreaking, because while they (and the fluid models from which they are derived) may be used to calculate the amplitude at which a wave will break, the process itself is a kinetic effect.

A phenomenological treatment of wavebreaking is included in the model through the modification of the factor μ used in Eqs. (4.2), to $\mu = \mu_h \mu_b$, where μ_b is the fraction of plasma electrons that contributes to the plasma wave, i.e. those that have not “broken”. In the simplest treatment:

$$\mu_b = \begin{cases} 1, & fc \leq v_\phi, \\ 0, & fc > v_\phi. \end{cases} \quad (4.5)$$

The wavebreaking amplitude corresponds to electrons having sufficient energy to escape the collective motion of the plasma wave. Electrons can then “spill over” into the next bucket of the wave, removing energy from the

oscillation. However, this behaviour can only occur when the electron reaches the edge of the wave bucket it occupies. It will therefore take some finite time after wavebreaking before all electrons escape the bucket and the wave is fully depleted. Assuming electrons are initially uniformly spread across all possible phases, the wave amplitude can be expected to decrease linearly in time, becoming fully damped-out after a plasma period, $T_p = 2\pi/\omega_p$.

If t_{WB} is the time at which the wavebreaking amplitude is reached (which can be found through simulation), μ_b is then given by:

$$\mu_b = \begin{cases} 1, & t \leq t_{\text{WB}}, \\ 1 - (t - t_{\text{WB}})/T_p, & t_{\text{WB}} < t < (t_{\text{WB}} + T_p) \\ 0, & t \geq (t_{\text{WB}} + T_p). \end{cases} \quad (4.6)$$

4.3 Warm plasma

4.3.1 Three-wave model

The evolution equations in Section 4.2 may be augmented to include thermal effects (e.g. [69]) by including the influence of damping and the modified plasma response at finite temperature.

The treatment of the case of warm plasma differs from that of cold plasma in two respects: firstly, collisional (inverse-bremsstrahlung) damping serves to reduce the ordered motion of electrons; secondly, the electron velocity \vec{v} is no longer singularly defined at every point, and is instead described by a distribution function, $\mathcal{F}(\vec{v})$.

As in Section 4.2.1, we limit our analysis to the planar, nonrelativistic case. We define $\nu_{ei}[\omega]$ as the electron-ion collision rate for a wave with frequency ω . For a Maxwellian distribution [72]:

$$\nu_{ei}[\omega] = \frac{Ze^2\omega_p^2}{6(2\pi)^{3/2}\varepsilon_0 m} \left(\frac{k_B T}{m}\right)^{-3/2} \left(\text{Ln} \left(\frac{4k_B T}{\hbar\omega} \right) - \gamma \right), \quad (4.7)$$

where Ze is the ion charge, k_B the Boltzmann constant, T the temperature, \hbar Plank's constant and γ Euler's constant.

We make the quasi-static approximation (as in Appendix A.2), assuming that the electron velocity, \vec{v} , may be described as the sum of a slowly-varying

component, \vec{u} , and the fast oscillatory response to the laser field, v_{os} , which may be treated as being in the steady-state with respect to the instantaneous laser field. For a collisionless plasma, the latter takes the form $v_{os} = \vec{a}c$, where \vec{a} is the reduced vector potential of the laser field. In the presence of collisions, the total electron velocity is then $\vec{v} = \vec{u} + (1 + i\nu_{ei}[\omega_l]/\omega_l)\vec{a}c$, where ω_l is the laser frequency. The electron velocity distribution is $\mathcal{F}(\vec{v}) = \mathcal{F}(\vec{u} + (1 + i\nu_{ei}[\omega_l]/\omega_l)\vec{a}c)$.

The motion of electrons, described in the cold plasma case by the Lorentz equation, Eq. (A.22), is now given by the Vlasov equation:

$$\begin{aligned}
& \frac{\partial \mathcal{F}(\vec{v})}{\partial t} + \vec{v} \cdot \nabla \mathcal{F}(\vec{v}) + \frac{d\vec{v}}{dt} \cdot \frac{\partial \mathcal{F}(\vec{v})}{\partial \vec{v}} = 0, \\
& \frac{\partial \mathcal{F}(\vec{v})}{\partial t} + \vec{v} \cdot \nabla \mathcal{F}(\vec{v}) - \frac{e}{m} \left(1 + i \frac{\nu_{ei}[\omega'_p]}{\omega'_p} \right) (\vec{E} + \vec{v} \times \vec{B}) \cdot \frac{\partial \mathcal{F}(\vec{v})}{\partial \vec{v}} = 0, \\
& \frac{\partial \mathcal{F}(\vec{u})}{\partial t} + \vec{u} \cdot \nabla \mathcal{F}(\vec{u}) \\
& \quad = \left(1 + i \frac{\nu_{ei}[\omega'_p]}{\omega'_p} \right) \left(\frac{e}{m} \nabla \varphi - \left(1 + i \frac{\nu_{ei}[\omega_l]}{\omega_l} \right) \frac{c^2}{2\gamma} \nabla |\vec{a}|^2 \right) \cdot \frac{\partial \mathcal{F}(\vec{u})}{\partial \vec{u}},
\end{aligned} \tag{4.8}$$

where ω'_p is the resonant frequency of the plasma wave, and φ the electric potential. We have removed oscillatory components as in Eq. (A.25), and have included the effects of collisional damping.

The plasma response can be calculated by following the method used by Landau [73]. The result is a shift in the resonant frequency of the plasma, the

Bohm-Gross shift, and an additional damping of the plasma wave, Landau damping.

The Bohm-Gross shift may be considered as a consequence of the electron pressure at finite temperatures, which adds an additional restoring force to any perturbation to the electron density, leading to an increase in the resonant plasma frequency. Landau damping is analogous to the process of wakefield acceleration. Electrons with thermal velocities close to the phase velocity of the plasma wave exchange energy with the wave. Whether electrons gain or lose energy depends on their phase relative to the wave, but the net effect is for electrons with velocity $v < v_\phi$ to gain energy, while those with $v > v_\phi$ lose energy. For Maxwellian distributions, then, this serves to damp the plasma wave.

For long wavelengths or low temperatures, i.e. regimes in which particle trapping may be neglected, good approximation to the results found by Landau can be made using more intuitive methods. Bohm and Gross [74] analysed the interaction of a Maxwellian distribution with a sinusoidal potential of wavenumber k , and found that the resonant frequency for the plasma wave is:

$$\omega'_p = \left(\omega_p^2 + 3 \frac{k_B T}{m} k^2 \right)^{\frac{1}{2}} . \quad (4.9)$$

The time-averaged Landau damping rate, ν_L may be calculated by analysing

the perturbation to the average electron velocity [75]. For a Maxwellian distribution, the result is:

$$\nu_L = \sqrt{\frac{\pi}{8}} \frac{\omega_p^2 \omega_p'^2}{k^3} \left(\frac{k_B T}{m} \right)^{-3/2} \exp \left(-\frac{\omega_p'^2}{2k} \left(\frac{k_B T}{m} \right)^{-1} \right). \quad (4.10)$$

Using these modified equations, and following the derivations in Appendix A.2, we arrive at the modified evolution equations for the three waves:

$$\begin{aligned} \left(\frac{\partial}{\partial t} - v_g \frac{\partial}{\partial z} + \left(\frac{\omega_p}{\omega_0} \right)^2 \frac{\nu_{ei}[\omega_0]}{2} \right) a_0 &= -\omega_p \frac{\omega_0 + \omega_1}{4\omega_0} f^* a_1, \\ \left(\frac{\partial}{\partial t} + c \frac{\partial}{\partial z} + \left(\frac{\omega_p}{\omega_1} \right)^2 \frac{\nu_{ei}[\omega_1]}{2} \right) a_1 &= \omega_p \frac{\omega_0 + \omega_1}{4\omega_0} f a_0, \\ \left(\frac{\partial}{\partial t} + \frac{(\nu_L + \nu_{ei}[\omega_p'])}{2} + i(\omega_0 - \omega_1 - \omega_p') \right) f &= \frac{\omega_p(\omega_0 + \omega_1)}{2(\omega_0 - \omega_1 + \omega_p')} a_0^* a_1. \end{aligned} \quad (4.11)$$

Assuming the plasma remains Maxwellian, the plasma is heated at a rate:

$$\frac{\partial}{\partial t} \left(\frac{k_B T}{m} \right) = \frac{c^2}{3} \left(\nu_{ei}[\omega_0] |a_0|^2 + \nu_{ei}[\omega_1] |a_1|^2 + (\nu_L + \nu_{ei}[\omega_p']) |f|^2 \right), \quad (4.12)$$

4.3.2 Coffey wavebreaking limit

In warm plasma, the wave amplitude at which individual electrons exceed the phase velocity of the plasma wave will depend on their position in the velocity distribution. As such, wavebreaking can be considered to occur at lower amplitudes than for the cold plasma case. However, we expect the behaviour in the warm wavebreaking regime to be significantly different from that discussed in Section 4.2.3 for cold plasma, as different parts of the distribution will exceed the phase velocity of the wave at different amplitudes. Indeed, some particles in the tail of the distribution have initial velocities higher than the phase velocity.

Coffey derived a wavebreaking limit for a warm (nonrelativistic) plasma using a waterbag model [76]. The initial electron velocities are chosen as a one-dimensional waterbag distribution with the same total energy as a Maxwellian distribution at the same temperature ($\int_{-\infty}^{\infty} v^2 \mathcal{F}(v) dv$ equal for the two distributions), such that:

$$\mathcal{F}(v)/\mathcal{F}(0) = \begin{cases} 0, & v < -\sqrt{3}v_{th}, \\ 1, & -\sqrt{3}v_{th} \leq v \leq \sqrt{3}v_{th} \\ 0, & v > \sqrt{3}v_{th}. \end{cases} \quad (4.13)$$

The waterbag boundaries are allowed to evolve as the plasma wave grows. When one of the boundaries reaches the phase velocity of the wave, the (in-

finite) gradient of the distribution will lead to strong Landau-type damping. The assumption is that at this amplitude strong damping will also occur for physical distributions. This wavebreaking amplitude, f_{WB} , is:

$$f_{\text{WB}} = \frac{v_\phi}{c} \left(1 - \frac{8}{3} \beta^{\frac{1}{4}} + 2\beta^{\frac{1}{2}} - \frac{1}{3} \beta \right)^{\frac{1}{2}}, \quad (4.14)$$

where $\beta = 3v_{th}^2/v_\phi^2 = 3(k_B T/m)\omega_p'/k$.

As for the model developed for cold plasma, we expect that some electrons continue to contribute to Raman amplification after the wavebreaking limit is reached. While the waterbag distribution leads to strong Landau-type damping at the Coffey wavebreaking amplitude, Landau damping at this wave amplitude for initially Maxwellian distributions is typically not sufficiently strong to suppress the Raman process, as will be shown in Section 4.7.3.

We re-evaluate the physical interpretation of wavebreaking for warm plasmas, moving away from the idea of a wave breaking, and towards the concept of trapped electrons modifying the plasma wave.




Figure 4.1: **Illustration of separatrix for warm plasma**
 Electron velocities in the wave frame, illustrating the separatrix about v_ϕ .

4.4 Particle trapping

Bohm and Gross [74] considered a distribution of electrons interacting with a potential, φ , moving at velocity v_ϕ . If we assume a steady state solution, the electron energy in the wave frame must be conserved, satisfying:

$$(v - v_\phi)^2 = (v_0 - v_\phi)^2 + 2e\varphi(x - v_\phi t)/m, \quad (4.15)$$

where v_0 is the electron velocity at $\varphi = 0$. We choose $\varphi = \varphi_0 \sin(k(x - v_\phi t))$, and note that electrons will have insufficient energy to reach a point x, t if $v_0^2 < -2e\varphi/m$. It is immediately clear that, for any sufficiently broad electron velocity distribution, the moving potential will lead to the formation of a separatrix centred about $v = v_\phi$, as shown in Fig. 4.1. Electrons will be trapped if $(v_0 - v_\phi)^2 < -2e\varphi_0/m$.

Bohm and Gross noted that, in the steady state, the density is inversely proportional to the electron velocity at that point. In the case where no

electrons are trapped, they arrive at Eq. (4.9), the plasma dispersion relation for warm plasma. In the case where electrons are trapped, they show that electron trapping serves to reduce the resonant plasma frequency. This is readily understood by considering that trapped electrons are localised around the troughs of the electron density modulation, reducing the restoring force on electrons in the wave. However, their solution assumes a distribution of trapped electrons peaked around $v = v_\phi$ (contrary to, e.g. [77]), and a single wavenumber for the plasma wave. In general, a full solution will result in a broadening of the wavenumber spectrum about k_p , in addition to the shift in the resonant frequency ω'_p .

As for the case of cold plasma, we assume that only the fundamental wavenumber satisfies the frequency and wavenumber matching conditions for Raman backscatter. Ideally, a full treatment should be carried out, by calculating both the resonant frequency and the relative amplitude of the fundamental wavenumber. However, such a treatment depends on the distribution function for trapped electrons, which cannot readily be calculated. Indeed, even if an assumed distribution were chosen, there are no useful tools to benchmark the results, because the distribution of trapped electrons in a PIC code will be unphysical due to the absence of collisions.

Such an investigation is therefore beyond the scope of this work, and we instead make use of a significantly simplified model. We note that electrons

trapped at the bottom of the trough make no resonant contribution to the density modulation, and assume that electrons at $v = \sqrt{v_\phi^2 + 2e\phi/m}$ make their full contribution. The contribution of untrapped electrons is essentially unchanged. While the velocities of electrons above the separatrix are reversed about v_ϕ compared to those below, their contribution to the wave depends only on their speed in the wave frame, which is equal. A linear interpolation gives the simple argument that the net contribution that trapped electrons make to the plasma wave modulation depth is halved. This gives the factor for the effective plasma wave modulation as $\mu = 1 - n_{\text{trapped}}/2n_0$.

4.4.1 Estimating the trapped electron density

Calculating the number of trapped electrons is not straightforward, as the electron velocity distribution will evolve from its initial form. We return to the warm wavebreaking model derived by Coffey: at the wavebreaking amplitude, the waterbag boundary has evolved from its initial value, from $\sqrt{3}v_{th}$ to v_ϕ . We proceed by making the assumption that an initially Maxwellian distribution would evolve in a similar manner: the distribution of electrons at v_ϕ at Coffey's warm wavebreaking limit is equal to the distribution of electrons at $\sqrt{3}v_{th}$ in the unperturbed Maxwellian distribution, i.e.

$$\mathcal{F}'(v_\phi) = \mathcal{F}_0(\sqrt{3}v_{th}), \quad (4.16)$$

where $\mathcal{F}_0(v)$ is the initial Maxwellian distribution, and $\mathcal{F}'(v)$ the distribution at wavebreaking.




Figure 4.2: **Maxwellian and waterbag distributions.**

The model assumes that for an initially Maxwellian distribution, $\mathcal{F}_0(v)$, the distribution at Coffey's warm wavebreaking limit, $\mathcal{F}'(v)$, satisfies $\mathcal{F}'(v_\phi) = \mathcal{F}_0(\sqrt{3}v_{th})$.

In order to calculate the distribution of electrons at v_ϕ for a plasma wave of arbitrary amplitude, we make use of a modified waterbag distribution with arbitrary boundaries:

$$\mathcal{F}(v, \alpha)/\mathcal{F}(0, \alpha) = \begin{cases} 0, & v < -\sqrt{3\alpha}v_{th}, \\ 1, & -\sqrt{3\alpha}v_{th} \leq v \leq \sqrt{3\alpha}v_{th} \\ 0, & v > \sqrt{3\alpha}v_{th}. \end{cases} \quad (4.17)$$

We choose α to give an initial waterbag distribution which will “break” at the current plasma wave amplitude, f :

$$fc = v_\phi \left(1 - \frac{8}{3}(\alpha\beta)^{\frac{1}{4}} + 2(\alpha\beta)^{\frac{1}{2}} - \frac{1}{3}\alpha\beta \right)^{\frac{1}{2}}. \quad (4.18)$$

Again, we assume that an initially Maxwellian distribution will evolve in a similar manner, to give:

$$\mathcal{F}'(v_\phi) = \mathcal{F}'_0(\sqrt{3\alpha}v_{th}). \quad (4.19)$$

In order to calculate the number of trapped electrons, we assume the number of electrons within $\pm\sqrt{2e\varphi_0/m}$ remains constant as the distribution

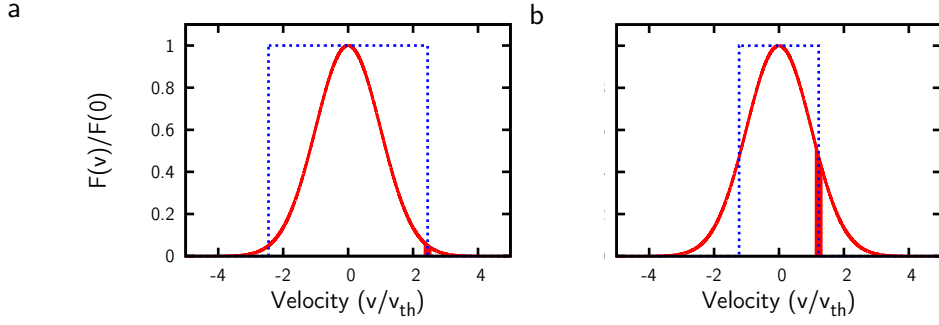


Figure 4.3: **Maxwellian and modified waterbag distributions.**

The model assumes that for an initially Maxwellian distribution, $\mathcal{F}_0(v)$, the distribution at the modified waterbag wavebreaking limit, $\mathcal{F}'(v)$, satisfies $\mathcal{F}'(v_\phi) = \mathcal{F}_0(\sqrt{3\alpha}v_{th})$. a) $\alpha = 2$ ($f < f_{WB}$), $\mathcal{F}'(v_\phi)$ is smaller, b) $\alpha = 0.5$ ($f > f_{WB}$), $\mathcal{F}'(v_\phi)$ is larger.

evolves. This gives:

$$\begin{aligned}
 n_{\text{trapped}}/n &= \frac{1}{2} \left[\text{erfc}(\sqrt{3\alpha/2}v_{th} - \sqrt{f}) - \text{erfc}(\sqrt{3\alpha/2}v_{th} + \sqrt{f}) \right], \\
 \mu &= 1 - \frac{1}{4} \left[\text{erfc}(\sqrt{3\alpha/2}v_{th} - \sqrt{f}) - \text{erfc}(\sqrt{3\alpha/2}v_{th} + \sqrt{f}) \right]. \quad (4.20)
 \end{aligned}$$

4.5 Limitations and comparison to other works

4.5.1 Limitations of the model

The model makes several approximations that limit its applicability.

The evolution equations are derived from a fluid model of plasma. As such, effects arising from the discretised nature of the plasma may only be included phenomenologically. Collisional and Landau damping are both included, while instabilities such as spontaneous backscatter are not. Such phenomenological treatments require that assumptions be made as to what form the plasma electron distribution takes. In this work, the distribution is assumed to remain Maxwellian as the plasma is heated. If the electron velocity distribution is not Maxwellian, it will impact on the resonant frequency of the plasma, the electron pressure, and the damping rates. Foremost amongst these will be the impact on the Landau damping rate, which is a direct interaction of thermal electrons with the plasma wave. The distribution of trapped electrons may also have a significant impact on the amplification efficiency. The model assumes a uniform velocity distribution for trapped electrons (although the same result is achieved for any linear distribution), and does not take into account the shift in the plasma wave resonance or the additional contribution to Landau damping due to trapped electrons.

In addition, the evolution equations are reduced to first order. This has

several consequences: dispersive effects and Raman forward scattering cannot be modelled; the model will be less accurate away from resonance; and the model is only valid in the Raman regime. However, none of these pose a significant problem. Dispersive effects have little influence in the planar case, and are only of real interest in models of higher dimensionality, in which focussing/guiding of the laser pulses may have significant impact. The growth rate for Raman forward scattering is typically much smaller than that for Raman backscatter [29], and experimental results show that Raman forward scattering is not significant for current experimental parameters [47, 51]. While the accuracy decreases away from resonance, we are interested in the efficient amplification of the laser pulses, and as a result our investigation will be limited to parameters near resonance. The Compton regime, while of interest, is beyond the scope of this investigation due to the typically high intensities required.

4.5.2 Other works

Some of these limitations have been resolved in other works. Spontaneous backscatter has been investigated using a three-wave model which includes a noise source to seed the instability [78]. Raman forward scatter has been modelled using an extended set of linearised equations [79]. The influence of anisotropic Maxwellian distributions has been investigated using a model similar to that described here [80].

Landau damping has been the subject of many works, including some which focus on the impact on Raman amplification. Landau damping will lead to a flattening of the electron distribution in the region of the phase velocity, causing the Landau damping rate to decrease [81]. Over long periods, the “damping” becomes oscillatory, as trapped electrons alternate between energy loss and gain [82], until this ordered behaviour is destroyed by collisions.

Yampolsky *et al* [83] extend the slowly-varying envelope model to include evolution equations for the saturation of Landau damping, and the nonlinear frequency shift due to trapped electrons. They find that saturation of Landau damping allows amplification in regimes where models using linear Landau damping predict no growth.

Lindberg *et al* [77] note that phase mixing of trapped electrons can lead to

transient damping above the linearised Landau damping rate. They develop an extended three-wave model that includes the evolution of a nonlinear electron distribution function, used to calculate the Landau damping rate and the frequency shift due to trapped electrons. Benchmarking against PIC simulations shows good agreement [70]. However, it should be noted that their model does not modify the Raman growth rates to account for spectral broadening of the plasma wave, and so good agreement may be limited to parameter spaces in which only a small fraction of electrons are trapped. Further, neither their model, nor the PIC code against which it is benchmarked, includes the effect of collisions, which act to reduce the nonlinearity of Landau damping.

The influence of collisions on the Landau damping rates is considered by Bilato and Brambilla [84].

The model developed in this chapter offers the advantages that it includes the influence of collisional heating, wavebreaking and spectral broadening, while maintaining low computational overheads. Its main disadvantage is that it relies on linearised Landau damping rates, which will become inaccurate as the electron distribution evolves. However, other models, such as PIC codes, typically overestimate the nonlinearity of this process due to the absence of collisions.

4.6 Numerical implementation

A simulation code, originally developed by Bernhard Ersfeld to investigate thermal effects [3], was extended to include the effects of harmonics and wavebreaking. The conventional approximations applied to the evolution equations ($\omega_0 + \omega_1 \approx 2\omega_{0,1}$, $\omega_0 - \omega_1 \approx \omega_p$, $v_g = c$, e.g. [42]) were also removed.

The slowly-varying envelope equations are solved using a set of FDTD algorithms, with the laser pulses introduced into the simulation domain at the boundaries. The code uses a leap-frog method, as discussed in Appendix A.1: the laser envelopes are evaluated at timestep n and spatial grid point m , while the electrostatic field is evaluated at $n + \frac{1}{2}$, $m + \frac{1}{2}$. This allows the midpoint value of the electrostatic field to be known explicitly for the laser solver, and vice versa, as shown in Figure (4.4). A semi-implicit update is used, as the two laser fields must be solved simultaneously; similarly, the co-dependence of the real and imaginary parts of f require that they be updated concurrently. Fortunately, the evolution equations can be evaluated by cross-substitution, which avoids the need for matrix inversion. Together, these methods act to avoid numerical instability in the simulation, especially important as we expect oscillatory solutions.

