Optimisation and Applications of Raman Plasma Amplifiers



James D. Sadler
Wolfson College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy

Michaelmas 2017

There is no more open door by which you can enter into the study of natural philosophy than by considering the physical phenomena of a candle.

Michael Faraday

We forget just how painfully dim the world was before electricity. A candle, a good candle, provides barely a hundredth of the illumination of a single 100 watt light bulb.

Bill Bryson

I do not think that the radio waves I have discovered will have any practical application.

Heinrich Hertz

In nature, nothing is perfect and everything is perfect. Trees can be contorted, bent in weird ways, and they're still beautiful.

Alice Walker

Acknowledgements

Becoming a scientist is a pursuit that requires effort, creativity, courage, networking, long term thinking and resilience in the face of adversity. I would like the thank the important people in my life for developing these qualities in me.

To my family, who gave their all and convinced me to pursue my passion as a career.
To my research group, who persistently paid attention to this work, gave up their time and showed me when I was wrong.
To Mr Howe, my physics school teacher, who chose to spend his final time of a terminal illness inspiring our class of future scientists.
To my supervisors Peter and Raoul, a pair of knowledgeable and forgiving scientists, who always had time to share their ideas and wisdom as my friends. Thank you for always being on hand to reap the rewards of good outcomes and show me the silver linings in bad outcomes.
To those at the Central Laser Facility and the Laboratory for Laser Energetics in the U.S.A., who showed me there was a world of science out there.
To my friends Charlie, Alex, Tom, Chris, David, Tom, Kate, Sam, Duncan, Aarne, Caroline and Clarke. Also to all of my office buddies, especially Tom and Chris for enough lunchtime chat to last a lifetime. A night out with you lot was never far away.
Finally to my girlfriend Adrianna. Thank you for inspiring me to write the second best product of my degree, as meeting you was the first.

Abstract

Laser pulses have the shortest duration, smallest spatial extent, highest power, longest range and highest electric field of any known physical probe. In just 5 decades, the laser has led to the detection of gravitational waves, quantum optics and computers, worldwide fibre-optic communications, innumerable manufacturing applications and a prime contender for replicating the sun's nuclear fusion on earth. Modern high power lasers achieve intensities exceeding $10^{21}\,\mathrm{Wcm}^{-2}$.

Laser powers have increased without bound due to the method of chirped pulse amplification. However, continuing this trend may require new amplifier materials with higher damage thresholds. This work assesses the feasibility of laser amplification in a plasma. Energy is transferred from a long, low power pulse to a short, high power pulse. This is due to the coupling of the light with plasma electron waves, the so called Raman instability. The resulting pulses could ignite a nuclear fusion capsule under the fast ignition scheme, or generate electron positron plasmas.

The energy efficiency and coherence of the amplifier are greatly influenced by the plasma conditions, investigated here using large scale computer simulations. The existing analytical solution is then generalised to three dimensions. The simulations show effective amplification, even in non-uniform plasmas.

A method for seeding the amplifier was investigated experimentally at the Central Laser Facility. A second experiment at the Laboratory for Laser Energetics in New York, U.S.A., found features of the Raman light that show promise for testing the full scheme in 2019.

Contents

List of Publications **iv**

Mathematical Conventions and Definition of the Symbols **viii**

1 Introduction and Theory of Plasmas **1**

1.1 Introduction 1

1.2 High-power laser systems 3

1.3 Ionisation 8

1.4 Neutral gases 10

1.5 Plasmas 12

1.5.1 The Debye length 13

1.5.2 Particle collisions 16

1.5.3 Stopping power 18

1.6 Plasma evolution at large scales 20

1.6