We here treat the evolution equations for the laser solver. The electro-




Figure 4.4: **Schematic of three-wave solver.**

Leapfrog update for laser field solver (LS) and electrostatic field solver (ES).

static solver is evaluated in an analogous manner. Taking the first-order Taylor expansion (see Appendix A.1) of the evolution equations derived in Section 4.3, we choose $\Delta t = \Delta z/v_g$ to give a simple advective solver for the laser fields:

$$\begin{aligned}
 [a_0]_m^{n+1} &= [a_0]_{m+1}^n - \Delta t \nu_0 [a_0]_{m+\frac{1}{2}}^{n+\frac{1}{2}} - \Delta t \gamma_0 [a_1]_{m+\frac{1}{2}}^{n+\frac{1}{2}} [f^*]_{m+\frac{1}{2}}^{n+\frac{1}{2}}, \\
 [a_1]_{m+1}^{n+1} &= [a_1]_m^n - \Delta t \nu_1 [a_1]_{m+\frac{1}{2}}^{n+\frac{1}{2}} + \Delta t \gamma_1 [a_0]_{m+\frac{1}{2}}^{n+\frac{1}{2}} [f]_{m+\frac{1}{2}}^{n+\frac{1}{2}},
 \end{aligned} \tag{4.21}$$

where $\gamma_{0,1} = \mu\omega_p \frac{\omega_0 + \omega_1}{4\omega_{0,1}}$, $\nu_{0,1} = \left(\frac{\omega_p}{\omega_{0,1}}\right)^2 \frac{\nu_{ei}[\omega_{0,1}]}{2}$. The leap-frog method means $[f]_{m+\frac{1}{2}}^{n+\frac{1}{2}}$ is known explicitly, so we drop the step notation. $[a_{0,1}]_{m+\frac{1}{2}}^{n+\frac{1}{2}}$ are obtained by interpolation:

$$\begin{aligned}
 [a_0]_{m+\frac{1}{2}}^{n+\frac{1}{2}} &= \frac{1}{2} ([a_0]_{m+1}^n + [a_0]_m^{n+1}), \\
 [a_1]_{m+\frac{1}{2}}^{n+\frac{1}{2}} &= \frac{1}{2} ([a_1]_m^n + [a_1]_{m+1}^{n+1}).
 \end{aligned}$$

Rearranging, we obtain:

$$\begin{aligned}
\left(1 + \frac{1}{2}\Delta t\nu_0\right) [a_0]_m^{n+1} &= \left(1 - \frac{1}{2}\Delta t\nu_0\right) [a_0]_{m+1}^n - \frac{1}{2}\Delta t\gamma_0 ([a_1]_m^n + [a_1]_{m+1}^{n+1}) f^* \\
\left(1 + \frac{1}{2}\Delta t\nu_1\right) [a_1]_{m+1}^{n+1} &= \left(1 - \frac{1}{2}\Delta t\nu_1\right) [a_1]_m^n + \frac{1}{2}\Delta t\gamma_1 ([a_0]_{m+1}^n + [a_0]_m^{n+1}) f
\end{aligned}
\tag{4.22}$$

These semi-implicit equations for the laser evolution are evaluated by cross-substitution:

$$\begin{aligned}
[a_0]_m^{n+1} &= \frac{\left(1 - \frac{1}{2}(\nu_0 - \nu_1) - \frac{1}{4}\Delta t^2\nu_0\nu_1 - \frac{1}{4}\Delta t^2\gamma_0\gamma_1|f|^2\right) [a_0]_{m+1}^n + \Delta t\gamma_0 [a_1]_m^n f^*}{\left(1 + \frac{1}{2}\Delta t(\nu_0 + \nu_1) + \frac{1}{4}\Delta t^2\nu_0\nu_1 + \frac{1}{4}\Delta t^2\gamma_0\gamma_1|f|^2\right)} \\
[a_1]_{m+1}^{n+1} &= \frac{\left(1 + \frac{1}{2}(\nu_0 - \nu_1) - \frac{1}{4}\Delta t^2\nu_0\nu_1 - \frac{1}{4}\Delta t^2\gamma_0\gamma_1|f|^2\right) [a_1]_m^n + \Delta t\gamma_1 [a_0]_m^n f}{\left(1 + \frac{1}{2}\Delta t(\nu_0 + \nu_1) + \frac{1}{4}\Delta t^2\nu_0\nu_1 + \frac{1}{4}\Delta t^2\gamma_0\gamma_1|f|^2\right)}
\end{aligned}
\tag{4.23}$$

4.7 Benchmarking

The three-wave model was benchmarked against the PIC code OSIRIS [2]. However, as discussed in Chapter 3, and illustrated in Fig. 4.5, numerical instabilities in PIC simulations may require that a high resolution be used in order to converge to a stable solution. The highest resolution used in the PIC simulations was 800 cells per pump wavelength, which, for some cases, was not sufficient to achieve a stable solution. However, these cases are included, as agreement may be gauged before the onset of instabilities. They also serve to emphasise the stability of the three-wave model.




Figure 4.5: **Illustration of high resolution required for PIC convergence.** Pump and probe envelopes from PIC and three-wave model, using different PIC resolutions. From left to right, 50, 100 and 200 cells per pump wavelength. Density of $1 \times 10^{19} \text{ cm}^{-3}$, pump intensity of $2.5 \times 10^{13} \text{ W cm}^{-2}$, other parameters as in Table (3.1).

4.7.1 Cold plasma

Benchmarking was first carried out for cold plasma, below the wavebreaking limit. Fig. 4.6 shows simulations run with varying pump intensity, while Fig. 4.7 shows the effect of varying plasma density. The three-wave model shows remarkable agreement with the PIC simulations. The disparity in Fig. 4.6d/4.7a is attributed to numerical instability in the PIC code.

Figs. 4.6,4.7 use optimal detuning, where the linearised three-wave equations are expected to give the best agreement. The interaction of a monochromatic pump with a detuned probe pulse is shown in Fig. 4.8, while Fig. 4.9 shows the effect of using a chirped pump. In both cases, agreement is good for the first peak, with some discrepancy observed in the tail of the amplified probe. The high level of agreement may be attributed to the reduced growth rate away from resonance, so the reduced accuracy of the linearised equations has little impact on the pulse evolution.

Fig. 4.10 shows the influence of including the harmonics of the plasma wave. While their inclusion gives improved agreement with the PIC simulations, the effect is small. This is in part due to benchmarking being limited to low amplitude plasma waves, in order to avoid numerical instability, detailed in Section 3.2.2. The high level of agreement for low amplitude plasma waves is mainly attributed to avoiding the common approximations used in

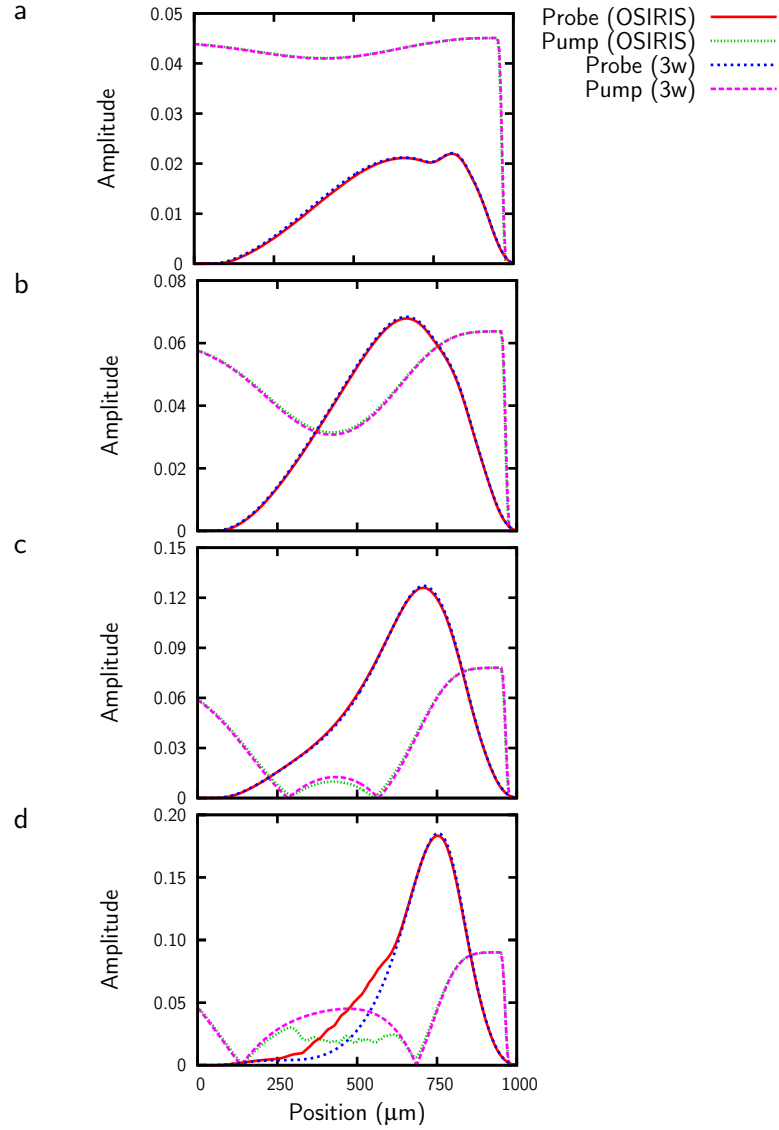


Figure 4.6: **Benchmarking for varying pump intensity.**

Pump and probe envelopes from PIC and three-wave model for a pump intensity of: a) 2.5×10^{13} , b) 5×10^{13} , c) 7.5×10^{13} and d) 1×10^{14} W cm⁻². Density of 1×10^{19} cm⁻³, other parameters as in Table (3.1).

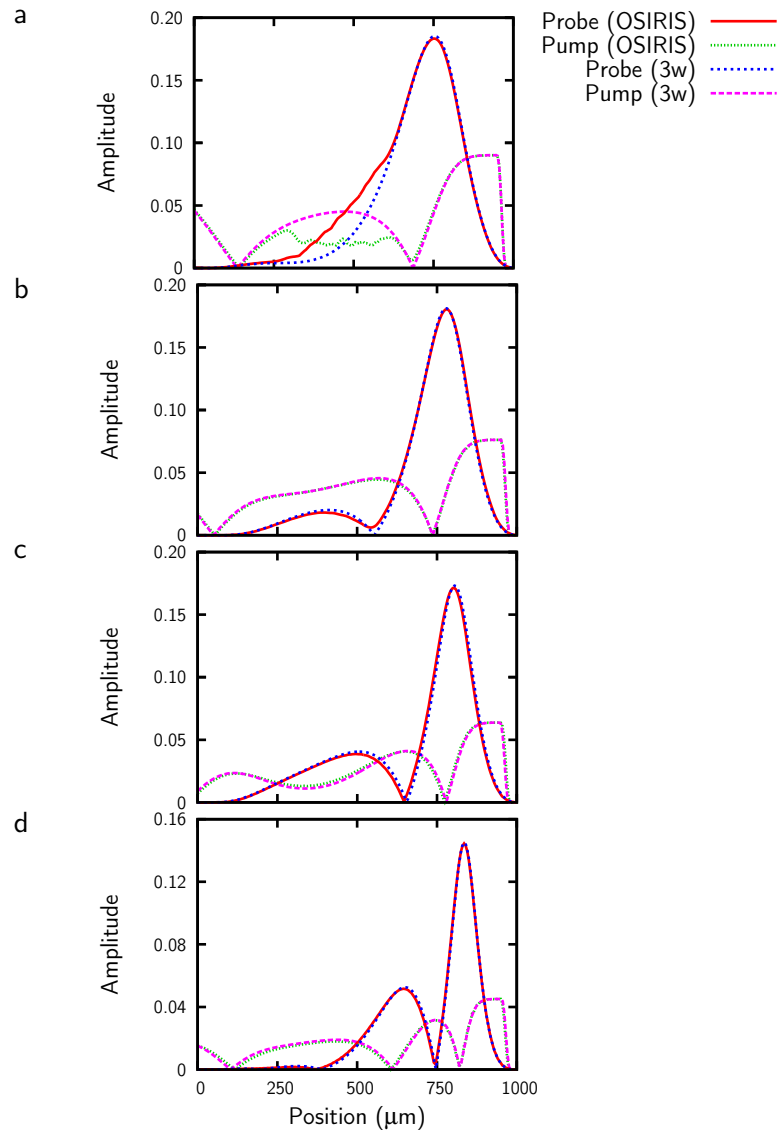


Figure 4.7: **Benchmarking for varying plasma density.**

Pump and probe envelopes from PIC and three-wave model for plasma density of: a) 1×10^{19} , b) 1.4×10^{19} , c) 2×10^{19} and d) $4 \times 10^{19} \text{ cm}^{-3}$. Other parameters as in Table (3.1).

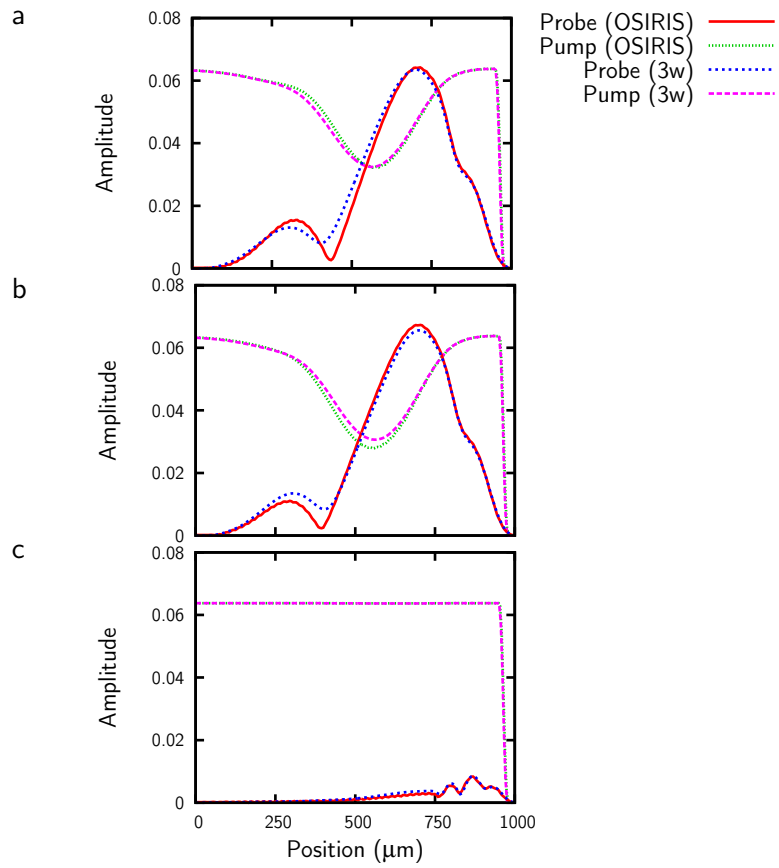


Figure 4.8: **Benchmarking for varying detuning.**

Pump and probe envelopes from PIC and three-wave model pump/probe detuning of: a) $1.05 \omega_p$, b) $0.95 \omega_p$, and c) $0.9 \omega_p$. Other parameters as in Table (3.1).

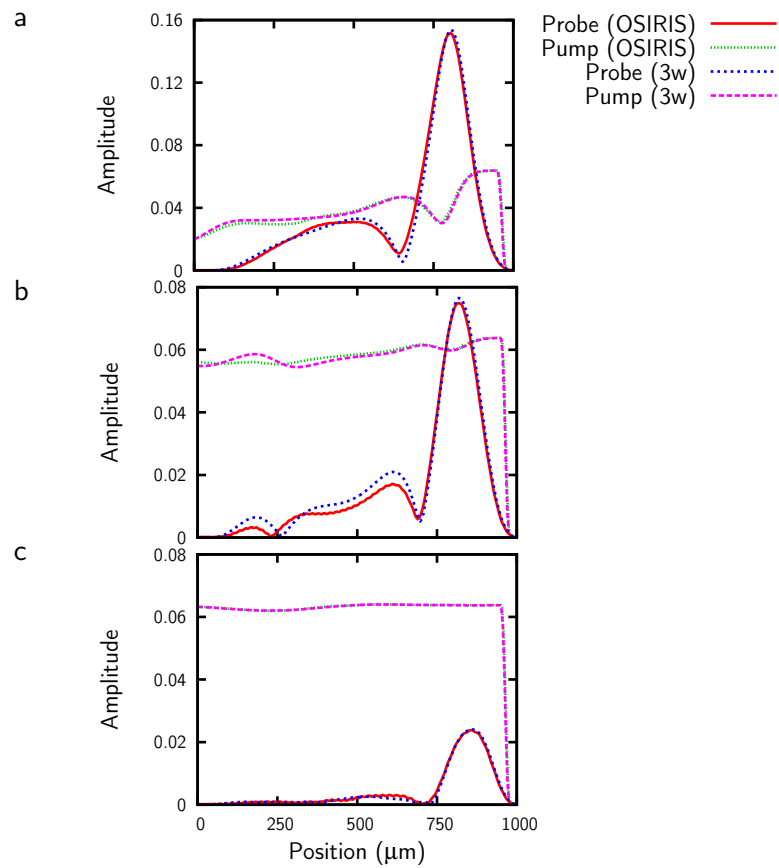


Figure 4.9: **Benchmarking for varying pump chirp.**

Pump and probe envelopes from PIC and three-wave model for chirped pump of bandwidth: a) 0.4%, b) 1.0% and c) 2.0%. Other parameters as in Table (3.1).




Figure 4.10: **Influence of harmonics on Raman amplification.**

Pump and probe envelopes from PIC and three-wave model. In the first figure, the influence of harmonics is included in the three wave model. In the second, it is neglected. Plasma density of $2 \times 10^{19} \text{ cm}^{-3}$, other parameters as in Table (3.1).

the three-wave evolution equations. These approximations reduce the accuracy of the three-wave model, (see, for example, the cold plasma case shown in Wang *et al* [70]). While the approximations simplify the equations, they do not reduce the computational overhead of the simulation.

4.7.2 Cold plasma: wavebreaking

Simulations have been carried out for cold plasma in the wavebreaking regime to benchmark the extensions of the three-wave model. Fig. 4.11 shows simulations carried out for different plasma densities, with the pump intensity increased to reach the wavebreaking limit. Fig. 4.12 uses a low plasma density, and increases the pump intensity to reach beyond the wavebreaking limit. The three wave model without the wavebreaking extensions is included for comparison.

The extended three-wave model shows good agreement with the PIC simulations, with close agreement over most of the first peak of the amplified probe. Some discrepancy may be seen for the peak amplitude in Figs. 4.12b,c, which is attributed to continued growth after the wavebreaking amplitude has been reached, due to strong ponderomotive forces driving collective behaviour after the plasma wave has broken. Such behaviour strongly depends on the electron velocity distribution after wavebreaking and cannot be readily reproduced by a fluid model. However, the effect only becomes noticeable for intensities approaching the Compton regime (in Fig. 4.12c, the bounce frequency $\approx 0.4\omega_p$), so in most cases may be neglected.

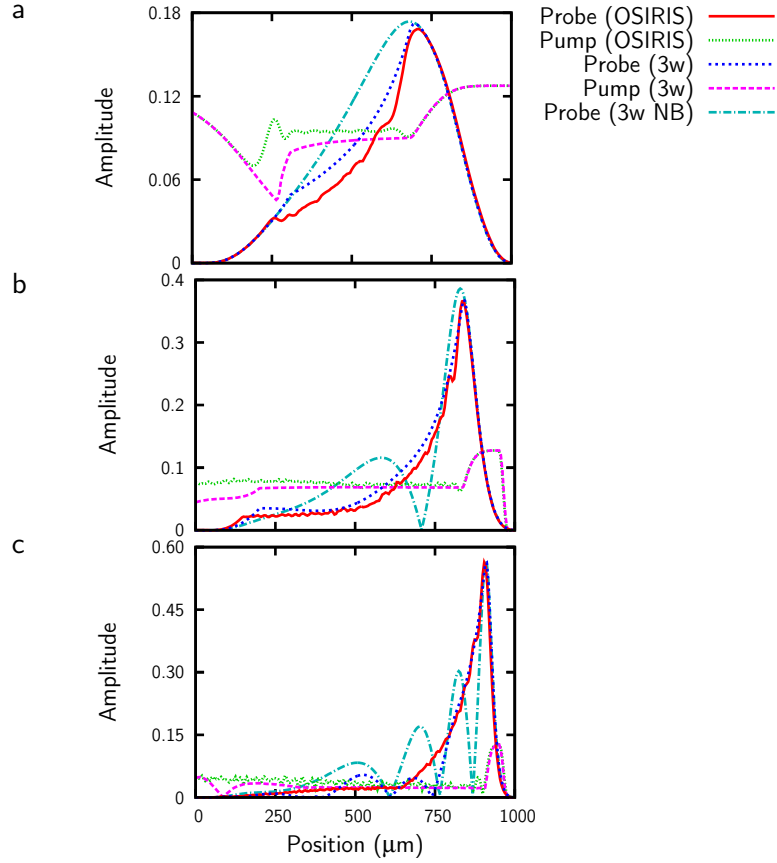


Figure 4.11: **Benchmarking of wavebreaking model near wavebreaking limit.**

Pump and probe envelopes from PIC and three-wave model in wavebreaking regimes. Both the extended three-wave model and the conventional model with no breaking (3w NB) are shown. Plasma density and pump intensity of: a) $5 \times 10^{18} \text{ cm}^{-3}$, $1 \times 10^{14} \text{ W cm}^{-2}$, b) $1 \times 10^{19} \text{ cm}^{-3}$, $2 \times 10^{14} \text{ W cm}^{-2}$, and c) $2 \times 10^{19} \text{ cm}^{-3}$, $4 \times 10^{14} \text{ W cm}^{-2}$. Other parameters as in Table (3.1).

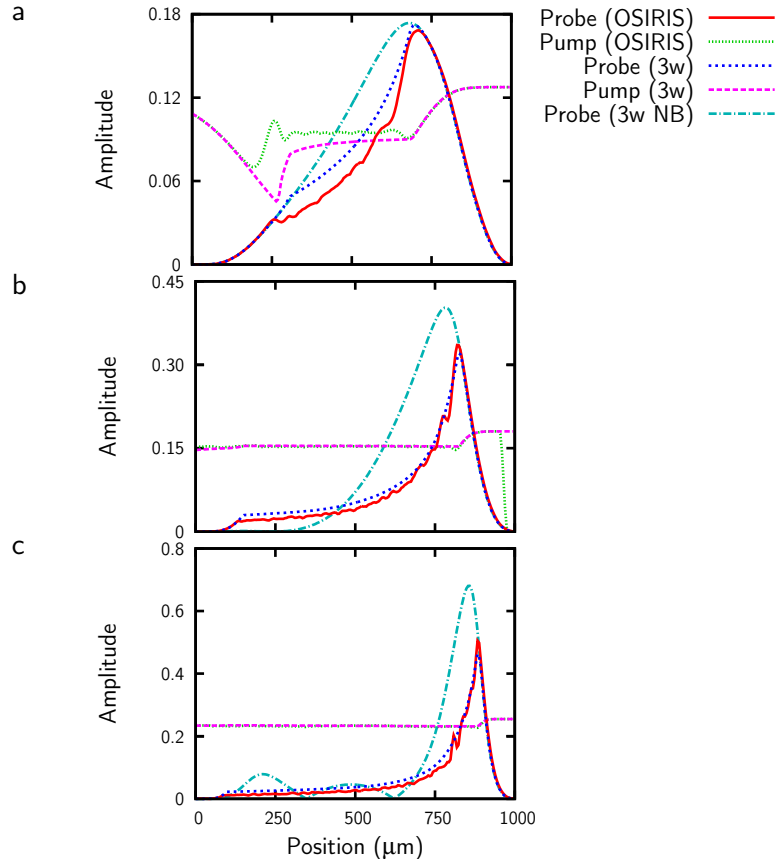


Figure 4.12: **Benchmarking of wavebreaking model beyond wavebreaking limit.**

Pump and probe envelopes from PIC and three-wave model in wavebreaking regimes. Both the extended three-wave model and the conventional model with no breaking (3w NB) are shown. Plasma density of $5 \times 10^{18} \text{ cm}^{-3}$, pump intensity: a) 1×10^{14} , b) 2×10^{14} and c) $4 \times 10^{14} \text{ W cm}^{-2}$. Other parameters as in Table (3.1).

4.7.3 Thermal plasma

As discussed in Chapter 3, PIC codes typically do not include collisional effects, and therefore exclude collisional damping, and thermalisation of the electron distribution. A modified version of the three-wave code, in which collisional damping and heating were excluded, was developed to allow direct comparison with PIC simulations. Eqs. (4.11,4.12) were modified as follows:

$$\begin{aligned} \left(\frac{\partial}{\partial t} - v_g \frac{\partial}{\partial z}\right) a_0 &= -\omega_p \frac{\omega_0 + \omega_1}{4\omega_0} \mu f^* a_1, \\ \left(\frac{\partial}{\partial t} + c \frac{\partial}{\partial z}\right) a_1 &= \omega_p \frac{\omega_0 + \omega_1}{4\omega_0} \mu f a_0, \\ \left(\frac{\partial}{\partial t} + \frac{\nu_L}{2} + i(\omega_0 - \omega_1 - \omega'_p)\right) f &= \frac{\omega_p(\omega_0 + \omega_1)}{2(\omega_0 - \omega_1 + \omega_p)} a_0^* a_1. \end{aligned} \quad (4.24)$$

$$\frac{\partial}{\partial t} \left(\frac{k_B T}{m}\right) = c^2 \nu_L |f|^2, \quad (4.25)$$

where T now reflects only the spread in electron velocity in the z direction.

The absence of collisional damping in current PIC codes also leads to an artificially high growth rate for instabilities, as discussed in Section 3.2.4. Benchmarking simulations of thermal plasmas are therefore limited to cases where Landau damping is sufficiently strong to suppress spontaneous Raman backscatter. This has been achieved by using a plasma density of $5 \times 10^{18} \text{ cm}^{-3}$ for all simulations of thermal plasma. This constraint is unfortunate, as it prevents independent analyses of Landau damping and particle trapping.

However, we may make use of the work of Wang *et al* [70], which used a modified PIC code to suppress spontaneous backscatter. Their simulations were carried out in regimes far below Coffey’s warm-wavebreaking limit (high density, high phase velocity), resulting in a much smaller fraction of trapped electrons. They observe the linearised Landau damping rates become inaccurate for temperatures above 200 eV. As we expect Landau damping to be characterised by the ratio $v_{th}/v_\phi \sim \sqrt{T}\omega_0/\omega_p$, we expect the linearised Landau damping rates to be accurate for $v_{th} < 0.2v_\phi$, corresponding to temperatures below 15 eV for the simulation parameters used here.

Fig. 4.13 shows a set of simulations carried out with different initial temperatures. For comparison, plots are shown using the wavebreaking model derived in Section 4.2.3 and the trapping model developed in Section 4.3.2. In the wavebreaking model, amplification saturates after the plasma wave reaches the wavebreaking amplitude; simulations were carried out using both the cold wavebreaking limit and Coffey’s warm wavebreaking limit.

As expected, the growth rate predicted by the wavebreaking model is too high when the cold wavebreaking limit is used, and too low when Coffey’s warm wavebreaking limit is used. The trapping model gives good agreement for the peak amplitude at lower temperatures (10 eV and below), but overestimates the pulse duration. This disparity is attributed to the shift in the resonant frequency of the plasma wave due to trapped electrons, which is not

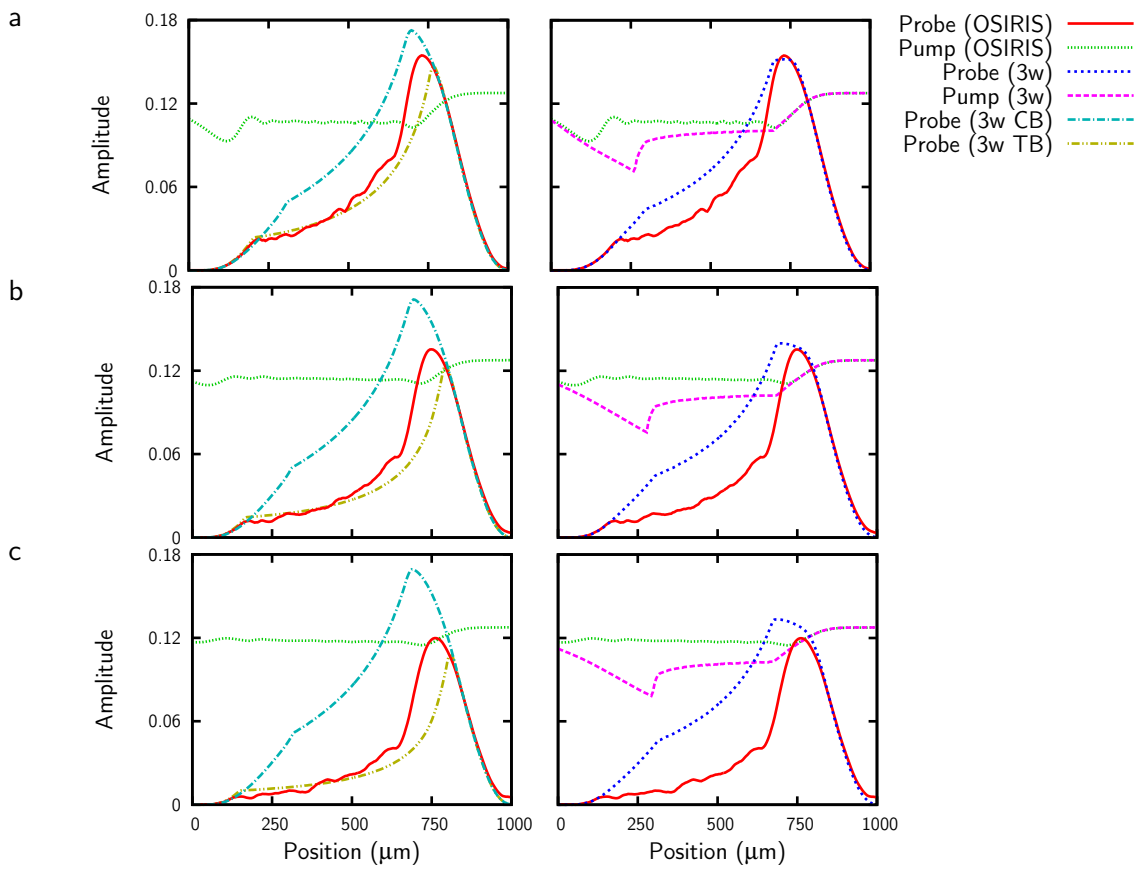


Figure 4.13: **Benchmarking of wavebreaking and trapping models in warm plasma.**
(continued overleaf)

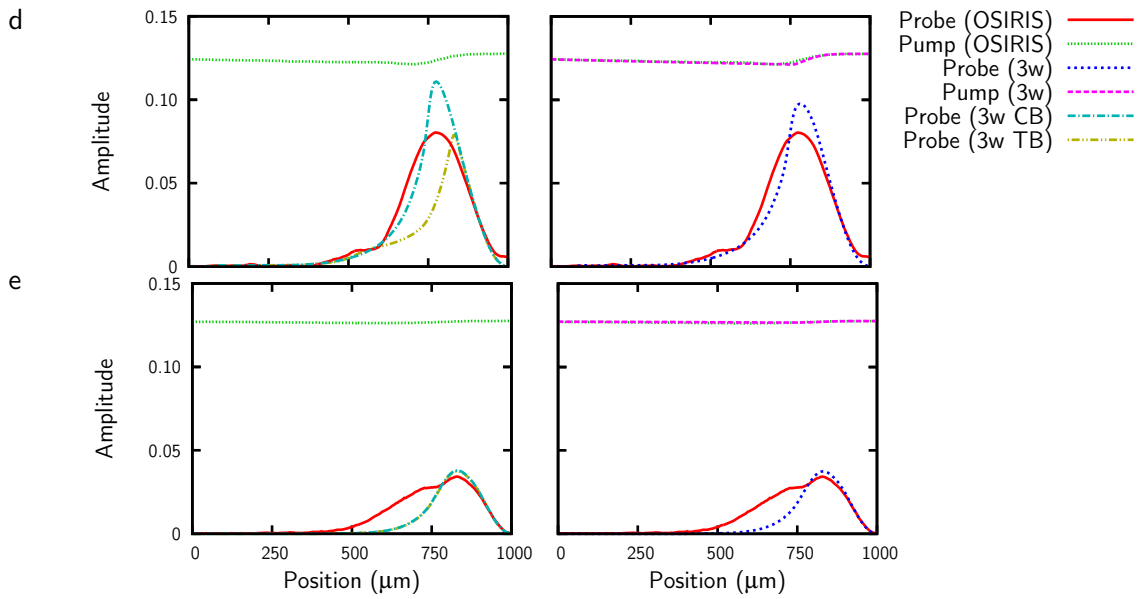


Figure 4.13: **Benchmarking of wavebreaking and trapping models in warm plasma.**

(continued from previous page) Pump and probe envelopes from PIC and three-wave model for varying temperature. The left hand plot shows the wavebreaking model using both the cold wavebreaking limit (3w CB) and Coffey's warm wavebreaking limit (3w TB). The right hand plot shows the trapping model. Initial temperature of: a) 1 eV, b) 5 eV, c) 10 eV, d) 25 eV and e) 50 eV. Plasma density of $5 \times 10^{18} \text{ cm}^{-3}$, other parameters as in Table (3.1).

accounted for in the model. At higher temperatures, the three-wave model overestimates the peak amplitude, but underestimates the pulse duration. These two effects are attributed to nonlinear Landau damping: the transient increase in the damping rate due to trapped electrons reduces the peak pulse intensity, while saturation of Landau damping allows increased growth towards the rear of the pulse. This behaviour agrees with the predictions for the applicability of the linearised Landau damping rates based on the simulations by Wang *et al.*

The same simulations were carried out using an initial probe intensity increased to that of the pump intensity, $1 \times 10^{14} \text{ W cm}^{-2}$, shown in Fig. 4.14. The simulations show better agreement for the amplified pulse duration. This is consistent with a shift in the resonant frequency of the plasma wave, as higher intensity pulses remain short as they are amplified, and so are less affected by detuning. The PIC simulations show absorption of the initial probe behind the first peak, again consistent with a shift in the plasma wave resonance.

Fig. 4.15 shows a comparison of the trapping model derived here, and that derived by Bohm and Gross for particle trapping [74]. It is that seen the Bohm Gross model predicts significantly lower growth than observed in the PIC simulation. This is attributed to the assumed distribution function for trapped electrons used in the Bohm-Gross model leading an increased shift

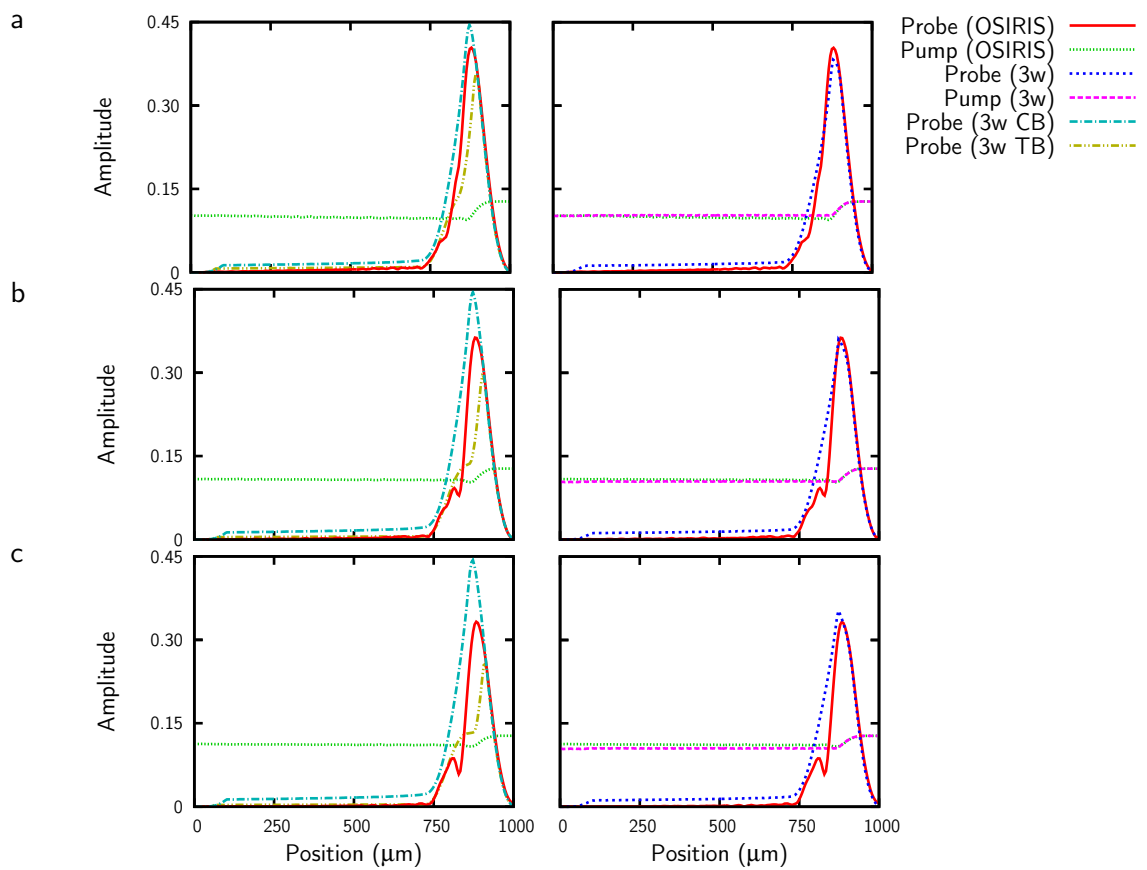


Figure 4.14: **Benchmarking of wavebreaking and trapping models in warm plasma (high intensity).**
 (continued overleaf)

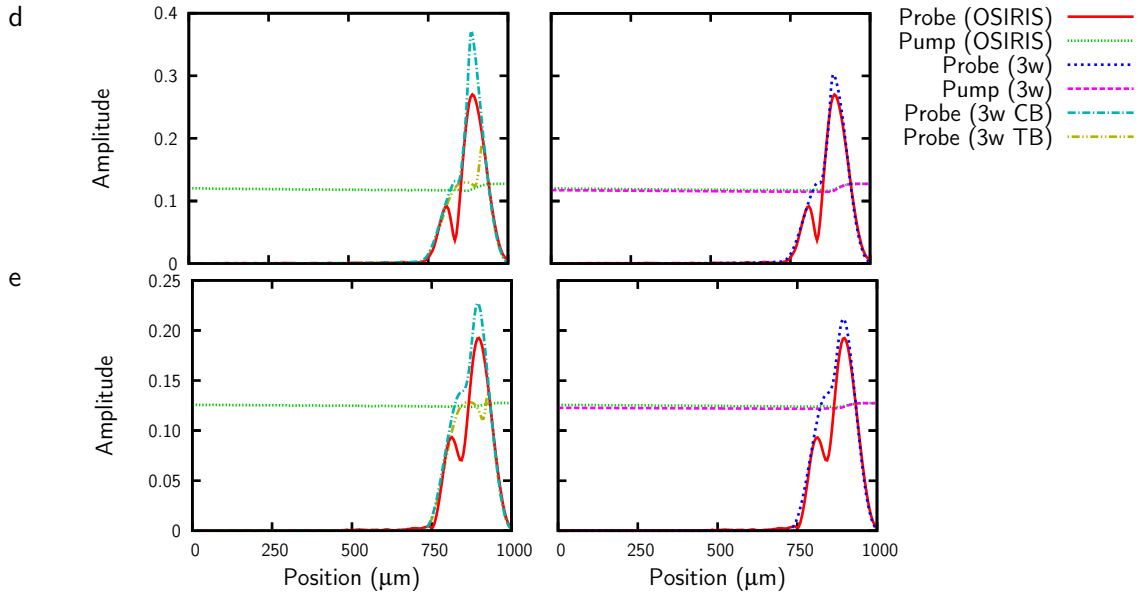


Figure 4.14: **Benchmarking of wavebreaking and trapping models in warm plasma (high intensity).**

(continued from previous page) Pump and probe envelopes from PIC and three-wave model for varying temperature. The left hand plot shows the wavebreaking model using both the cold wavebreaking limit (3w CB) and Coffey's warm wavebreaking limit (3w TB). The right hand plot shows the trapping model. Initial temperature of: a) 1 eV, b) 5 eV, c) 10 eV, d) 25 eV and e) 50 eV. Plasma density of $5 \times 10^{18} \text{ cm}^{-3}$, probe intensity of $1 \times 10^{14} \text{ W cm}^{-2}$, other parameters as in Table (3.1).

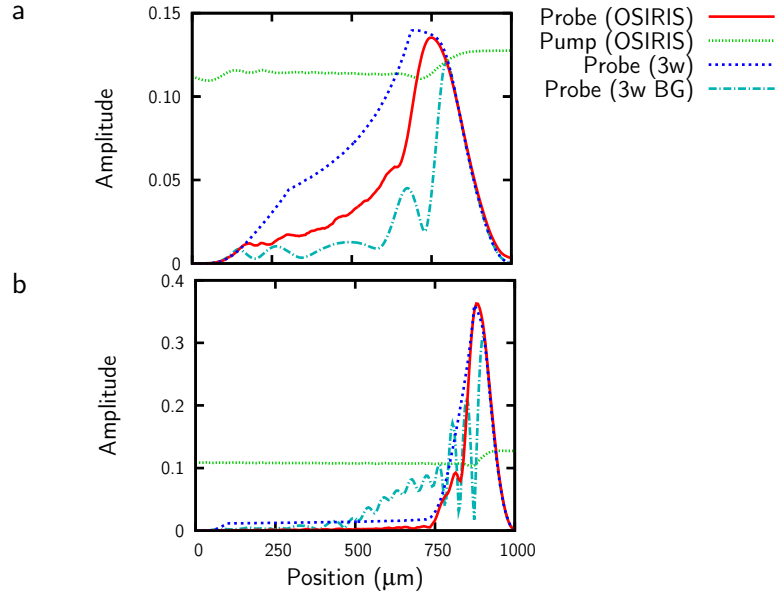


Figure 4.15: **Comparison with Bohm-Gross model**

Benchmarking Pump and probe envelopes from PIC and three-wave model, using both the trapping model derived in this work and that derived by Bohm and Gross (3w BG). Probe intensity a) $1 \times 10^{12} \text{ W cm}^{-2}$, b) $1 \times 10^{14} \text{ W cm}^{-2}$. Plasma density $5 \times 10^{18} \text{ cm}^{-3}$, initial temperature 5 eV, other parameters as in Table (3.1).

in plasma wave resonance. However, it must again be emphasised that the absence of collisions in the PIC code may lead to a nonphysical distribution of trapped electrons.

4.8 Case studies

While the goal of this work is model development, we here present some simulation results to illustrate potential applications of the code.

4.8.1 Collisional damping

As discussed in Chapter 3, PIC codes typically do not include collisional effects. Fig. 4.16 shows a comparison of the collisionless model developed for benchmarking against PIC codes (described in Section 4.7) and the full thermal three-wave model. As can be seen, the inclusion of collisional effects leads to a significant reduction in the intensity of the amplified probe. In addition, collisional effects modify the optimal parameters for amplification: in the collisionless model, lower temperatures give increased growth, as Landau damping and particle trapping are reduced, while for the full model, higher amplification is observed for higher temperatures, due to the decrease in collisional damping. The optimal temperature range depends on the parameters used, and a model including collisions is useful for the planning of experimental work.

The observed reduction in amplification associated with the inclusion of collisional effects is mainly due to the additional shift in the plasma resonance frequency as the plasma is heated, rather than the direct action of collisional

damping. The optimal detuning for amplification varies significantly for the collisional and collisionless models, as shown in Fig. 4.17, mainly due to pre-heating of the plasma by the pump ahead of the Raman interaction. As the plasma temperature varies in both space and time, the probe pulse observes a temperature gradient as it interacts with the pump. This results in a “thermal chirp”, as the resonant plasma frequency varies with both time and position. This allows a pulse to retain its short duration as it is amplified [52], analogous to using a chirped pump [43]. However, as with a chirped pump, this is accompanied by a decrease in peak intensity [68]. The use of a chirped pump may be used to counteract the thermal chirp, which may lead to a significant increase in amplification, demonstrated in Fig. 4.17.

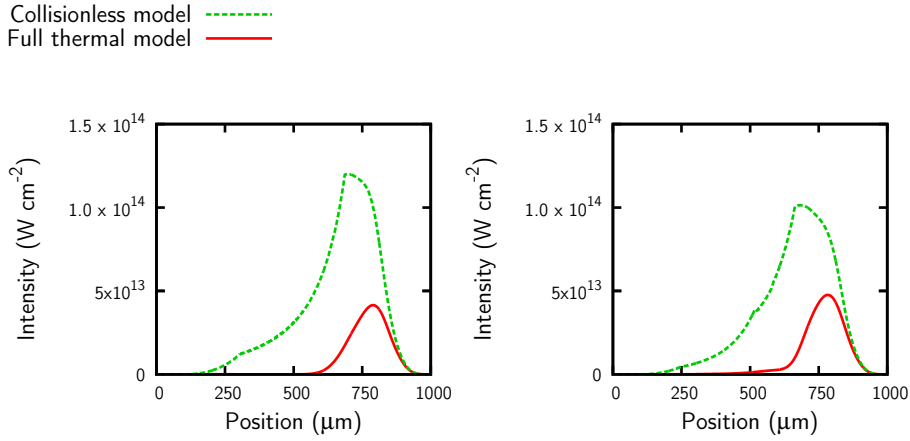


Figure 4.16: **Influence of collisions.** Probe envelopes, showing full thermal model and collisionless (PIC-like) model, for 5 eV and 15eV. Other parameters as in Table (3.1).




Figure 4.17: **Influence of thermal effects with and without collisions.** Parameter scans showing peak amplified probe intensity, comparing cold plasma model, full thermal model and collisionless (PIC-like) model, for varying pump/probe detuning, and varying pump chirp. Other parameters as in Table (3.1).




Figure 4.18: **Influence of wavebreaking.**

Snapshots of probe intensity during amplification in the wavebreaking regime. Pump intensity is shown for the final snapshot. Simulation carried out for cold plasma of density $5 \times 10^{18} \text{ cm}^{-3}$, initial pump of intensity $1 \times 10^{15} \text{ W cm}^{-2}$, initial probe of intensity $1 \times 10^{15} \text{ W cm}^{-2}$ and FWHM duration 100 fs. Other parameters as in Table (3.1).

4.8.2 Wavebreaking

The inclusion of wavebreaking allows some interesting regimes to be studied.

Assuming pump depletion is low, the wave breaks when $\int a dt$ reaches some critical value. In the general case, illustrated in Fig. 4.18, this gives near-superradiant scaling [39], as the growth rate behind the head of the probe is only approximately linear. Malkin *et al* [42] discuss a regime in which the plasma wave breaks after the first peak, suppressing Burnham-Chiao ringing [85]. In this case the growth should be superradiant, due to the scaling of the pump-depletion regime.

4.8.3 Noise

The amplification of noise is a concern for Raman amplification experiments, as the exponential growth seen in the linear regime allows small signals to be amplified to comparable intensities to the probe pulse, causing premature depletion of the pump and the introduction of seed precursors. While the code does not model spontaneous backscatter, it may be used to give some indication of how susceptible a system is to noise, by comparing growth rates for probe pulses of different intensity. An ideal system would give low growth for noise, but large growth for seeded amplification [68]. Heating and damping have been suggested as processes to suppress spontaneous backscatter [44].

Fig. 4.19 shows effect of varying probe intensity for different initial conditions. As expected, for a monochromatic pump, the inclusion of thermal effects reduces the growth rate. For these parameters, a higher initial temperature increases the growth rate, as the “thermal chirp” is reduced. However, the reduction in amplification for small signals is typically of the order of 100, which may not be sufficient to suppress spontaneous backscatter for long propagation lengths.

The use of a chirped pump may be used to further reduce the growth rate. The use of an up-chirped pump will counteract the effect of the thermal chirp,

leading to an increased growth rate for small signals. If a down-chirp is used, the pump and thermal chirps will add together, leading to a further reduction in the amplification of small signals. There is a compromise between the growth rate for seeded (large signal) amplification, and the amplification of small signal noise. The optimal values will vary for the experimental parameters.

4.8.4 Impact of wavebreaking and trapping

The simplicity of the three-wave model allows the impact of individual processes to be evaluated, as they may be removed from the model and a comparison made. Fig. 4.20 shows a comparison of amplification in cold plasma, for the full model, and ignoring harmonics and wavebreaking, and warm plasma, including and excluding the impact of particle trapping. As expected, the impact of both wavebreaking and particle trapping is reduced at high plasma densities, where the fractional modulation depth of the plasma wave is lower. Similarly, increasing the intensity of the initial seed leads to an increase in these effects, due to the increase in the amplitude of the excited plasma wave. It is worth noting that the influence of spectral broadening is negligible in the parameter space investigated by Wang *et al* [70].

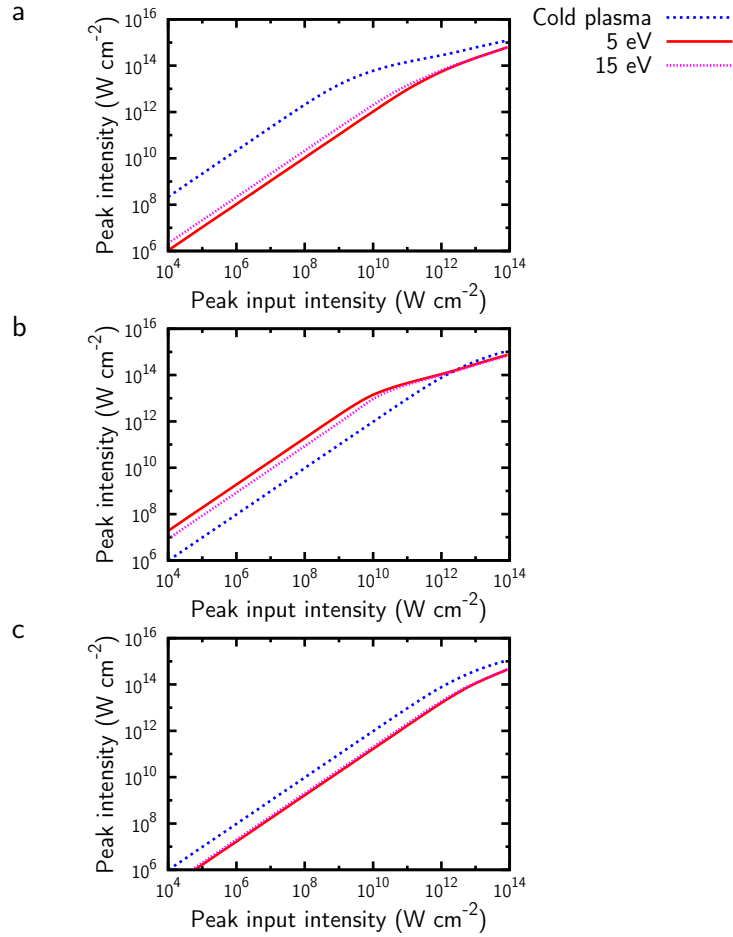


Figure 4.19: **Influence of thermal effects on Raman growth rates.** Peak amplified probe intensity, comparing cold plasma model, and thermal model at 5 eV and 15 eV, using a a) monochromatic, b) up chirped, and c) down chirped pump. Pump bandwidth in the chirped pump case is $0.05\omega_0$. Other parameters as in Table (3.1).

Cold plasma without breaking - - -
 Cold plasma with breaking - · - · -
 Thermal plasma without trapping —
 Thermal plasma with trapping - · - · -




Figure 4.20: **Influence of wavebreaking and particle trapping on Raman growth rates.**

Peak amplified probe intensity, comparing cold plasma model, including and excluding wavebreaking, and thermal model at 5 eV, including and excluding particle trapping. Probe intensity of a) 1×10^{12} W cm⁻², b) 1×10^{14} W cm⁻². Other parameters as in Table (3.1).

4.9 Conclusion

The three-wave model for Raman amplification is extended, in the cold plasma case, to include the effects of wavebreaking and harmonics, and, in the warm-plasma case, to include the influence of broadening of the wavenumber spectrum due to particle trapping. Benchmarking was carried out against particle-in-cell simulations.

Agreement for cold plasma below the wavebreaking limit is excellent, attributed to avoiding approximations commonly used in the three-wave model. The inclusion of harmonics improves agreement with PIC simulations, although the effect is small. In wavebreaking regimes, strong agreement is achieved for the first peak, with some disparity seen for the peak amplitude as parameters approach the onset of the Compton regime, which attributed to the direct action of the ponderomotive force on plasma electrons after the plasma wave has broken.

For warm plasmas, the applicability of the model is limited by the use of linearised Landau damping rates, which are shown to become less accurate for thermal velocities $v_{th} > 0.2v_\phi$. Below this threshold, the inclusion of spectral broadening is shown to greatly improve the agreement of the three-wave model with PIC simulations. Good agreement is seen for the peak intensity of the amplified pulse, but the model tends to overestimate the pulse duration

at low growth rates, attributed to the model not including the nonlinear shift in the plasma wave resonance frequency due to trapped electrons. At higher intensities, the interaction is less sensitive to this detuning as the amplified laser pulse remains short, with the result that detuning effects do not have a significant impact, and strong agreement is seen.

This illustrates that the three-wave model remains a powerful tool for the simulation of Raman amplification, as additional effects not explicitly modelled by the three-wave formalism may readily be included in a phenomenological manner. Particle trapping is shown to have an additional impact on growth rates beyond the modification of Landau damping rates and resonant frequency of the excited plasma wave, due to broadening of the wavenumber spectrum.

The use of an analytical treatment to estimate the number of trapped electrons avoids the need to model the evolution of the entire electron distribution function. While such a treatment limits the applicability of the model, as nonlinear Landau damping and detuning due to particle trapping are not modelled, it incurs a significantly lower computational overhead than modelling the evolution of the distribution function.

The low computational overhead allows parameter scans or large simulations to be carried out, which would not be feasible using PIC codes. Further, collisional effects, which are typically not included in PIC models due to the

complexity and computational overhead of their implementation, may be investigated. It is shown that collisional effects can have a significant impact not only on the efficiency, but on the optimal parameters for amplification, making them an important consideration for experimental works.

The model is not currently suitable for the simulation of experiments carried out at Strathclyde. Although the initial conditions of the experiment (800 nm laser, plasma density $2 \times 10^{18} \text{ W cm}^{-2}$, temperature of 5 eV) satisfy $v_{th} < 0.2v_\phi$, collisional heating and Landau damping will quickly lead to regimes in which the linearised Landau damping rates are not valid.

4.10 Future work

The applicability of the model could be extended through the inclusion of nonlinear Landau damping rates and the shift in plasma resonance due to trapped electrons. Conversely, it may be more practical to extend the model developed by e.g. Lindberg *et al* [77] to account for spectral broadening in the three-wave evolution equations. Further analytical work could be carried out on the impact of the velocity distribution of trapped electrons, and the dependence of this distribution on collisional effects.

Implementing a second-order solver for the plasma wave should allow the Compton regime to be modelled, and would allow simulations in which the probe bandwidth exceeds the pump/probe detuning. The use of a second-order laser update (discussed in Appendix A.3) would allow dispersive effects to be modelled. Implementing a three-dimensional model would further exploit the low computational overhead of the model, and increase the applicability of simulations to experimental work. The assumption that plasma heating may be treated adiabatically is questionable for long interaction times (for example in centimetre-scale capillaries), but an extended treatment does not make sense in a 1D (planar) scheme, as heat transport will predominantly be in the transverse direction.

Chapter 5

Development of the aPIC model

Chapter Abstract

The aPIC (averaged-PIC) code is an envelope-PIC model created by MinSup Hur for the simulation of Raman amplification in plasma [4]. It offers many of the advantages of a conventional PIC code, but has a substantially smaller computational overhead.

In this chapter, the aPIC model is described, and comparisons are made with conventional envelope-PIC models. Modifications made to the code to allow simulation of the experimental work at Strathclyde are described. These include the implementation of chirped laser pulses, the inclusion of single-pulse ponderomotive contributions (which drive Raman forward scattering and wake-field generation), and a dispersive model for the laser solver.

Benchmarking shows strong agreement with conventional PIC codes, with some disparity for off-resonance amplification, attributed to the fully-explicit particle push.

A case study using the aPIC code illustrates a new regime of broad bandwidth amplification, achieved through the exploitation of nonlinear plasma waves.

5.1 Introduction

The experimental work carried out at Strathclyde involves centimetre-scale plasmas in regimes which may be strongly influenced by kinetic effects. Long propagation distances prevent simulation with conventional PIC models, due to the prohibitive computational overhead, as discussed in Chapter 3. Three-wave models offer the advantage of low computational overhead, but, as the plasma response is derived from a fluid treatment of plasma, they cannot explicitly model particle effects. While, as shown in Chapter 4, such effects may be included phenomenologically, this requires that a suitable analytical model exist for the process.

The aPIC code [4] offers the advantages of a PIC code, but has a far lower computational overhead, achieved through the use of several approximations. It relies on an envelope representation of the laser fields, coupled with an electrostatic solver. However, in its original form, the code was not suitable for modelling Strathclyde experiments, which used chirped pulses, in regimes where plasma wakes may act to influence the interaction.

Conventional envelope-PIC models are described to allow comparisons to be drawn with the aPIC model. The aPIC code is discussed, including the derivation of the governing equations, the limitations introduced through the use of approximations, and the numerical implementation of the model.

The extensions made to the code are detailed, and benchmarking carried out against the fully resolved PIC code OSIRIS. The applicability of the code is illustrated through a case study, in which a new regime of broad bandwidth amplification is identified. Conclusions are drawn and future work identified.

5.2 Envelope PIC

5.2.1 Separation of scales

As discussed in Chapter 3, one of the biggest limitations of conventional PIC codes is the computational overhead. However, in many cases this overhead is predominantly due to the modelling of the laser fields. In the case of wakefield acceleration, for example, the laser wavelength is typically an order of magnitude smaller than that of the plasma wave [86]. In order to resolve the laser fields, a large number of cells are required, while the plasma structure alone can be resolved using far fewer cells.

The three-wave model derived in Appendix A.2 uses an envelope representation of the laser pulses and plasma wave. However, one of the consequences of the envelope representation of the plasma wave is that it cannot explicitly model particle effects. As shown in Chapter 4, these effects may be included phenomenologically, but this is only possible if a suitable analytical model exists for the process.

In the case where the laser field is responsible for the smallest simulation scale, an envelope model can be used to represent the laser field, while an explicit treatment for particles is maintained through a PIC model, allowing simulation of processes such as wakefield acceleration or Landau damping.

The use of an envelope model allows the laser field to be modelled at

any point without approximation, with variations in frequency accounted for through the use of complex envelopes. However, this does not lead directly to reduced computational overheads. The electron response will still vary on the scale of the laser frequency and wavelength, while an envelope representation chosen to match forward-propagating pulses will result in the envelope of reflected or backscattered components varying on a scale of half the laser wavelength¹. In order to take advantage of the envelope model, it is necessary to make approximations for both the electron response and laser evolution.

¹Depending on the choice of carrier wave, the envelope of reflected components will either vary at twice the wavenumber or twice the frequency of the forward-propagating field.

5.2.2 Approximations required for envelope-PIC

The assumption for the electron response typically takes the form of the quasi-static approximation, as used in the three-wave model derived in Appendix A.2. This assumes that the electron motion may be represented as the sum of a slowly-varying component and a fast oscillatory response to the laser field, with the latter dependent only on the instantaneous laser field, expressed as a vector potential.

Only the slowly-varying component of the electron response is modelled explicitly, with the particle represented by its average position over a laser cycle. The laser field modifies the slowly-varying electron distribution through the ponderomotive force, which acts to push charged particles away from areas of high intensity, and through the variation in Lorentz factor due to the fast electron oscillation in the laser field. The plasma acts on the laser through the modification of the plasma susceptibility.

Neglecting numerical solutions corresponding to reflection or backscatter of the laser field is readily achieved in the model by reducing the wave equation to first order. However, simply neglecting second-order derivatives will also remove any dispersive behaviour. In cases where dispersion is important, the quasi-static approximation (also referred to as the “frozen field” approximation) for the laser evolution may be used, which approximates the disper-

sive behaviour of the full second-order equation while preventing backscatter or reflection. Several laser solvers are discussed in detail in Appendix A.3.

Together, these approximations have the result that only the slowly-varying plasma response and the variation of the laser envelope need to be modelled. This allows a significant reduction in the simulation resolution, leading to significant computational savings.

5.2.3 Envelope-PIC equations

The evolution of the laser field is modelled using the wave equation, and the electric and magnetic fields due to particles are calculated using the standard Yee solver [87]. The plasma susceptibility is calculated from the particle positions, and used to update the laser fields. The ponderomotive force due to the laser field is added to the Lorentz force for the particle push.

For simplicity, we assume that, over the short timescales of interest, the ions are stationary, and so deal only with the electron population. Starting from Ampere's law, we follow the derivations in Appendix A.2 to arrive at the wave equation (Eq. (A.17)):

$$\left(\frac{\partial^2}{\partial t^2} - c^2 \nabla^2\right) \vec{a} = -\frac{ne^2}{\varepsilon_0 \gamma m} \vec{a}.$$

The current term on the right-hand side may be found numerically by summing over all particles in a cell:

$$\begin{aligned} \left(\frac{\partial^2}{\partial t^2} - c^2 \nabla^2\right) \vec{a} &= -\frac{1}{\Delta x \Delta y \Delta z} \frac{1}{\varepsilon_0} \sum_j \frac{q_j^2}{\gamma_j m_j} \vec{a} \\ &= -\frac{1}{\Delta x \Delta y \Delta z} \frac{e^2}{m \varepsilon_0} \eta \sum_j \frac{1}{\gamma_j} \vec{a}, \end{aligned} \quad (5.1)$$

where $\Delta x \Delta y \Delta z$ is the cell volume, γ_j , m_j and q_j are the Lorentz factor, charge and mass of the j th particle, and η is the ratio between the plasma

density and the equilibrium macroparticle density in the PIC code, satisfying $\eta = m_j/m = -q_j/e$.

The Lorentz factor has contributions from both the slowly-varying momentum, \vec{p} , and the fast response to the laser field ²:

$$\gamma = \sqrt{1 + \frac{\vec{p}^2}{m^2 c^2} + |\vec{a}|^2}. \quad (5.2)$$

The electric and magnetic fields due to the slowly-varying motion of particles are calculated in the usual way using the Yee equations. The particle push for the slowly-varying momentum is then calculated using the sum of the Lorentz force and the ponderomotive force:

$$\frac{d\vec{p}}{dt} = -e \left(\vec{E} + \frac{\vec{p}}{\gamma m} \times B \right) - \frac{mc^2}{4\gamma} \nabla |\vec{a}|^2. \quad (5.3)$$

²This relation is only exactly satisfied when \vec{p} and \vec{a} are orthogonal. However, this is a reasonable approximation in regimes where the quasi-static approximation is valid

5.2.4 Applicability

Envelope-PIC models have been used to simulate wakefield acceleration, with good agreement to conventional PIC codes [63].

However, as discussed in Appendix A.2, the quasi-static approximation for electron response breaks down if the motion of electrons parallel to \vec{a} over a laser cycle becomes comparable to the scale of variation of the laser envelope. Self-focussing and filamentation may lead to the narrowing of the laser pulse, while the interaction of accelerated electrons with the laser pulse may lead to large excursion lengths, as the electrons observe a downshifted laser frequency. In the bubble regime, this interaction may resonantly drive betatron oscillations, leading to a significant enhancement of the radiation emitted by accelerated electrons [50].

The quasi-static approximation for the laser field breaks down in regimes where reflection or backscatter occur, or in situations in which there is a counterpropagating laser pulse, e.g. to cause electron injection into the wakefield [88].

5.3 Original aPIC model

5.3.1 Introduction

The envelope-PIC model does not readily lend itself to the study of Raman backscatter, as counterpropagating components of the laser field clearly cannot be suppressed. The smallest spatial scales are those of the laser beat and the plasma wave it drives, rather than the laser wavelength. PIC models typically make use of the Yee solver, which, as discussed in Chapter 3, may be considered an inefficient choice for simulating Raman backscatter; the plasma wave varies rapidly in space but slowly in time, but the Courant condition requires the same high resolution be used for both.

The aPIC model [4] uses several techniques and approximations to overcome these difficulties. While the laser beat has a small spatial scale, the two laser pulses taken individually vary slowly compared with the laser wavelength, and so may be modelled by separate envelope solvers, with the beat calculated from the two. The constraints of the Courant condition are avoided by assuming that the electron evolution is sufficiently slow to be treated electrostatically, thus allowing a Poisson solver to be used.

The laser pulses are solved on a coarse grid, with the electron motion solved on a fine grid. The same long timestep is used for both, satisfying the Courant condition for the laser solver. The laser interacts with the *average*

current inside a coarse grid cell, hence the name “aPIC”. The ability to use a longer timestep results in significant computational speedup (typically by a factor of 10 to 100). Furthermore, the use of the quasi-static approximation for the electron response, typical in envelope-PIC models, readily allows the implementation of a moving-window geometry for counterpropagating pulses. This allows simulations over long distances to be carried out, such as for the capillary experiments at Strathclyde, which would otherwise be unfeasible.

5.3.2 Poisson solver

The Poisson equation is not subject to the Courant condition for stability as it has no time dependence (although the timestep must still be sufficiently small to resolve the electron evolution). However, the assumption that the electron behaviour may be treated as electrostatic has the caveat that the model does not include the magnetic field due to particle motion, nor the retardation of the electric field of relativistic electrons.

For study of Raman amplification, these assumptions are not prohibitive. The magnetic field due to particle motion may be neglected if the $v \times B$ force due to the induced magnetic field is small. In the planar case, there is no contribution, and the contribution will be small if the transverse variation of electron velocity is small. This assumption is generally valid in regimes where the quasi-static approximation is valid, which requires the transverse motion of electrons be small compared with the variation of the laser field. In the Raman process, the longitudinal velocity is typically much less than the speed of light, due to the plasma wave breaking when electrons exceed the phase velocity of the wave, so retarded fields can readily be neglected. Indeed, the approximation only requires that the evolution of relativistic structures is slow. For example, in wakefield acceleration, individual electrons may be rapidly accelerated, but the wake and accelerated electron bunch typically

evolve slowly, and therefore could be modelled with reasonable accuracy using an electrostatic treatment.

The main limitation of the Poisson solver arises not from the assumption that the process can be treated as electrostatic, but that the solver is only computationally viable for PIC modelling in the planar case. For planar systems, the potential difference can be calculated using only the number of particles between two points. For higher-dimensional systems, the potential at a point must be calculated from the position of all particles. While some approximations could be envisaged, for example interacting with only the total cell charge for distant particles, an electrostatic treatment would likely be far more computationally intensive than the conventional Yee solver.

5.3.3 aPIC equations

The governing equations are similar to those derived in Section 5.2, except, as in Appendix A.2, we split the laser field into forward and backward propagating components. From Eq. (A.17), we substitute for separate pump and probe envelopes, to obtain:

$$\left(\frac{\partial^2}{\partial t^2} - c^2 \frac{\partial^2}{\partial z^2}\right) a_0 e^{i\phi_0} + \left(\frac{\partial^2}{\partial t^2} - c^2 \frac{\partial^2}{\partial z^2}\right) a_1 e^{i\phi_1} = -\frac{ne^2}{\varepsilon_0 \gamma m} (a_0 e^{i\phi_0} + a_1 e^{i\phi_1}). \quad (5.4)$$

We seek to separate the evolution equations for pump and probe. An envelope equation for the pump may be obtained by multiplying by $e^{-i\phi_0}$, and then averaging over a laser beatlength, $\lambda_b = 2\pi/(k_0 + k_1)$. For low growth rates, this causes the oscillating probe term to sum to zero. Applying a similar treatment for the probe, we obtain:

$$\begin{aligned} \left(\frac{\partial^2}{\partial t^2} - c^2 \frac{\partial^2}{\partial z^2}\right) a_0 + 2i\omega_0 \left(\frac{\partial}{\partial t} - c \frac{\partial}{\partial z}\right) a_0 &= -\frac{e}{mc\varepsilon_0 \lambda_b} \int_{\lambda_b} (a_0 + a_1 e^{-i\phi_b}) dz, \\ \left(\frac{\partial^2}{\partial t^2} - c^2 \frac{\partial^2}{\partial z^2}\right) a_1 + 2i\omega_1 \left(\frac{\partial}{\partial t} + c \frac{\partial}{\partial z}\right) a_1 &= -\frac{e}{mc\varepsilon_0 \lambda_b} \int_{\lambda_b} (a_0 e^{i\phi_b} + a_1) dz, \end{aligned} \quad (5.5)$$

where $\phi_b = \phi_0 - \phi_1$ is the ponderomotive phase. Unlike Appendix A.2, we

choose the carrier phase to satisfy the vacuum dispersion relation³, $\omega_{0,1}^2 = c^2 k_{0,1}^2$.

The current term on the right-hand side of the evolution equation for each laser pulse has two components: a self-current (the plasma response to the laser pulse in question), and a source current driven by the counterpropagating pulse. Dispersive effects may be neglected by removing the self-currents and the second derivatives of the envelope fields, to give the final form of the envelope equations:

$$\begin{aligned} \left(\frac{\partial}{\partial t} - c \frac{\partial}{\partial z} \right) a_0 &= \frac{i}{2\omega_0} \frac{ne^2}{m\varepsilon_0\lambda_b} \int_{\lambda_b} \frac{a_1 e^{-i\phi_b}}{\gamma} dz, \\ \left(\frac{\partial}{\partial t} + c \frac{\partial}{\partial z} \right) a_1 &= \frac{i}{2\omega_1} \frac{ne^2}{m\varepsilon_0\lambda_b} \int_{\lambda_b} \frac{a_0 e^{i\phi_b}}{\gamma} dz. \end{aligned} \quad (5.6)$$

Eq. (5.6) is implemented in the aPIC code by summing over all electrons in a coarse grid cell, setting the cell size to a laser beatlength:

$$\begin{aligned} \left(\frac{\partial}{\partial t} - c \frac{\partial}{\partial z} \right) a_0 &= \frac{i}{2\omega_0} \frac{e^2}{\varepsilon_0 m} \eta \sum_j \frac{a_{1j} e^{i\phi_j}}{\gamma_j}, \\ \left(\frac{\partial}{\partial t} + c \frac{\partial}{\partial z} \right) a_1 &= \frac{i}{2\omega_1} \frac{e^2}{\varepsilon_0 m} \eta \sum_j \frac{a_{0j} e^{-i\phi_j}}{\gamma_j}, \end{aligned} \quad (5.7)$$

where $a_{0j,1j}$ and ϕ_j are the envelope fields and ponderomotive phase interpo-

³Any carrier phase may be chosen, with the complex envelope chosen to give the correct initial conditions. Choosing the vacuum values in this case allows simplification of the governing equations.

lated to the position of the j th particle, with γ_j its Lorentz factor:

$$\gamma_j = \sqrt{1 + \frac{p_j^2}{m^2 c^2} + \frac{1}{2} (|a_0|^2 + |a_1|^2 + 2\Re[a_0 a_1^* e^{i\phi_j}])}. \quad (5.8)$$

Ignoring the contribution due to the gradient of the individual envelope fields, which is small compared with the influence of the beat term, the particle push is given by: ⁴

$$\frac{dp_j}{dt} = -eE + \frac{mc^2}{2\gamma} (k_0 + k_1) \Im[a_0 a_1^* e^{i\phi_j}]. \quad (5.9)$$

⁴For circularly polarised pulses, it is possible to model the ponderomotive force and Lorentz factor using a long timestep due to the low phase velocity of the ponderomotive beat. For linearly polarised pulses, the ponderomotive force and Lorentz factor vary on the scale of the plasma frequency, and so it is only possible to model their average values if the computational advantages of the envelope model are to be retained.

5.3.4 Laser field update

In the original aPIC code, the laser field update was performed using an implicit solver, the centred-grid backward differential:

$$\frac{a_m^{n+1} - a_m^n}{\Delta t} \mp c \frac{a_{m+1}^{n+1} - a_{m-1}^{n+1}}{2\Delta z} = -\frac{i}{2\omega} j_m^n, \quad (5.10)$$

where a_m^n is the quantity a at grid cell m at timestep n . a and j are the envelope amplitude and source term from Eq. (5.7), and \mp relates to the pump and probe, respectively. a_m^{n+1} is found via a matrix inversion technique.

However, while such implicit methods have no stability condition, they introduce some degree of numerical diffusion, a nonphysical effect causing the laser field to spread out. For short pulses, the laser envelopes vary on short length scales, and so the effect of numerical diffusion can be significant. For parameters of interest, a small timestep ($c\Delta t \sim \Delta z/10$) is required for the laser field update to prevent nonphysical results.

To minimise the computational overhead of a small timestep, the original aPIC code implemented two techniques, supercycling of the laser update and a fine grid for the probe envelope. Supercycling allows the laser update to be run several times every timestep, i.e. N repetitions, with correspondingly reduced increments of $\Delta t/N$. As the particle updates are responsible for the bulk of the computational overhead, carrying out multiple laser updates

for each particle update allows a reduction of numerical diffusion without a significant increase in computation time. Using a smaller grid size also reduces numerical diffusion. As the probe pulse is typically many times shorter than the pump, a fine grid is implemented only for the probe. This fine grid is used for the laser update only, again resulting in only a small additional overhead.

5.3.5 Pseudo-2D model

A 2D aPIC model has been implemented [57] as a series of parallel 1D solvers, coupled by diffraction. The assumption implicit in this model is that the transverse variation of the system is small, and that the slowly-varying transverse electron motion may be neglected. These are reasonable approximations for the Raman process, where the initial plasma density profile can be assumed to be in equilibrium with the long pump pulse. As the probe pulse duration is typically several times shorter than its width, the transverse drift of electrons may be neglected over the timescales of interest.

Starting from Eq. (A.15), and retaining transverse components of the Laplacian, while keeping the assumption that the plasma dispersion may be ignored (neglect self-currents and second derivatives of the envelopes in t and z), Eq. (5.6) becomes:

$$\begin{aligned} \left(\frac{\partial}{\partial t} - c \frac{\partial}{\partial z} \right) a_0 &= \frac{i}{2\omega_1} \frac{ne^2}{m\varepsilon_0\lambda_b} \int_{\lambda_b} \frac{a_1 e^{-i\phi_b}}{\gamma} dz - \frac{ic}{2k_1} \nabla_{\perp}^2 a_0, \\ \left(\frac{\partial}{\partial t} + c \frac{\partial}{\partial z} \right) a_1 &= \frac{i}{2\omega_1} \frac{ne^2}{m\varepsilon_0\lambda_b} \int_{\lambda_b} \frac{a_0 e^{i\phi_b}}{\gamma} dz - \frac{ic}{2k_1} \nabla_{\perp}^2 a_1. \end{aligned} \quad (5.11)$$

Here ∇_{\perp}^2 is the transverse Laplacian, $\nabla^2 = \partial^2/\partial z^2 + \nabla_{\perp}^2$.

5.4 Extensions to the aPIC model

5.4.1 Chirps

In general, envelope models may easily be extended to include chirped pulses by the use of complex envelopes (see Appendix A.2). Separate evolution equations for the pump and probe envelopes may be obtained by averaging over a beat length, as in Eq. (5.5), which removes non-resonant components. However, in the case of chirped pulses, the beat length varies. This makes the calculation more difficult in the aPIC model, because currents are calculated numerically by summing over all particles within a beat length.

While the chirp rate of the laser pulses is known, there are many cases where the instantaneous laser frequency is not explicitly known, for example due to finite-bandwidth laser pulses or nonlinear frequency shifts. If the currents need to be averaged over a beatlength, it would greatly reduce the applicability of the aPIC model.

By way of analogy, we consider the general case of electrons radiating in a laser field. For an infinitely long laser pulse interacting with uniformly distributed electrons, radiation in the counter-propagating direction will completely destructively interfere. Radiation in the co-propagating direction leads to dispersion of the laser pulse. If the electron density is modulated with a periodicity equal to the laser wavelength, the degree of destructive

interference is reduced. While the radiation emitted from the density peaks still interferes destructively with that from the troughs, there are more electrons in the peaks, and so there is net emission in the counter-propagating direction. The interference will manifest as an oscillation of the radiation intensity, with a periodicity equal to the bunching length, but with an overall increase in intensity.

Raman amplification can be considered in the same manner, with the difference that the density perturbation moves, which changes the required periodicity of the density perturbation if net growth is to occur, and changes the frequency of the emitted counter-propagating radiation. The aPIC model smooths out the oscillations in radiation intensity on the scale of the bunching length by averaging over a beat length, giving only the net gain. However, in the case where the cell length is not equal to the beat length, there will be some oscillation of the scattered field, depending on the local cell. However, as with the fully resolved case, the general trend should show the correct gain.

To test that the aPIC code gives the correct net gain, regardless of the periodicity over which averaging is carried out, the code was modified to allow the use of an arbitrary cell size. Figure 5.1 shows an amplified pulse for different cell sizes. As the difference between cell size and beatlength increases, an additional short-wavelength oscillation of the envelope may be




Figure 5.1: **Influence of aPIC averaging length on probe laser field.** Amplified probe after 2500 steps for different grid sizes. Note that due to the difference in timestep, the snapshots correspond to different times, for easier comparison.

observed. However, apart from this numerical artefact, the growth rates are unaffected.

As this oscillatory behaviour is negligible for small mismatch, the code was extended to allow chirped pulses, achieved by modifying the initial pump and probe envelopes. However, numerical diffusion introduced by the original implicit solver leads to artificial damping of the modulated envelopes, which could only be suppressed through the use of a small timestep ($c\Delta t \ll \Delta z$). This has been rectified through the development of a fully explicit laser update, described in Section 5.4.3.

5.4.2 Wake fields

Laser pulses propagating through plasma will drive a wake [89]. This may become an important consideration for Raman amplification, as the presence of the wake can act to reduce the growth rate, while driving the wake will cause the laser pulse to lose energy.

Modifying the code to include wake terms is readily achieved by including the envelope gradients in Eq.(5.9), which were neglected in the original model:

$$\frac{dp_j}{dt} = -eE - \frac{mc^2}{4\gamma} (\nabla|a_0|^2 + \nabla|a_1|^2 - 2(k_0 + k_1) \Im[a_0 a_1^* e^{i\phi_j}]). \quad (5.12)$$

However, while this will drive a wake, modifying the Raman growth rates, it will not model energy lost by the laser field as it is transferred to the wake. Photons are not absorbed during the process, so this loss of energy cannot be modelled simply by decreasing the laser amplitude. Instead, the loss of energy manifests as a decrease in the laser frequency. This occurs through the process of photon deceleration, as the laser pulse interacts with the co-moving density gradient of the wake. This process may be included in the model by extending the laser solver to include dispersion, discussed in Section 5.4.4.

5.4.3 Advective laser update

Numerical diffusion due to the implicit update, described in Section 5.3.4, led to artificial damping of chirped laser pulses. It was found that even with supercycling, the timestep typically had to satisfy $c\Delta t < \Delta z$, removing the advantage of an implicit solver.

For this purpose, a fully explicit advective solver, as described in Appendix A.3.1 was implemented:

$$a_m^{n+1} = a_{m\mp 1}^n - \frac{i}{\omega} \Delta t j_m^n, \quad (5.13)$$

where the timestep is chosen such that $\Delta t = \Delta z/c$. As before, m and n are the spatial-step and timestep number, and \mp correspond to probe and pump, respectively.

This solution is numerically stable and does not suffer from numerical diffusion. While the timestep cannot be varied independently of the spatial step, in practice the timestep is larger than that required by the original implicit solver. The advective solver is therefore less computationally intensive, while removing the need for supercycling or a fine grid, making it an excellent choice for the dispersionless update.

5.4.4 Dispersive laser update

A dispersive update is necessary to allow the self-consistent modelling of a laser pulse driving a wake, as discussed in Section 5.4.2. It is also an important consideration for higher dimensional systems, such as the pseudo-2D model described in Section 5.3.5. For the capillary experiments at Strathclyde, the plasma channel is expected to play a role in guiding the laser pulse. Focussing and defocussing arise through the relative positions of the phase-fronts of the laser pulse, hence simulation of these processes requires a dispersive model.

As discussed in Section 5.2, solutions to the full second-order wave equation generally do not allow a reduced resolution to be used in envelope-PIC models, as backward propagating components of the laser pulse require a high resolution. The paraxial approximation is commonly used to allow dispersive behaviour while ignoring backward propagating solutions. However, as discussed in Appendix A.3.2, this approximation relies on prescribed values for the group velocity and wavenumber of the laser pulse, and so is unsuitable for nonlinear regimes where the laser frequency or plasma susceptibility vary significantly. This is a concern for simulations of Raman amplification, as nonlinear or off-resonance amplification may lead to a significant shift in the laser frequency. For this reason, an approximation to the wave equation

derived by Gordon *et al* [90], detailed in Appendix A.3.3, has been implemented for the laser solver. It avoids the use of prescribed values for the group velocity and wavenumber.

Simulations were carried out for a single laser pulse propagating in homogeneous plasma. In this case, a full wave equation solver, detailed in Appendix A.3.4, requires only a low resolution, as there should be no reflection or backscatter. Comparisons have been made between the full solver, paraxial approximation and Gordon approximation.

Figure 5.2 shows simulation results approximately 7 and 35 ps after entering the plasma, with parameters as for the probe in Table 3.1. As can be seen, the agreement between the three models for the group velocity is excellent. However, plots of the real part of the envelope field show that the phase velocities predicted by the Gordon and paraxial approximations are lower than that of the full solver. For the parameters used, the position of a point of constant phase predicted by the Gordon approximation lags the full solver by 30 nm after 35 ps. Following the simulation evolution shows that, over the same time, the paraxial approximation develops a lag of 200 nm. Despite this difference in phase, it may be seen from the periodicity of the envelope modulation that the agreement for the wavenumber is good for all three models.

The error in the phase was found to vary linearly with the plasma density.

While the error for the Gordon approximation appears small, 30 nm after ~ 1 cm of propagation, the error in the light frame is significant. Using the homogeneous plasma dispersion relation, for the parameters used, a point of constant phase should precess by approximately 60 nm over 35 ps in plasma. The error in the phase velocity predicted by the Gordon model, relative to the vacuum speed of light, is therefore 50%.

As the Gordon approximation gives the correct wavenumber and group velocity for the laser pulse without relying on prescribed values, it is a strong candidate for simulations of wakefield acceleration in scenarios where focussing effects are not important. However, despite the increased accuracy compared with the paraxial approximation, the error in the phase makes it a poor choice for modelling the guiding of a laser pulse. The model is similar to those derived in [87, 91], in that it neglects the second-derivative in time in the light frame. Further work is required to determine if the additional approximation made in the Gordon model (c.f. Appendix A.3.3) contributes to this error in phase, or if it is primarily a consequence of neglecting the second-order temporal derivative.

The error in phase was found to have a negligible effect on the Raman growth rates, making the Gordon approximation suitable for the planar model. The numerical implementation designed for the aPIC code is fully explicit, and so will not suffer from numerical diffusion, making it suitable

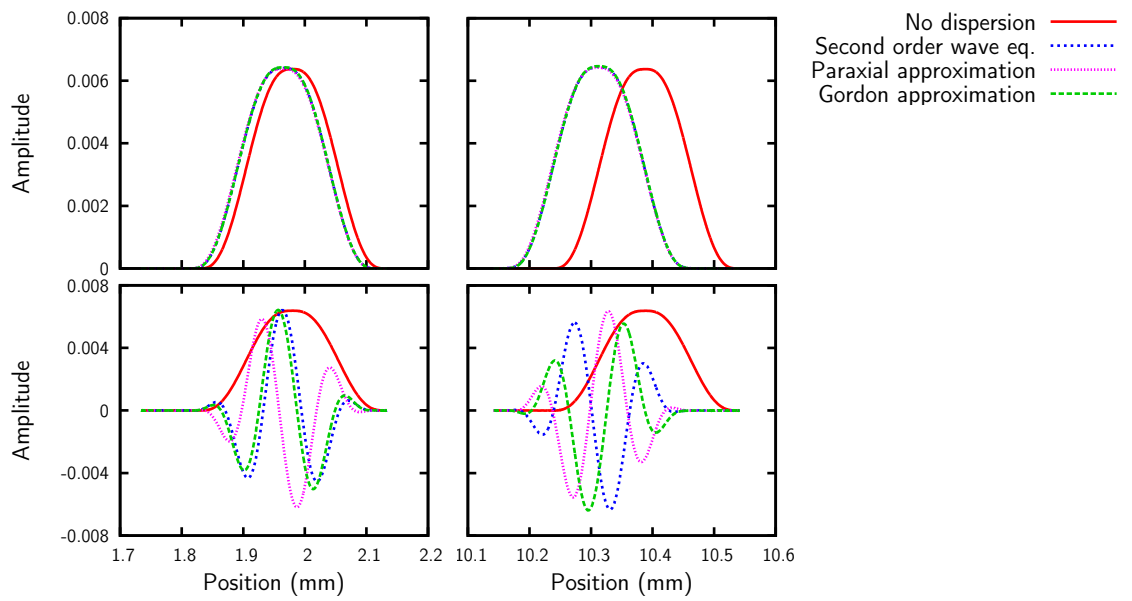


Figure 5.2: **Comparison of laser solvers.**

Absolute and real amplitudes of a laser pulse ~ 7 and ~ 35 ps after entering the plasma. Parameters as for probe in Table 3.1.

for chirped pumps. However, further work is required to find a suitable dispersive model for modelling guiding of a laser pulse in the pseudo-2D model.

5.5 Stability and accuracy

5.5.1 Numerical stability

As with conventional PIC codes, which are discussed in Chapter 3, the resolution of aPIC simulations must be sufficiently high to model the physical processes involved. Figure 5.3 shows the influence of increasing the number of cells used in the electrostatic solver, keeping the total macroparticle number constant. Simulation parameters are shown in Table 3.1. Just as in the case of conventional PIC codes, insufficient resolution leads to numerical instability (c.f. Figs. 3.3,3.4). However, increasing the resolution only increases the accuracy of the model to a point, beyond which the results diverge from the correct value. This behaviour is consistent with numerical scattering [56], in which the sharp gradients in the electrostatic potential which arise due to a low number of particles per cell can act to scatter particles.

The effect of increasing the resolution while keeping the number of particles per cell constant is shown in Figure 5.4, which shows the pump amplitude between 650 and 700 μm . The accuracy of the model is again seen to improve to a point, beyond which it diverges. This is attributed to the use of too coarse a temporal resolution, allowing particles to traverse many cells per timestep. This reduces the accuracy of the code, as the electrons respond to the field in their initial cell, rather than the average field over the step. Using

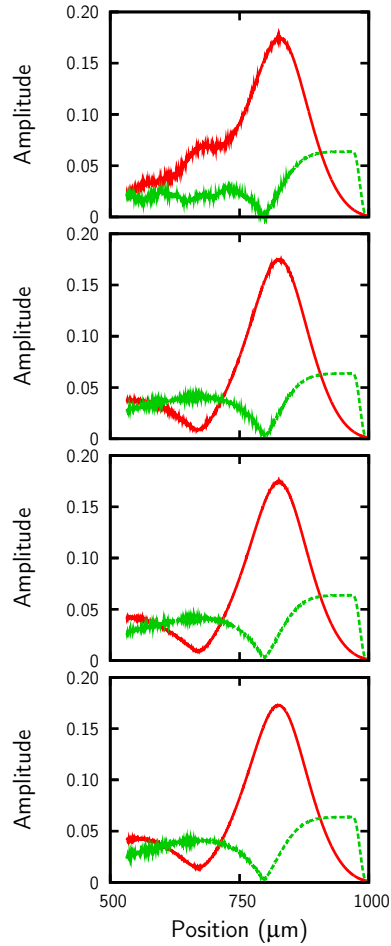


Figure 5.3: **Influence of simulation resolution on numerical stability (constant macroparticle number).**

Pump and probe envelopes for aPIC simulations with varying spatial resolution for the electrostatic solver. Resolution, from top to bottom, of 50, 100, 200 and 400 cells per pump wavelength. Timestep of $0.98 \lambda_b/c$, with other parameters as in Table 3.1. Total macroparticle number is kept constant as resolution is increased.

a smaller timestep, as shown in Figure 5.5, is seen to improve the accuracy of the model.




Figure 5.4: **Influence of simulation resolution on numerical stability (scaled macroparticle number).**

Probe envelope over a short range for aPIC simulations with varying spatial resolution for the electrostatic solver. Resolution, from left to right, of 100, 200 and 400 cells per pump wavelength. Timestep of $0.98 \lambda_b/c$, with other parameters as in Table 3.1, but with macroparticle number per cell kept constant as resolution is increased.




Figure 5.5: **Influence of temporal resolution on numerical stability (constant macroparticle number).**

Probe envelope over a short range for aPIC simulations with varying temporal resolution. Timestep, from left to right, of 0.98 , 0.49 and $0.245 \lambda_b/c$. Other parameters as in Table 3.1.

5.5.2 Benchmarking

Benchmarking was again carried out against OSIRIS. Figure 5.6 shows the pump and probe envelopes for OSIRIS and aPIC for different pump/probe detuning. The OSIRIS results are shown for the full plasma length, while the aPIC results have made use of a $500 \mu\text{m}$ moving-window. The agreement for the resonant case is excellent. Some fine structure may be seen in the aPIC results, attributed to a shift in the probe frequency, similar to that observed in Section 5.4.1 for chirped pulses. However, for the non-resonant case, some disparity is observed; for the case of detuning $< \omega_p$, the aPIC simulations predict lower growth than OSIRIS, while for detuning $> \omega_p$, aPIC simulations predict larger growth.

Increasing the spatial and temporal resolution while keeping the number of particles per cell constant, as shown in Figure 5.7, shows that this disparity is not a consequence of the resolution. As the resolution is increased the aPIC result converges, but the converged value does not agree with that predicted by OSIRIS. Simulations carried out with the advective solver show this disparity is not due to the phase error introduced by the dispersive laser solver, discussed Section 5.4.4. The disparity is attributed to the particle push used in aPIC, which is fully explicit. Analogy may be drawn with the case of a charged particle moving in a magnetic field, which is known to

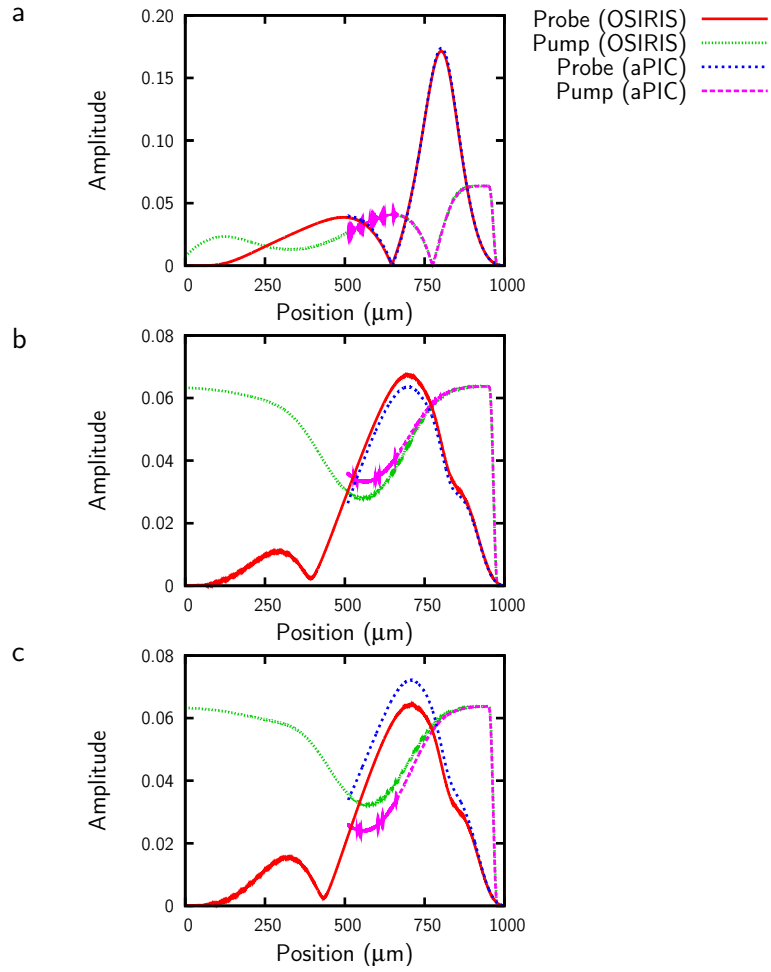


Figure 5.6: **Benchmarking for varying detuning.** Pump and probe envelopes from OSIRIS and aPIC for a detuning of a) ω_p , b) $0.95\omega_p$, and c) $1.05\omega_p$. Timestep of $0.245\lambda_b/c$, with other parameters as in Table 3.1.




Figure 5.7: **Influence of simulation resolution.**

Plots of amplified probe for varying resolution. “1x” corresponds to a timestep of $0.98 \lambda_b/c$, with other parameters as in Table 3.1. Temporal and spatial resolution are increased, with macroparticle number per cell kept constant.

converge to an incorrect value if a fully explicit solver is used [92].

This error can therefore only be removed by modifying the particle-push method used in the aPIC model. While the disparity is not large, improving the particle push to use a semi-implicit method should be considered for future work, to improve the accuracy of the code.

5.6 Model demonstration: nonlinear plasma waves

The goal of this work has been the development of the aPIC code to allow simulation of the experimental work at Strathclyde, e.g. [49, 51]. However, we here include a short case study to illustrate the functionality of the model. We illustrate a new regime of broad-bandwidth amplification, achieved through the exploitation of nonlinear plasma waves.

Two laser pulses are collided in plasma, a forward propagating, broad-bandwidth, chirped probe pulse, and a backward propagating, monochromatic pump pulse, both centred about the same frequency, ω_0 . In the linear regime, the probe will lose energy at $\omega_0 + \omega_p$, and gain energy at $\omega_0 - \omega_p$, as it interacts with the pump through stimulated Raman backscatter. The resonant plasma wave excited as the probe gains energy will have a negative phase velocity, while that excited as the probe loses energy will have positive phase velocity.

The plasma response will vary with position, as the laser pulses have finite duration. At the centre, however, the entirety of the two laser pulses will interact. If the amplitude of the excited plasma wave is small, the plasma can support the linear superposition of the two resonantly excited plasma waves. The spectral gain of the probe does not depend on whether the initial

probe has positive or negative chirp.

The situation becomes different if the plasma wave amplitude becomes high. The plasma can no longer support a superposition of the two excited waves, and so the the plasma response will depend on the order in which the beatwaves are excited. The spectrum of the probe after the interaction will therefore depend on the chirp.

Simulations have been carried out using the aPIC code, for a plasma density of $1 \times 10^{18} \text{ cm}^{-3}$ and a temperature of 5 eV, using 1 ps duration laser pulses, centred about 800 nm, with a probe bandwidth of 60 nm. For low amplitude laser pulses ($a = 0.005$), the driven plasma wave remains predominantly linear, and the probe gain spectrum shows only small variation as the chirp is reversed, as seen in Fig. 5.8. However, at higher intensities, ($a = 0.01$), the interaction becomes highly nonlinear, despite the laser intensity remaining below the threshold for the Compton regime. In this case, broad bandwidth energy transfer is achieved to/from the probe, depending on whether the chirp is positive or negative, i.e. corresponding to the direction of energy transfer at the front of the probe, with energy transfer in the opposite direction suppressed.

This behaviour can be explained by considering the excited plasma wave. The plasma cannot support a superposition of the two plasma waves, and so the varying phase velocity of the laser beat acts to reverse the phase velocity

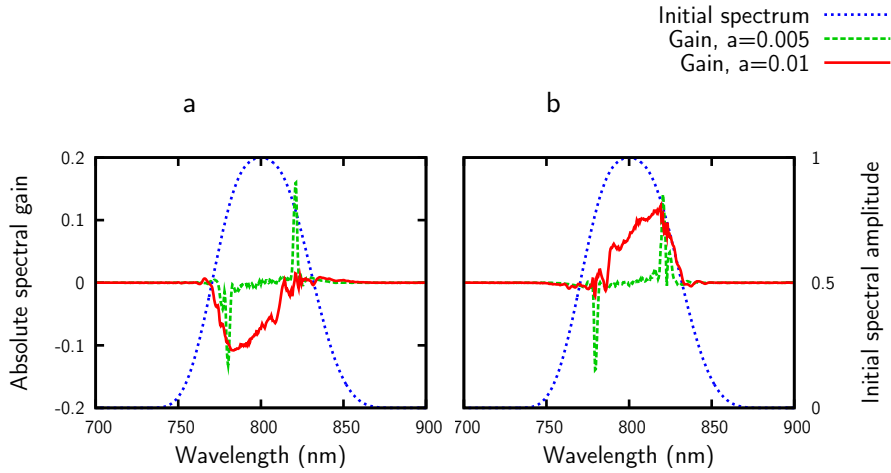


Figure 5.8: **Influence of probe chirp on amplification.**

Plots of absolute probe gain over the interaction for a) down- and b) up-chirped probes. Initial pulse spectrum shown for comparison. Amplitudes are normalised to the peak spectral amplitude of the initial pulse.

of the excited plasma wave. As the phase velocity of the excited plasma wave varies, it allows amplification over a broad bandwidth. In the linear case, the excited plasma waves do not interact, and so their individual phase velocities remain unchanged. This behaviour is similar to autoresonance, which is known to occur for Raman amplification in nonuniform plasma [93].

The electron phase space in the interaction region, a sample of which is illustrated in Fig. 5.9 for the case of a down-chirped probe, shows many interesting nonlinear effects. The formation of a separatrix, as discussed in Chap 4, is readily apparent. At the centre of the interaction, a plasma wave is first excited with a positive phase velocity. As the amplitude increases, some electrons gain sufficient velocity to exceed the phase velocity of the structure, and cross the separatrix. However, as the probe frequency at that

point decreases, the phase velocity of the wave reverses. The electrons above the separatrix then form the bulk of the electron distribution.

The interaction to the rear of the probe, shown in the left-most plot in Figure 5.9, is characterised by an accelerating (increasingly negative) phase velocity. Electrons that initially have sufficient velocity to cross the separatrix are seen to become trapped as the structure accelerates. The interaction to the front of the probe, shown in the right-most plot, is characterised by a decelerating (decreasingly positive) phase velocity. This results in an increase in the number of electrons with sufficient velocity to cross the separatrix. However, as electrons cross the separatrix, they gain energy from the electrostatic wave, acting to suppress it. When this effect becomes large, the potential between adjacent electrostatic buckets may be eroded sufficiently as to allow them to merge, which can be seen to give rise to vortex-type behaviour at late times.

The nonlinear broadening of the Raman resonance observed here due to nonlinear plasma waves may be of practical use, allowing the amplification of broad bandwidth pulses using a monochromatic pump without relying on pump depletion [42] or entering the Compton Regime [17]. It is also an important consideration for the analysis of spectral data from experimental works, such as in the measurement of plasma density by colliding pulses [49].




Figure 5.9: **Phase space evolution.**

Sample electron phase space, showing snapshots of the electron distribution (velocity against position) at different times and positions. Phase spaces are shown every 133 fs over 2 μm intervals, to the rear, centre and front of the interaction region, separated by 80 μm .

5.7 Conclusion

The aPIC code has been extended to allow simulation of the experiments carried out at Strathclyde.

Chirped pulses have been implemented through the use of complex envelopes. The original aPIC model operates by averaging of currents over a ponderomotive beat, which is not possible with the inclusion of chirped pulses, as the beatlength is not explicitly known. However, simulations show that a difference between averaging length and beatlength leads only to an additional short wavelength oscillation of the laser envelope. Growth rates are not affected beyond this artefact. This result also shows the aPIC code should be valid in nonlinear regimes where the laser frequency evolves from its original value. A fully explicit advective solver has been implemented to prevent damping of chirped laser pulses due to numerical diffusion.

The code has also been extended to include ponderomotive contributions due to a single laser pulse. These are necessary to model wakefield generation and Raman forward scattering, both of which may act to limit the efficiency of the Raman process. In order to model these processes self-consistently, a new dispersive laser solver was implemented. Comparisons against a solver based on the full wave equation (unsuitable in regimes where backscatter or reflection may occur) show the dispersive approximation chosen gives excel-

lent agreement for the group velocity of the laser pulse, but introduces an error in the phase of the pulse. The solver is therefore suitable for planar cases or regimes in which guiding of the laser pulse is negligible.

Benchmarking against OSIRIS shows good agreement, with some disparity seen for amplification away from resonance. This was attributed to the use of a fully-explicit particle push, which causes simulations to converge to an incorrect value with increasing resolution.

A case study was carried out using the aPIC code, in which a new regime of broad bandwidth amplification was identified. The increase in amplification bandwidth was achieved through the excitation of a high amplitude plasma wave, such that the plasma could not support a superposition of plasma waves with different wavenumbers. While this new regime may have practical applications, it is of particular relevance for the analysis of experimental data.

5.8 Future work

The accuracy of the model could be improved through the implementation of a semi-implicit particle update, removing the disparity in growth rate observed for off-resonance amplification.

The modifications made to the code can readily be applied to the existing pseudo-2D model, which would allow the impact of processes such as gain-focussing to be investigated. The development of an improved dispersive solver for the laser update would enable accurate simulation of the guiding of a laser pulse in a plasma channel, which is relevant to the experimental work at Strathclyde.

The moving window geometry would readily allow preheating of the plasma by the pump laser to be emulated, as the temperature of fresh plasma created at the window edge could be modified. However, this approach is limited, as it does not allow collisional heating within the window.

However, the aPIC code is now sufficiently advanced that emphasis can shift from code development towards using it as a practical tool for simulations of Raman amplification in plasma. The use of the code as a simulation tool will therefore represent the majority of future work on this topic [94].

Chapter 6

Conclusions and future work

6.1 Conclusions

Three different models have been investigated for the simulation of Raman amplification in plasma: a Particle-in-Cell code; a three-wave model, derived from a fluid treatment of plasma, extended to include the influence of wave-breaking and particle trapping; and the aPIC model, originally developed by MinSup Hur [4].

Conventional (Yee) PIC models solve Faraday's and Ampere's laws on a grid using finite-difference schemes, and require the finest approximations of the models investigated. The resolution must be sufficiently high to model the physical processes of interest, and avoid the growth of numerical instabilities, which results in large computational overheads. Simulations in this work are therefore limited to the planar case. A new numerical instability relevant to Raman amplification is identified, which further increases the resolution required to achieve convergent results. The omission of collisions in current PIC codes limits their applicability, both through the absence of processes that can impact on Raman amplification, such as heating, and through artificially high instability growth rates, which occur due to the absence of collisional damping. The absence of collisional effects will also limit the accuracy with which Landau damping is simulated.

This contrasts with three-wave models, which offer low computational

overheads, but at the expense of not including particle effects. The influence of such effects may be included phenomenologically. A one-dimensional three-wave model that includes heating and damping effects has been extended, to include the influence of harmonics and wavebreaking of the plasma wave in cold plasma, and the effects of particle trapping on the effective amplitude of the wave in warm plasma. However, such modifications require the existence of a suitable analytical model. The main limitation of the model developed in this work is the linearisation of Landau damping rates, which become inaccurate when the perturbation to the assumed electron velocity distribution becomes large.

The aPIC model makes use of several approximations and incorporates a moving-window geometry to provide many of the advantages of a PIC code, but with a significantly lower computational overhead. The model has been extended to improve numerical accuracy, allow the use of chirped laser pulses, model wake generation and Raman forward scattering, and approximate the dispersion of the laser pulses. The main limitation for the simulation of Raman amplification is found to be the approximations that are required for the laser solver, which give inaccurate phase velocities of the plasma wave due to dispersion. This limits the applicability of the model to the planar case, by preventing accurate modelling of laser guiding and self-focussing in the existing pseudo-2d model. A numerical error arises due to the fully

explicit particle push used in the code, although the impact is small. As for conventional PIC models, aPIC does not include collisional effects.

6.2 Future work

The addition of collisions to conventional PIC models would improve both their accuracy and applicability, although at the cost of further increasing computational overheads. The development of a moving-window geometry suitable for Raman amplification would allow significant reductions in the computational overhead of simulations. Increased availability of computational resources will also act to ease this constraint.

The three wave model could be further extended to model Landau damping in nonlinear regimes. Other works, e.g. [77], use a three-wave model which simulates the evolution of the electron velocity distribution to accurately model the Landau damping rate and frequency shift. However, such an approach leads to a significant increase in computational overhead, although this remains significantly less than that of PIC models. As is the case with PIC codes, the model is currently limited by the absence of collisional effects. Future work to include collisional effects in such models would remove this limitation. Alternatively, an analytical model for nonlinear Landau damping could be developed, which would retain the low computational overhead of

the model developed in this work. The three-wave model could be extended to higher dimensionality, but this would require the implementation of a dispersive laser solver. The inclusion of heat-transport mechanisms in such a model would also include the accuracy of the plasma temperature due to pre-heating by the pump laser.

The aPIC model could be improved through the development of a laser solver that more accurately models the phase velocity. Implementation in the pseudo-2D model would allow self-focussing and guiding effects to be accurately simulated. The development of an improved particle-push should be considered to increase the accuracy of the code. The impact of pre-heating of the plasma due to collisional damping of the pump could readily be approximated by making use of the moving-window geometry, although other collisional effects are not easily included.

6.3 Applicability to experiments

For the experimental parameters used at Strathclyde, simulations carried out with current conventional PIC codes would incur prohibitive computational overheads, due to the centimetre-scale interaction lengths, while the absence of collisional effects would reduce the accuracy of simulations in this parameter space.

The three-wave model requires further development to be applicable in the regimes of interest, in which Landau damping is not well described by the linearised rates. The availability of higher-dimensionality models would also be a great advantage to model processes relevant to the interaction.

The aPIC model has been found to be the most suitable, as simulations can be carried out with currently available computational resources, and it models Landau damping in nonlinear regimes. The model does not include the effects of collisions, but is no more inaccurate in this regard than conventional PIC models. While the error in dispersion limits the accuracy of the model to the planar case, the development of a laser solver with improved dispersion properties is an area of active research.

Appendix

A.1 Finite-difference schemes

A.1.1 Taylor expansions

If a is some continuous function, we can write its discretised value at time $n\Delta t$ and position $m\Delta z$, where n and m are integers, as a_m^n . We may write the Taylor expansions for a_m^{n+1} and a_m^{n-1} as:

$$\begin{aligned}a_m^{n+1} &= a_m^n + \Delta t \left[\frac{\partial a}{\partial t} \right]_m^n + \frac{\Delta t^2}{2} \left[\frac{\partial^2 a}{\partial t^2} \right]_m^n + O(\Delta t^3), \\a_m^{n-1} &= a_m^n - \Delta t \left[\frac{\partial a}{\partial t} \right]_m^n + \frac{\Delta t^2}{2} \left[\frac{\partial^2 a}{\partial t^2} \right]_m^n + O(\Delta t^3).\end{aligned}\tag{A.1}$$

Second order approximations of the first and second time derivatives can respectively be found by taking the difference and sum of these expansions and rearranging:

$$\left[\frac{\partial a}{\partial t} \right]_m^n = \frac{a_m^{n+1} - a_m^{n-1}}{2\Delta t} + O(\Delta t^3),\tag{A.2}$$

$$\left[\frac{\partial^2 a}{\partial t^2} \right]_m^n = \frac{a_m^{n+1} - 2a_m^n + a_m^{n-1}}{\Delta t^2} + O(\Delta t^3).\tag{A.3}$$

Similarly, expansions may be made for the spatial derivatives:

$$\left[\frac{\partial a}{\partial z}\right]_m^n = \frac{a_{m+1}^n - a_{m-1}^n}{2\Delta z} + O(\Delta z^3), \quad (\text{A.4})$$

$$\left[\frac{\partial^2 a}{\partial z^2}\right]_m^n = \frac{a_{m+1}^n - 2a_m^n + a_{m-1}^n}{\Delta z^2} + O(\Delta z^3). \quad (\text{A.5})$$

Derivatives correct to higher order may be found by making use of values at more points, e.g. at $m - 2$, $m + 2$, $n - 2$, $n + 2$. While this is viable for ordinary differential equations, solving a finite-difference time domain problem for $n + 2$ will likely give rise to intractable equations, so we here limit our treatment to second-order.

A.1.2 Numerical schemes

Many numerical implementations rely only on first-order derivatives. We consider the simple case of the set of partial-differential equations:

$$\begin{aligned}\frac{\partial a}{\partial t} &= \frac{\partial b}{\partial z}, \\ \frac{\partial b}{\partial t} &= \frac{\partial a}{\partial z},\end{aligned}\tag{A.6}$$

where both a and b are functions of t and z . In many cases it is preferable to use only values at the current timestep to extrapolate to the next, in order to reduce the memory demands of the simulation. We discretise a and b , and take Taylor expansions using only values at n , $n + 1$. This can equivalently be viewed as either a first-order Taylor expansion in time, or a second-order Taylor expansion in time with the temporal derivative evaluated at the midpoint $n + \frac{1}{2}$.

We may gain either explicit or implicit equations for a_m^{n+1} , analogous to the forward Euler and backward Euler methods for ordinary differential equations:

$$\frac{a_m^{n+1} - a_m^n}{\Delta t} = \frac{b_{m+1}^n - b_{m-1}^n}{2\Delta z},\tag{A.7}$$

$$\frac{a_m^{n+1} - a_m^n}{\Delta t} = \frac{b_{m+1}^{n+1} - b_{m-1}^{n+1}}{2\Delta z}.\tag{A.8}$$

Equivalent equations are developed for b . These equations suffer from additional numerical error, which can be attributed to either the mismatch in the time at which the derivatives are evaluated, or due to taking the Taylor expansion only to first-order.

More advanced alternatives exist for both explicit and implicit methods to avoid this issue. The Crank-Nicholson method [95] is an implicit method which avoids time-mismatch by interpolating to a half-step, allowing a time-centred second-order equation:

$$\left[\frac{\partial a}{\partial t}\right]_m^{n+\frac{1}{2}} = \frac{1}{2} \left[\frac{\partial b}{\partial z}\right]_m^n + \frac{1}{2} \left[\frac{\partial b}{\partial z}\right]_m^{n+1}. \quad (\text{A.9})$$

An alternative explicit algorithm is the leap-frog method, which gives fully-explicit time-centred equations. For the case of Eqs. (A.6), this may be achieved by evaluating a at integer timesteps and b at half-integer timesteps. a is advanced using the mid-point value of b , and vice-versa:

$$\begin{aligned} \frac{a_m^{n+1} - a_m^n}{\Delta t} &= \frac{b_m^{n+\frac{1}{2}} - b_m^{n-\frac{1}{2}}}{\Delta t}, \\ \frac{b_m^{n+\frac{3}{2}} - b_m^{n+\frac{1}{2}}}{\Delta t} &= \frac{a_m^{n+1} - a_m^n}{\Delta t}. \end{aligned} \quad (\text{A.10})$$

Explicit methods offer the advantage that the equations are readily solved, and the computational effort for a single timestep is small. However, they

suffer the limitation that they are only numerically stable when the Courant-Friedrichs-Lewy condition [65] (typically referred to simply as the Courant condition) is satisfied. This constraint for the maximum timestep (for Eqs. (A.6), $\Delta t < \Delta z$) may lead to a significant increase in the number of timesteps required.

Implicit methods are typically solved via matrix-inversion (some implicit methods rely on iteration), which is computationally intensive, and introduces numerical diffusion. However, they have the advantage that they are unconditionally stable.

A.2 Envelope equations for stimulated Raman scattering

A.2.1 Laser evolution

We start from Ampere's law:

$$\nabla \times \vec{B} = \mu_0 \vec{J} + \frac{1}{c^2} \frac{\partial \vec{E}}{\partial t}. \quad (\text{A.11})$$

Substituting the electric and magnetic fields for the reduced vector potential, $\vec{a} = e\vec{A}/mc$, and electric potential, φ , such that $\vec{E} = -\nabla\varphi - (mc/e)\partial\vec{a}/\partial t$, $\vec{B} = (mc/e)\nabla \times \vec{a}$, and rearranging, gives:

$$\left(\frac{\partial^2}{\partial t^2} - c^2 \nabla^2 \right) \vec{a} = \frac{e}{\varepsilon_0 mc} \vec{J} - \frac{e}{mc} \frac{\partial}{\partial t} \nabla \varphi, \quad (\text{A.12})$$

where we have used the vector identity $\nabla \times (\nabla \times \vec{a}) = \nabla(\nabla \cdot \vec{a}) - \nabla^2 \vec{a}$, and chosen the Coulomb gauge, $\nabla \cdot \vec{a} = 0$.

We follow Kruer [29], using conservation of charge and Poisson's equation

to give a relation for the current:

$$\begin{aligned}
\frac{\partial \rho}{\partial t} + \nabla \cdot \vec{J} &= 0, \\
\nabla^2 \varphi &= -\frac{\rho}{\varepsilon_0}, \\
\rightarrow \nabla \cdot \left(\vec{J} - \varepsilon_0 \frac{\partial}{\partial t} \nabla \varphi \right) &= 0.
\end{aligned} \tag{A.13}$$

We split the current into a divergenceless part, \vec{J}_\perp , and a curl-free part, \vec{J}_\parallel .

It follows that:

$$\vec{J}_\parallel = \varepsilon_0 \frac{\partial}{\partial t} \nabla \varphi. \tag{A.14}$$

Substituting into Eq. (A.12) gives:

$$\left(\frac{\partial^2}{\partial t^2} - c^2 \nabla^2 \right) \vec{a} = \frac{e}{\varepsilon_0 m c} \vec{J}_\perp. \tag{A.15}$$

We make the quasi-static approximation, treating the electron response to the laser field as being in the steady state with respect to the instantaneous laser field.¹ Assuming that the ions may be treated as stationary, the current

¹This requires that the motion of electrons parallel to \vec{a} over a laser cycle be small compared to the scale of variation of \vec{a} . Formally, the assumption is $(\vec{v} \cdot \nabla) \vec{a} + \vec{v} \times (\nabla \times \vec{a}) = 0$, which is automatically satisfied in the planar case.

is then:

$$\begin{aligned} J &= -ne\vec{v} \\ &= -ne\left(\vec{u} + \frac{c}{\gamma}\vec{a}\right), \end{aligned} \tag{A.16}$$

where \vec{v} is the total velocity, \vec{u} is the slowly-varying component of the velocity, and $(c/\gamma)\vec{a}$ the fast electron response to the laser field.

The maximum rate at which the slowly-varying component of the electron velocity may vary is given by the plasma frequency. For tenuous plasma, this value is much smaller than the laser frequency, and as such we assume that only currents due to the fast oscillation of electrons in the laser field will contribute coherently to the laser fields. Currents due to the slowly-varying motion of electrons are neglected. Using the fact $\nabla \cdot a = 0$, we may substitute for \vec{J}_\perp in Eq. (A.15):

$$\left(\frac{\partial^2}{\partial t^2} - c^2\nabla^2\right)\vec{a} = -\frac{ne^2}{\varepsilon_0\gamma m}\vec{a}. \tag{A.17}$$

For an infinite plane wave in unperturbed plasma, one arrives at the familiar plasma dispersion relation, $\omega^2 = c^2k^2 + \omega_p^2$, where ω_p is the plasma frequency.

We limit our analysis to the planar case, allowing us to neglect the transverse components of the Laplacian, and assume electron velocities may be

treated as nonrelativistic ($\gamma \approx 1$). As we are interested in the case of Raman backscatter, we write the laser field as a sum of two counterpropagating waves. Implementing an envelope model gives $\vec{a} = (a_0 e^{i\phi_0} + a_1 e^{i\phi_1}) \hat{u}/2 + c.c.$, where a_0, a_1 are the envelope amplitudes of the pump and probe pulses, $\phi_0 = \omega_0 t + k_0 z$, $\phi_1 = \omega_1 t - k_1 z$ are the phases of the associated carrier waves, \hat{u} is the polarisation vector, and *c.c.* denotes the complex-conjugate of the preceding expression. $a_{0,1}$ may be chosen as complex to allow for detuning from $\phi_{0,1}$, satisfying $|a_{0,1}| = (e/mc\omega_{0,1})\sqrt{2I_{0,1}/\varepsilon_0 c}$. $\hat{u} = \hat{u}_x$ for linearly polarised light, $\hat{u} = (\hat{u}_x + i\hat{u}_y)/\sqrt{2}$ for circular polarisation.² The envelope phases are chosen to satisfy $\omega_{0,1}^2 = \omega_p^2 + c^2 k_{0,1}^2$. Writing the plasma density as the perturbation, δn , to the quiescent density, n_0 , the resulting equations are:

$$\begin{aligned} \left(\frac{\partial^2}{\partial t^2} - c^2 \frac{\partial^2}{\partial z^2} + \omega_p^2 \right) a_0 e^{i\phi_0} + \left(\frac{\partial^2}{\partial t^2} - c^2 \frac{\partial^2}{\partial z^2} + \omega_p^2 \right) a_1 e^{i\phi_1} \\ = -\omega_p^2 \frac{\delta n}{n_0} (a_0 e^{i\phi_0} + a_1 e^{i\phi_1}). \end{aligned} \quad (\text{A.18})$$

We seek separate equations for the evolution of pump and probe. The pump may be isolated by expanding the left-hand side, multiplying by $e^{-i\phi_0}$ and neglecting non-resonant components. As we assume slowly-varying en-

²We have here chosen $|a| \sim E_{rms}$, which allows a single set of equations for both linear and circular polarisation. The well-known relation $|a_{0,1}| = eE_{0,1}/mc\omega_{0,1}$ is therefore only valid for the linearly polarised case.

velopes, the periodic term $a_1 e^{i(\phi_1 - \phi_0)}$ on the left-hand side will average to zero over a ponderomotive beatlength, and may therefore be neglected. The periodic term $\delta n a_1 e^{i(\phi_1 - \phi_0)}$ on the right-hand side is, however, retained, as the plasma density perturbation δn is non-uniform on the scale of a beatlength, i.e. the excited plasma wave. Similarly isolating for the probe, the resulting equations are:

$$\begin{aligned} \left(\frac{\partial^2}{\partial t^2} - c^2 \frac{\partial^2}{\partial z^2} \right) a_0 + 2i \left(\omega_0 \frac{\partial}{\partial t} - c^2 k_0 \frac{\partial}{\partial z} \right) a_0 &= -\omega_p^2 \frac{\delta n}{n_0} (a_0 + a_1 e^{i(\phi_1 - \phi_0)}), \\ \left(\frac{\partial^2}{\partial t^2} - c^2 \frac{\partial^2}{\partial z^2} \right) a_1 + 2i \left(\omega_1 \frac{\partial}{\partial t} + c^2 k_1 \frac{\partial}{\partial z} \right) a_1 &= -\omega_p^2 \frac{\delta n}{n_0} (a_0 e^{i(\phi_0 - \phi_1)} + a_1). \end{aligned} \tag{A.19}$$

We neglect the effects of dispersion by removing second derivatives of the envelopes on the left-hand side and the self-currents on the right-hand side, and approximate $c^2 k_0 / \omega_0 \approx c^2 k_1 / \omega_1 = v_g$. We introduce an envelope representation of the longitudinal electrostatic field, E_z , associated with the excited plasma wave: $f e^{i(\phi_1 - \phi_0)} / 2 + c.c. = -e E_z / m c \omega_p$. From Poisson's equation, we have:

$$\begin{aligned} \frac{\delta n}{n_0} &= \frac{c}{\omega_p} \frac{\partial}{\partial z} f e^{i(\phi_1 - \phi_0)} / 2 + c.c. \\ &\approx -i \frac{\omega_0 + \omega_1}{\omega_p} f e^{i(\phi_1 - \phi_0)} / 2 + c.c., \end{aligned} \tag{A.20}$$

where we have assumed the plasma wave envelope varies slowly on the scale of a ponderomotive beatlength. Substituting, and again neglecting nonresonant terms, we obtain:

$$\begin{aligned} \left(\frac{\partial}{\partial t} - v_g \frac{\partial}{\partial z} \right) a_0 &= -\omega_p \frac{\omega_0 + \omega_1}{4\omega_0} f^* a_1, \\ \left(\frac{\partial}{\partial t} + v_g \frac{\partial}{\partial z} \right) a_1 &= \omega_p \frac{\omega_0 + \omega_1}{4\omega_1} f a_0. \end{aligned} \quad (\text{A.21})$$

A.2.2 Plasma Response

For the plasma response, we start with the Lorentz force:

$$\frac{\partial(\gamma\vec{v})}{\partial t} + \vec{v} \cdot \nabla(\gamma\vec{v}) = -\frac{e}{m} \left(\vec{E} + \vec{v} \times \vec{B} \right), \quad (\text{A.22})$$

where we have again assumed stationary ions. Substituting for the reduced vector potential we obtain:

$$\begin{aligned} \frac{\partial(\gamma\vec{v})}{\partial t} + \vec{v} \cdot \nabla(\gamma\vec{v}) &= c \frac{\partial\vec{a}}{\partial t} + \frac{e}{m} \nabla\varphi - c\vec{v} \times (\nabla \times \vec{a}) \\ &= c \frac{\partial\vec{a}}{\partial t} + \frac{e}{m} \nabla\varphi - c(\nabla\vec{a}) \cdot \vec{v} + c(\vec{v} \cdot \nabla) \vec{a}. \end{aligned} \quad (\text{A.23})$$

Using the quasi-static approximation for the electron response to the laser field, we substitute $\gamma\vec{v} = \gamma\vec{u} + c\vec{a}$, to give:

$$\begin{aligned} \frac{\partial(\gamma\vec{u})}{\partial t} + \vec{v} \cdot \nabla(\gamma\vec{u}) &= \frac{e}{m} \nabla\varphi - c(\nabla\vec{a}) \cdot \vec{v}, \\ \frac{\partial(\gamma\vec{u})}{\partial t} + \vec{u} \cdot \nabla(\gamma\vec{u}) + \frac{c}{\gamma} \vec{a} \cdot \nabla(\gamma\vec{u}) &= \frac{e}{m} \nabla\varphi - c(\nabla\vec{a}) \cdot \vec{u} - \frac{c^2}{\gamma} (\nabla\vec{a}) \cdot \vec{a} \\ &= \frac{e}{m} \nabla\varphi - c(\nabla\vec{a}) \cdot \vec{u} - \frac{c^2}{2\gamma} \nabla|\vec{a}|^2. \end{aligned} \quad (\text{A.24})$$

Consistent with the quasi-static approximation, we assume that contributions from $\vec{a} \cdot \nabla(\gamma\vec{u})$, $(\nabla\vec{a}) \cdot \vec{u}$ will sum to zero over a laser cycle. This allows us to

average out oscillatory components to give

$$\frac{\partial(\gamma\vec{u})}{\partial t} + \vec{u} \cdot \nabla(\gamma\vec{u}) = \frac{e}{m} \nabla\varphi - \frac{c^2}{2\gamma} \nabla|\vec{a}|^2. \quad (\text{A.25})$$

The first term on the right-hand side is the electrostatic force, while the second term is the ponderomotive force, which acts to drive charged particles away from regions of high intensity. It is worth noting that for the nonrelativistic case, the force is conservative, i.e. the net change in electron momentum depends only on the initial and final positions, not the path taken. Further, the force depends only on the gradient of the square of the laser amplitude - it does not depend on the laser polarisation.³

Particle number in a fully ionised plasma is conserved, hence:

$$\frac{\partial n}{\partial t} + \nabla \cdot (\vec{u}n) = 0. \quad (\text{A.26})$$

We assume a small perturbation to the equilibrium state, and non-relativistic electron velocities. Substituting $n = n_0 + \delta n$, we then neglect small values (products of \vec{u} , δn) and relativistic effects ($\gamma \approx 1$). Taking the divergence

³One common misunderstanding arises in the context of “linearly polarised” pulses. A linearly polarised laser pulse, such that $\vec{E} = E \cos(\phi)\hat{u}_x$, $\vec{B} = B \cos(\phi)\hat{u}_y$, cannot exert a force in the y direction. The resolution of this apparent contradiction is simple - such a pulse is only possible if it extends infinitely in the y plane, in which case $\partial|a|^2/\partial y = 0$. If $\vec{u}_y \cdot \nabla|\vec{B}| \neq 0$, then $\vec{u}_y \times \vec{B} \neq 0$ (as $\nabla \cdot \vec{B} = 0$), giving rise to a $\vec{v} \times \vec{B}$ force in the y direction.

of Eq. (A.25) and the time derivative of Eq. (A.26), and using Poisson's equation, we obtain the linearised equation for the plasma response:

$$\begin{aligned} \frac{\partial^2 \delta n}{\partial t^2} + n_0 \left(\frac{n_0 e}{m} \nabla^2 \varphi - \frac{c^2}{2\gamma} \nabla^2 |\vec{a}|^2 \right) &= 0, \\ \left(\frac{\partial^2}{\partial t^2} + \omega_p^2 \right) \frac{\delta n}{n_0} &= \frac{c^2}{2} \nabla^2 |\vec{a}|^2. \end{aligned} \quad (\text{A.27})$$

One may reach the same result by considering the plasma electrons act as simple harmonic oscillators [66], driven by the ponderomotive force.

As with the laser evolution, we limit our analysis to the planar case, and substitute for the laser fields, to obtain:

$$\left(\frac{\partial^2}{\partial t^2} + \omega_p^2 \right) \frac{\delta n}{n_0} = \frac{c^2}{4} \frac{\partial^2}{\partial z^2} (|a_0|^2 + |a_1|^2 + 2\Re[a_0^* a_1 e^{i(\phi_1 - \phi_0)}]). \quad (\text{A.28})$$

Here we have assumed that only the difference terms of the laser beat contribute to the plasma response.⁴

We neglect the derivatives of the laser envelopes (which give rise to wake-field terms), as the spatial variation of the laser envelopes is much smaller than that of the laser beat. Again substituting an envelope field for the

⁴For linearly polarised pulses, the phase velocity of the sum terms is $\gg c$, so we assume the resulting plasma response may be neglected, and take the average value over these oscillations for the difference term. For circularly polarised pulses, we have chosen (through our definition of \hat{u}) laser pulses polarised in the opposite sense i.e. one with clockwise polarisation, one counter-clockwise, which will give only difference terms in the laser beat. Laser pulses polarised in the same sense only give sum terms, which typically prevents any significant plasma response.

plasma wave (Eq. (A.20)), we obtain:

$$\left(\frac{\partial^2}{\partial t^2} - 2i(\omega_0 - \omega_1) \frac{\partial}{\partial t} - (\omega_0 - \omega_1)^2 + \omega_p^2 \right) f = -i \frac{\omega_p}{2} (\omega_0 + \omega_1) a_0^* a_1. \quad (\text{A.29})$$

We may factorise Eq. (A.29) to obtain roots corresponding to the Stokes and anti-Stokes excitation:

$$\left(\frac{\partial}{\partial t} - i(\omega_0 - \omega_1 + \omega_p) \right) \left(\frac{\partial}{\partial t} - i(\omega_0 - \omega_1 - \omega_p) \right) f = -i \frac{\omega_p(\omega_0 + \omega_1)}{2} a_0^* a_1. \quad (\text{A.30})$$

Near the Stokes resonance ($\omega_0 - \omega_1 \approx \omega_p$), we expect the anti-Stokes contribution to be negligible. This is equivalent to assuming a slowly-varying envelope for the plasma wave, as we are neglecting the time derivative of the wave evolution. The resulting first order equation is:

$$\left(\frac{\partial}{\partial t} - i(\omega_0 - \omega_1 - \omega_p) \right) f = \frac{\omega_p(\omega_0 + \omega_1)}{2(\omega_0 - \omega_1 + \omega_p)} a_0^* a_1. \quad (\text{A.31})$$

One consequence of reducing the plasma wave evolution to first order is that it prevents energy transfer from probe to pump in the case where the laser bandwidths are sufficiently large that the local probe frequency is higher than the local pump frequency. However, this is to be expected, as the

approximation $\omega_0 - \omega_1 \approx \omega_p$ is no longer valid.

A.3 Wave equation solvers

Several approximations exist to the wave equation. We here derive some of those used in this work.

We start with the wave equation,

$$\left(\frac{\partial^2}{\partial t^2} - c^2 \nabla^2\right) \vec{a} = \vec{j}, \quad (\text{A.32})$$

where \vec{a} and \vec{j} are the reduced vector potential and normalised current. For simplicity we consider only a single pulse. a and j will be used to denote the amplitudes of the associated envelope fields, satisfying $\vec{a} = ae^{i(\omega t - kz)}\hat{u}/2 + c.c.$, $\vec{j} = je^{i(\omega t - kz)}\hat{u}/2 + c.c.$, with \hat{u} the polarisation vector and $c.c.$ denoting the complex conjugate.

A.3.1 Advective solver

In the special case of the first order wave equation,

$$\left(\frac{\partial}{\partial t} - v\frac{\partial}{\partial z}\right)a = -\frac{i}{2\omega}j, \quad (\text{A.33})$$

obtained by expanding Eq. (A.32) and neglecting second order derivatives, we may consider it as a convective derivative,

$$\frac{D}{Dt}a = -\frac{i}{2\omega}j. \quad (\text{A.34})$$

This is an ordinary differential equation, and as such has no stability condition. It is readily solved numerically, by advancing the field one grid cell per timestep, with the timestep satisfying $v\Delta t = \Delta z$:

$$a_m^{n+1} = a_{m\mp 1}^n - \frac{i}{\omega}\Delta t j_m^n. \quad (\text{A.35})$$

The same solution can be achieved by taking first-order Taylor expansions of Eq. (A.33) and setting $v\Delta t = \Delta z$.

A.3.2 Paraxial approximation

The paraxial approximation is commonly used due to its simplicity, e.g. [86].

Taking the wave equation, we may transform into the laser frame, $\zeta = z - v_g t$, $\varsigma = t$, where v_g is the group velocity of the laser pulse, to give:

$$\left(\frac{\partial^2}{\partial \varsigma^2} - 2v_g \frac{\partial^2}{\partial \zeta \partial \varsigma} - (c^2 - v_g^2) \frac{\partial^2}{\partial \zeta^2} \right) \vec{a} = \vec{j}. \quad (\text{A.36})$$

We neglect second derivatives in the laser frame, and approximate $\frac{\partial^2}{\partial \zeta \partial \varsigma} \approx -i \frac{\omega v_g}{c^2} \frac{\partial}{\partial \varsigma}$:

$$2i \frac{\omega v_g^2}{c^2} \left(\frac{\partial}{\partial t} + v_g \frac{\partial}{\partial z} \right) \vec{a} = \vec{j}. \quad (\text{A.37})$$

This equation varies by only a small numerical factor from that obtained by neglecting second derivatives of t and z in the lab frame (e.g. Eq. (A.21), although it should be noted those equations use the plasma value of k for the carrier wave). The limitation of the model is that k and v_g are prescribed, and are not updated as the laser pulse or plasma susceptibility evolve. They are typically chosen to match their initial values in unperturbed plasma, but this will become increasingly inaccurate in nonlinear regimes.

A.3.3 Gordon approximation

Gordon *et al* [90] developed a first-order approximation to the full wave equation which does not rely on a prescribed value for the group velocity or wavenumber of the laser pulse.

Taking Eq. (A.32) and transforming into the light frame, $\xi = z - ct$, $\tau = t$, we obtain:

$$\left(\frac{\partial^2}{\partial \tau^2} - 2c \frac{\partial^2}{\partial \xi \partial \tau} \right) \vec{a} = \vec{j}. \quad (\text{A.38})$$

We neglect only the second derivative in τ , equivalent to assuming $\partial a / \partial \tau \ll c \partial a / \partial \xi$ ⁵, which suppresses backward propagating components [91]. Cowan *et al* [87] and Mora and Antonsen [91] solve this equation directly. We follow Gordon, transforming back into the lab frame, to give:

$$-2c \frac{\partial}{\partial z} \left(\frac{\partial}{\partial t} + c \frac{\partial}{\partial z} \right) \vec{a} = \vec{j}. \quad (\text{A.39})$$

Expanding for envelope fields, we obtain:

$$2i\omega \left(\frac{\partial}{\partial t} + c \frac{\partial}{\partial z} \right) \left(1 + \frac{i}{k} \frac{\partial}{\partial z} \right) a = j. \quad (\text{A.40})$$

⁵It is commonly stated that this is a reasonable approximation as the laser varies slowly in the co-moving frame. This would in fact justify dropping first derivatives. Instead, this approximation assumes that the *rate of change* in the co-moving frame varies slowly.

Multiplying both sides by a factor $1 - (i/k)\partial/\partial z$, and neglect the resulting $(1/k^2)\partial^2/\partial z^2$ term as small:

$$\left(\frac{\partial}{\partial t} + c\frac{\partial}{\partial z}\right)a = -\frac{1}{2\omega}\left(i + \frac{1}{k}\frac{\partial}{\partial z}\right)j. \quad (\text{A.41})$$

This equation should provide a good approximation to the models implemented by Cowan *et al* [87] and Mora and Antonsen [91], which neglect only the second derivative in τ . The Gordon approximation has the advantage that it can be easily implemented in existing first-order solvers, as the left hand side remains first order. The dispersive elements may be included simply by modifying the current term on the right-hand side.

Substituting differentials for their finite-difference expansions (see Appendix A.1) gives:

$$\begin{aligned} \frac{a_m^{n+1} - a_m^{n-1}}{2\Delta t} + c\frac{a_{m+1}^n - a_{m-1}^n}{2\Delta z} \\ = -\frac{i}{2\omega}j_m^n - \frac{c}{2\omega^2}\left(\frac{j_m^{n+1} - j_m^{n-1}}{2\Delta z}\right). \end{aligned} \quad (\text{A.42})$$

Rearranging for unknown values provides the final form:

$$\begin{aligned} a_m^{n+1} = a_m^{n-1} - \frac{c\Delta t}{\Delta z}(a_{m+1}^n - a_{m-1}^n) - \frac{i\Delta t}{\omega}j_m^n \\ - \frac{c\Delta t}{\omega^2}\left(\frac{j_m^{n+1} - j_m^{n-1}}{2\Delta z}\right). \end{aligned} \quad (\text{A.43})$$

A.3.4 Full wave equation solver

A full finite-difference solver has also been implemented for the wave equation. It makes use of no approximations beyond those necessary for Taylor expansions, and is correct to third order.

Expanding Eq. (A.32), the wave equation is:

$$\frac{\partial^2 a}{\partial t^2} - c^2 \frac{\partial^2 a}{\partial z^2} + 2i\omega \left(\frac{\partial a}{\partial t} + c \frac{\partial a}{\partial z} \right) = j. \quad (\text{A.44})$$

We can rewrite this as a discretised relation. Euler solvers are first-order in time, and therefore unsuitable for solving second derivatives, which will always be second order in any numerical scheme derived from Taylor expansions. Crank-Nicholson uses temporal averaging to infer the spatial derivatives at $n + \frac{1}{2}$, making it second order; however, j cannot be interpolated in this way, as j_m^{n+1} is not known, so the resulting equations are not time-centred.

By substituting differentials for their finite-difference expansions from Appendix A.1, we make use of values at $n - 1$ to ensure the second-order equations are time centred. As envelope models make use of a coarser resolution, the increased memory overheads incurred through this treatment are

negligible.

$$\begin{aligned} \frac{a_m^{n+1} - 2a_m^n + a_m^{n-1}}{(\Delta t)^2} - c^2 \frac{a_{m+1}^n - 2a_m^n + a_{m-1}^n}{(\Delta z)^2} \\ + i\omega \left(\frac{a_m^{n+1} - a_m^{n-1}}{\Delta t} + c \frac{a_{m+1}^n - a_{m-1}^n}{\Delta z} \right) = j_m^n. \end{aligned} \quad (\text{A.45})$$

Separating known and unknown values gives:

$$\begin{aligned} \left(\frac{1}{(\Delta t)^2} + i \frac{\omega}{\Delta t} \right) a_m^{n+1} = 2 \left(\frac{1}{(\Delta t)^2} - \frac{c^2}{(\Delta z)^2} \right) a_m^n - \left(\frac{1}{(\Delta t)^2} - i \frac{\omega}{\Delta t} \right) a_m^{n-1} \\ + \left(\frac{c^2}{(\Delta z)^2} - i \frac{c\omega}{\Delta z} \right) a_{m+1}^n + \left(\frac{c^2}{(\Delta z)^2} + i \frac{c\omega}{\Delta z} \right) a_{m-1}^n \\ + j_m^n. \end{aligned} \quad (\text{A.46})$$

Multiplying by $1 - i\omega\Delta t$:

$$\begin{aligned} \left(\frac{1}{(\Delta t)^2} + \omega^2 \right) a_m^{n+1} = 2 \left(\left(\frac{1}{(\Delta t)^2} - \frac{c^2}{(\Delta z)^2} \right) - i \left(\frac{\omega}{\Delta t} - \frac{c^2\omega\Delta t}{(\Delta z)^2} \right) \right) a_m^n \\ - \left(\left(\frac{1}{(\Delta t)^2} + \omega^2 \right) - 2i \frac{\omega}{\Delta t} \right) a_m^{n-1} \\ + \left(\left(\frac{c^2}{(\Delta z)^2} - \frac{\omega^2 c \Delta t}{\Delta z} \right) - i \left(\frac{c^2\omega\Delta t}{(\Delta z)^2} + \frac{c\omega}{\Delta z} \right) \right) a_{m+1}^n \\ + \left(\left(\frac{c^2}{(\Delta z)^2} + \frac{\omega^2 c \Delta t}{\Delta z} \right) - i \left(\frac{c^2\omega\Delta t}{(\Delta z)^2} + \frac{c\omega}{\Delta z} \right) \right) a_{m-1}^n \\ + (1 - i\omega\Delta t) j_m^n. \end{aligned} \quad (\text{A.47})$$

We rewrite the equation in terms of the real and imaginary parts, $a =$

$\Re a + i\Im a, j = \Re j + i\Im j$, to reach the final form:

$$\begin{aligned}
\left(\frac{1}{(\Delta t)^2} + \omega^2\right) \Re a_m^{n+1} = & 2\left(\frac{1}{(\Delta t)^2} - \frac{c^2}{(\Delta z)^2}\right) \Re a_m^n + 2\left(\frac{\omega}{\Delta t} + \frac{c^2\omega\Delta t}{(\Delta z)^2}\right) \Im a_m^n \\
& - \left(\omega^2 + \frac{1}{(\Delta t)^2}\right) \Re a_m^{n-1} + 2\frac{\omega}{\Delta t} \Im a_m^{n-1} \\
& + \left(\frac{c^2}{(\Delta z)^2} - \frac{\omega^2 c\Delta t}{\Delta z}\right) \Re a_{m+1}^n + \left(\frac{c^2\omega\Delta t}{(\Delta z)^2} + \frac{c\omega}{\Delta z}\right) \Im a_{m+1}^n \\
& + \left(\frac{c^2}{(\Delta z)^2} + \frac{\omega^2 c\Delta t}{\Delta z}\right) \Re a_{m-1}^n + \left(\frac{c^2\omega\Delta t}{(\Delta z)^2} + \frac{c\omega}{\Delta z}\right) \Im a_{m-1}^n \\
& + \Re j_m^n + \omega\Delta t \Im j_m^n,
\end{aligned} \tag{A.48}$$

$$\begin{aligned}
\left(\frac{1}{(\Delta t)^2} + \omega^2\right) \Im a_m^{n+1} = & 2\left(\frac{1}{(\Delta t)^2} - \frac{c^2}{(\Delta z)^2}\right) \Im a_m^n - 2\left(\frac{\omega}{\Delta t} - \frac{c^2\omega\Delta t}{(\Delta z)^2}\right) \Re a_m^n \\
& - \left(\omega^2 + \frac{1}{(\Delta t)^2}\right) \Im a_m^{n-1} - 2\frac{\omega}{\Delta t} \Re a_m^{n-1} \\
& + \left(\frac{c^2}{(\Delta z)^2} - \frac{\omega^2 c\Delta t}{\Delta z}\right) \Im a_{m+1}^n - \left(\frac{c^2\omega\Delta t}{(\Delta z)^2} + \frac{c\omega}{\Delta z}\right) \Re a_{m+1}^n \\
& + \left(\frac{c^2}{(\Delta z)^2} + \frac{\omega^2 c\Delta t}{\Delta z}\right) \Im a_{m-1}^n - \left(\frac{c^2\omega\Delta t}{(\Delta z)^2} + \frac{c\omega}{\Delta z}\right) \Re a_{m-1}^n \\
& + \Im j_m^n - \omega\Delta t \Re j_m^n.
\end{aligned} \tag{A.49}$$

Bibliography

- [1] R. G. Hemker, F. S. Tsung, V. K. Decyk, W. B. Mori, S. Lee, and T. Katsouleas, “Development of a parallel code for modeling plasma based accelerators,” in *Proceedings of the Fifth IEEE Particle Accelerator Conference*, Vol. 5 (1999) pp. 3672–3674.
- [2] R. A. Fonseca, L. O. Silva, F. S. Tsung, V. K. Decyk, W. Lu, C. Ren, W. B. Mori, S. Deng, S. Lee, T. Katsouleas, and J. C. Adam, “OSIRIS: A three-dimensional, fully relativistic particle in cell code for modeling plasma based accelerators.” in *International Conference on Computational Science 2002*, Lecture Notes in Computer Science, Vol. 2331, edited by P. M. A. Sloot, C. J. K. Tan, J. Dongarra, and A. G. Hoekstra (Springer, 2002) pp. 342–351.
- [3] B. Ersfeld, J. Farmer, G. Raj, and D. A. Jaroszynski, “Effects of energy absorption on Raman amplification in plasma,” in *Harnessing rela-*

tivisititc plasma waves as novel radiation sources from terahertz to X-rays and beyond, SPIE Proceedings, Vol. 7359 (2009) p. 73590P.

- [4] M. S. Hur, G. Penn, J. S. Wurtele, and R. Lindberg, “Slowly varying envelope kinetic simulations of pulse amplification by Raman backscattering,” *Physics of Plasmas* **11**, 5204 (2004).
- [5] T. H. Maiman, “Stimulated optical radiation in ruby,” *Nature* **187**, 493 (1960).
- [6] C. Raulin, D. Petzoldt, and S. Hellwig, “Removal of wrinkles and acne scars with the ultra-pulse CO2 laser,” *Hautarzt* **47**, 443 (1996).
- [7] L. D. Smullin and G. Fiocco, “Optical echoes from the moon,” *Nature* **194**, 1267 (1962).
- [8] M. H. Anderson, J. R. Ensher, M. R. Matthews, C. E. Wieman, and E. A. Cornell, “Observation of Bose-Einstein condensation in a dilute atomic vapor,” *Science* **269**, 198 (1995).
- [9] S. P. D. Mangles, C. D. Murphy, Z. Najmudin, A. G. R. Thomas, J. L. Collier, A. E. Dangor, E. J. Divall, P. S. Foster, J. G. Gallacher, C. J. Hooker, D. A. Jaroszynski, A. J. Langley, W. B. Mori, P. A. Norreys, F. S. Tsung, R. Viskup, B. R. Walton, and K. Krushelnick, “Monoen-

- ergetic beams of relativistic electrons from intense laser-plasma interactions,” *Nature* **431**, 535 (2004).
- [10] R. Weiss, “Electromagnetically coupled broadband gravitational antenna,” Quarterly Progress Report, Research Laboratory of Electronics, MIT **105**, 54 (1972).
- [11] J. Nuckolls, L. Wood, A. Thiessen, and G. Zimmerman, “Laser compression of matter to super-high densities: Thermonuclear (CTR) applications,” *Nature* **239**, 139 (1972).
- [12] J. Schwinger, “On gauge invariance and vacuum polarization,” *Phys. Rev.* **82**, 664 (1951).
- [13] D. A. Jaroszynski, R. Bingham, and R. A. Cairns, eds., *Laser-Plasma Interactions*, Scottish Graduate Series (CRC Press/Taylor & Francis, 2009).
- [14] D. E. Spence, P. N. Kean, and W. Sibbett, “60-fsec pulse generation from a self-mode-locked Ti:sapphire laser,” *Opt. Lett.* **16**, 42 (1991).
- [15] D. Strickland and G. Mourou, “Compression of amplified chirped optical pulses,” *Optics Communications* **56**, 219 (1985).
- [16] A. Dubietis, G. Jonušauskas, and A. Piskarskas, “Powerful femtosecond

- pulse generation by chirped and stretched pulse parametric amplification in BBO crystal,” *Optics Communications* **88**, 437 (1992).
- [17] G. Shvets, N. J. Fisch, A. Pukhov, and J. Meyer-ter Vehn, “Superradiant amplification of an ultrashort laser pulse in a plasma by a counter-propagating pump,” *Phys. Rev. Lett.* **81**, 4879 (1998).
- [18] I. Blumenfeld, C. E. Clayton, F.-J. Decker, M. J. Hogan, C. Huang, R. Ischebeck, R. Iverson, C. Joshi, T. Katsouleas, N. Kirby, W. Lu, K. A. Marsh, W. B. Mori, P. Muggli, E. Oz, R. H. Siemann, D. Walz, and M. Zhou, “Energy doubling of 42 GeV electrons in a metre-scale plasma wakefield accelerator,” *Nature* **445**, 741 (2007).
- [19] A. H. Compton, “A quantum theory of the scattering of x-rays by light elements,” *Phys. Rev.* **21**, 483 (1923).
- [20] C. V. Raman and K. S. Krishnan, “A new type of secondary radiation,” *Nature* **121**, 501 (1928).
- [21] G. Landsberg and L. Mandelstam, “Eine neue erscheinung bei der lichtzerstreuung in krystallen,” *Naturwissenschaften* **16**, 557 (1928).
- [22] V. C. Farmer, ed., *The Infrared Spectra of Minerals* (Mineralogical Society, London, 1974).

- [23] G. Blumberg, M. V. Klein, and S.-W. Cheong, “Charge and spin dynamics of an ordered stripe phase in $\text{La}_{1\frac{2}{3}}\text{Sr}_{\frac{1}{3}}\text{NiO}_4$ investigated by Raman spectroscopy,” *Phys. Rev. Lett.* **80**, 564 (1998).
- [24] M. G. Cottam and D. J. Lockwood, *Light scattering in magnetic solids* (Wiley, New York, 1986).
- [25] R. H. Stolen and E. P. Ippen, “Raman gain in glass optical waveguides,” *Applied Physics Letters* **22**, 276 (1973).
- [26] A. J. Stentz, “Applications of Raman lasers and amplifiers in fiber communication systems,” in *Nonlinear Optical Engineering*, Proc. SPIE No. 3263, edited by M. S. Mangir (1998) pp. 91–94.
- [27] I. S. Grudinin and L. Maleki, “Ultralow-threshold Raman lasing with CaF_2 resonators,” *Opt. Lett.* **32**, 166 (2007).
- [28] M. E. Fermann, M. Hofer, F. Haberl, and S. P. Craig-Ryan, “Femtosecond fibre laser,” *Electronics Letters* **26**, 1737 (1990).
- [29] W. L. Kruer, *The physics of laser plasma interactions* (Addison-Wesley, California, 1987).
- [30] A. T. Lin and J. M. Dawson, “Stimulated Compton scattering of electromagnetic waves in plasma,” *Physics of Fluids* **18**, 201 (1975).

- [31] W. L. Kruer, S. C. Wilks, B. B. Afeyan, and R. K. Kirkwood, “Energy transfer between crossing laser beams,” *Physics of Plasmas* **3**, 382 (1996).
- [32] S. H. Glenzer, B. J. MacGowan, P. Michel, N. B. Meezan, L. J. Suter, S. N. Dixit, J. L. Kline, G. A. Kyrala, D. K. Bradley, D. A. Callahan, E. L. Dewald, L. Divol, E. Dzenitis, M. J. Edwards, A. V. Hamza, C. A. Haynam, D. E. Hinkel, D. H. Kalantar, J. D. Kilkenny, O. L. Landen, J. D. Lindl, S. LePape, J. D. Moody, A. Nikroo, T. Parham, M. B. Schneider, R. P. J. Town, P. Wegner, K. Widmann, P. Whitman, B. K. F. Young, B. Van Wonterghem, L. J. Atherton, and E. I. Moses, “Symmetric inertial confinement fusion implosions at ultra-high laser energies,” *Science* **327**, 1228 (2010).
- [33] R. E. Turner, D. W. Phillion, E. M. Campbell, and K. Estabrook, “Time-resolved observations of stimulated Raman scattering from laser-produced plasmas,” *Physics of Fluids* **26**, 579 (1983).
- [34] K. Tanaka, L. M. Goldman, W. Seka, M. C. Richardson, J. M. Soures, and E. A. Williams, “Stimulated Raman scattering from uv-laser-produced plasmas,” *Phys. Rev. Lett.* **48**, 1179 (1982).
- [35] D. W. Phillion, D. L. Banner, E. M. Campbell, R. E. Turner, and K. G.

- Estabrook, “Stimulated Raman scattering in large plasmas,” *Physics of Fluids* **25**, 1434 (1982).
- [36] D. W. Forslund, J. M. Kindel, W. B. Mori, C. Joshi, and J. M. Dawson, “Two-dimensional simulations of single-frequency and beat-wave laser-plasma heating,” *Phys. Rev. Lett.* **54**, 558 (1985).
- [37] M. N. Rosenbluth, “Parametric instabilities in inhomogeneous media,” *Phys. Rev. Lett.* **29**, 565 (1972).
- [38] M. R. Amin and R. A. Cairns, “Two-dimensional analysis of beat wave current drive with intense microwave pulses,” *Nuclear Fusion* **30**, 327 (1990).
- [39] R. H. Dicke, “Coherence in spontaneous radiation processes,” *Phys. Rev.* **93**, 99 (1954).
- [40] D. A. Jaroszynski, P. Chaix, N. Piovella, D. Oepts, G. M. H. Knippels, A. F. G. van der Meer, and H. H. Weits, “Superradiance in a short-pulse free-electron-laser oscillator,” *Phys. Rev. Lett.* **78**, 1699 (1997).
- [41] R. Bonifacio, F. Casagrande, G. Cerchiono, L. D. S. Souza, P. Pierini, and N. Piovella, “Physics of the high-gain FEL and superradiance,” *Rivista del nuovo cimento* **13** (1990).

- [42] V. M. Malkin, G. Shvets, and N. J. Fisch, “Fast compression of laser beams to highly overcritical powers,” *Phys. Rev. Lett.* **82**, 4448 (1999).
- [43] B. Ersfeld and D. A. Jaroszynski, “Superradiant linear Raman amplification in plasma using a chirped pump pulse,” *Phys. Rev. Lett.* **95**, 165002 (2005).
- [44] R. L. Berger, D. S. Clark, A. A. Solodov, E. J. Valeo, and N. J. Fisch, “Inverse bremsstrahlung stabilization of noise in the generation of ultrashort intense pulses by backward Raman amplification,” *Physics of Plasmas* **11**, 1931 (2004).
- [45] R. M. G. M. Trines, F. Fiuza, R. Bingham, R. A. Fonseca, L. O. Silva, R. A. Cairns, and P. A. Norreys, “Simulations of efficient Raman amplification into the multipetawatt regime,” *Nat Phys* **7**, 87 (2011).
- [46] W. Cheng, Y. Avitzour, Y. Ping, S. Suckewer, N. J. Fisch, M. S. Hur, and J. S. Wurtele, “Reaching the nonlinear regime of Raman amplification of ultrashort laser pulses,” *Phys. Rev. Lett.* **94**, 045003 (2005).
- [47] J. Ren, W. Cheng, S. Li, and S. Suckewer, “A new method for generating ultraintense and ultrashort laser pulses,” *Nat Phys* **3**, 732 (2007).
- [48] “Alpha-X project,” <http://phys.strath.ac.uk/alpha-x/> (2002-present).

- [49] D. A. Jaroszynski, R. Bingham, E. Brunetti, B. Ersfeld, J. Gallacher, B. van der Geer, R. Issac, S. P. Jamison, D. Jones, M. de Loos, A. Lyachev, V. Pavlov, A. Reitsma, Y. Saveliev, G. Vieux, and S. M. Wiggins, “Radiation sources based on laser-plasma interactions,” *Philos. Trans. R. Soc. A-Math. Phys. Eng. Sci.* **364**, 689 (2006).
- [50] S. Cipiccia, M. R. Islam, B. Ersfeld, R. P. Shanks, E. Brunetti, G. Vieux, X. Yang, R. C. Issac, S. M. Wiggins, G. H. Welsh, M.-P. Anania, D. Maneski, R. Montgomery, G. Smith, M. Hoek, D. J. Hamilton, N. R. C. Lemos, D. Symes, P. P. Rajeev, V. O. Shea, J. M. Dias, and D. A. Jaroszynski, “Gamma-rays from harmonically resonant betatron oscillations in a plasma wake,” *Nat Phys* **7**, 867 (2011).
- [51] G. Vieux, A. Lyachev, X. Yang, B. Ersfeld, J. P. Farmer, E. Brunetti, R. C. Issac, G. Raj, G. H. Welsh, S. M. Wiggins, and D. A. Jaroszynski, “Chirped pulse Raman amplification in plasma,” *New Journal of Physics* **13**, 063042 (2011).
- [52] J. P. Farmer, B. Ersfeld, and D. A. Jaroszynski, “Raman amplification in plasma: Wavebreaking and heating effects,” *Physics of Plasmas* **17**, 113301 (2010).
- [53] D. J. Spence, A. Butler, and S. M. Hooker, “First demonstration of

- guiding of high-intensity laser pulses in a hydrogen-filled capillary discharge waveguide,” *Journal of Physics B: Atomic, Molecular and Optical Physics* **34**, 4103 (2001).
- [54] C. J. Woolley, K. O’Keeffe, H. K. Chung, and S. M. Hooker, “Time-resolved plasma temperature measurements in a laser-triggered hydrogen-filled capillary discharge waveguide,” *Plasma Sources Science and Technology* **20**, 055014 (2011).
- [55] J. Vieira, F. Fiuza, R. Fonseca, L. Silva, C. Huang, W. Lu, M. Tzoufras, F. Tsung, V. Decyk, W. Mori, J. Cooley, and T. Antonsen, “One-to-one full-scale simulations of laser-wakefield acceleration using QuickPIC,” *Plasma Science, IEEE Transactions on* **36**, 1722 (2008).
- [56] E. Cormier-Michel, B. A. Shadwick, C. G. R. Geddes, E. Esarey, C. B. Schroeder, and W. P. Leemans, “Unphysical kinetic effects in particle-in-cell modeling of laser wakefield accelerators,” *Phys. Rev. E* **78**, 016404 (2008).
- [57] M. S. Hur and J. S. Wurtele, “Two-dimensional simulations of the amplification and focusing of intense laser pulses in the kinetic regime of Raman backward amplification in plasmas,” *Computer Physics Communications* **180**, 651 (2009).

- [58] D. L. Bobroff and H. A. Haus, “Impulse response of active coupled wave systems,” *Journal of Applied Physics* **38**, 390 (1967).
- [59] G. L. Lamb, “Analytical descriptions of ultrashort optical pulse propagation in a resonant medium,” *Rev. Mod. Phys.* **43**, 99 (1971).
- [60] A. B. Langdon and C. K. Birdsall, *Plasma physics via computer simulation* (McGraw-Hill, New York, New York :, 1985) pp. –.
- [61] K. Nanbu, “Theory of cumulative small-angle collisions in plasmas,” *Phys. Rev. E* **55**, 4642 (1997).
- [62] F. Peano, M. Marti, L. O. Silva, and G. Coppa, “Statistical kinetic treatment of relativistic binary collisions,” *Phys. Rev. E* **79**, 025701 (2009).
- [63] K. Paul, D. L. Bruhwiler, B. Cowan, J. R. Cary, C. Huang, F. S. Tsung, W. B. Mori, E. Cormier-Michel, C. G. R. Geddes, E. Esarey, S. Martins, R. A. Fonseca, and L. O. Silva, “Benchmarking the codes VORPAL, OSIRIS and QuickPIC with laser wakefield acceleration simulations,” in *Proc. Advanced. Accel. Concepts Workshop*, Vol. LBNL-2293E (2008) pp. 315–320.
- [64] K. Yee, “Numerical solution of initial boundary value problems involving Maxwell’s equations in isotropic media,” *Antennas and Propagation*,

- IEEE Transactions on, Antennas and Propagation*, IEEE Transactions on DOI - 10.1109/TAP.1966.1138693 **14**, 302 (1966).
- [65] R. Courant, K. Friedrichs, and H. Lewy, “On the partial difference equations of mathematical physics,” *IBM Journal of Research and Development* **11**, 215 (1967).
- [66] J. M. Dawson, “Nonlinear electron oscillations in a cold plasma,” *Phys. Rev.* **113**, 383 (1959).
- [67] R. Betti, C. D. Zhou, K. S. Anderson, L. J. Perkins, W. Theobald, and A. A. Solodov, “Shock ignition of thermonuclear fuel with high areal density,” *Phys. Rev. Lett.* **98**, 155001 (2007).
- [68] V. M. Malkin, G. Shvets, and N. J. Fisch, “Detuned Raman amplification of short laser pulses in plasma,” *Phys. Rev. Lett.* **84**, 1208 (2000).
- [69] C.-S. Lui, “Parametric instabilities in laser plasma interaction,” in *From Particles to Plasmas*, edited by J. Van Dam (Addison-Wesley, California, 1989) Chap. 9, p. 141.
- [70] T.-L. Wang, D. Michta, R. R. Lindberg, A. E. Charman, S. F. Martins, and J. S. Wurtele, “Feasibility study for using an extended three-wave model to simulate plasma-based backward Raman amplification in one spatial dimension,” *Physics of Plasmas* **16**, 123110 (2009).

- [71] A. G. Khachatryan and S. S. Elbakian, “An explicit solution for non-linear plasma waves of arbitrary amplitude,” arXiv , physics/9809011v1 (1998).
- [72] Y. Shima and H. Yatom, “Inverse bremsstrahlung energy absorption rate,” Phys. Rev. A **12**, 2106 (1975).
- [73] L. D. Landau, “On the vibration of the electronic plasma,” J. Phys. USSR **10**, 26 (L D 1946).
- [74] D. Bohm and E. P. Gross, “Theory of plasma oscillations. A. Origin of medium-like behavior,” Phys. Rev. **75**, 1851 (1949).
- [75] T. H. Stix, *The Theory of Plasma Waves* (McGraw-Hill, New York, 1962).
- [76] T. P. Coffey, “Breaking of large amplitude plasma oscillations,” Physics of Fluids **14**, 1402 (1971).
- [77] R. R. Lindberg, A. E. Charman, and J. S. Wurtele, “Reduced kinetic description of weakly-driven plasma waves,” Physics of Plasmas **15**, 055911 (2008).
- [78] A. A. Balakin, G. M. Fraiman, N. J. Fisch, and V. M. Malkin, “Noise suppression and enhanced focusability in plasma Raman amplifier with multi-frequency pump,” Physics of Plasmas **10**, 4856 (2003).

- [79] V. M. Malkin, Y. A. Tsidulko, and N. J. Fisch, “Stimulated Raman scattering of rapidly amplified short laser pulses,” *Phys. Rev. Lett.* **85**, 4068 (2000).
- [80] B. Ersfeld, J. Farmer, G. Raj, and D. A. Jaroszynski, “The role of absorption in Raman amplification in warm plasma,” *Physics of Plasmas* **17**, 083301 (2010).
- [81] J. M. Dawson, “On Landau damping,” *Physics of Fluids* **4**, 869 (1961).
- [82] T. O’Neil, “Collisionless damping of nonlinear plasma oscillations,” *Physics of Fluids* **8**, 2255 (1965).
- [83] N. A. Yampolsky and N. J. Fisch, “Effect of nonlinear Landau damping in plasma-based backward Raman amplifier,” *Physics of Plasmas* **16**, 072105 (2009).
- [84] R. Bilato and M. Brambilla, “On the nature of “collisionless” Landau damping,” *Communications in Nonlinear Science and Numerical Simulation* **13**, 18 (2008).
- [85] D. C. Burnham and R. Y. Chiao, “Coherent resonance fluorescence excited by short light pulses,” *Phys. Rev.* **188**, 667 (1969).
- [86] E. Esarey, P. Sprangle, J. Krall, and A. Ting, “Overview of plasma-

- based accelerator concepts,” Plasma Science, IEEE Transactions on **24**, 252 (1996).
- [87] B. M. Cowan, D. L. Bruhwiler, E. Cormier-Michel, E. Esarey, C. G. R. Geddes, P. Messmer, and K. M. Paul, “Characteristics of an envelope model for laser-plasma accelerator simulation,” J. Comput. Phys. **230**, 61 (2011).
- [88] E. Esarey, R. F. Hubbard, W. P. Leemans, A. Ting, and P. Sprangle, “Electron injection into plasma wakefields by colliding laser pulses,” Phys. Rev. Lett. **79**, 2682 (1997).
- [89] T. Tajima and J. M. Dawson, “Laser electron accelerator,” Phys. Rev. Lett. **43**, 267 (1979).
- [90] D. F. Gordon, W. B. Mori, and J. Antonsen, T. M., “A ponderomotive guiding center particle-in-cell code for efficient modeling of laser-plasma interactions,” Plasma Science, IEEE Transactions on **28**, 1135 (2000).
- [91] P. Mora and J. T. M. Antonsen, “Kinetic modeling of intense, short laser pulses propagating in tenuous plasmas,” Physics of Plasmas **4**, 217 (1997).
- [92] L. Brieda, “Particle push in magnetic field (Boris

- method),” in *Scientific Computing and Plasma Physics Blog*,
<http://www.particleincell.com/2011/vxb-rotation/> (2011).
- [93] O. Yaakobi, L. Friedland, R. R. Lindberg, A. E. Charman, G. Penn, and J. S. Wurtele, “Spatially autoresonant stimulated Raman scattering in nonuniform plasmas,” *Physics of Plasmas* **15**, 032105 (2008).
- [94] G. Vieux, B. Ersfeld, J. P. Farmer, M. S. Hur, R. C. Issac, and D. A. Jaroszynski, “Plasma density measurements using chirped pulse broadband raman amplification,” *Applied Physics Letters* **103**, 121106 (2013).
- [95] J. Crank and P. Nicolson, “A practical method for numerical evaluation of solutions of partial differential equations of the heat-conduction type,” *Proc. Camb. Phil. Soc.* **43**, 50 (1947).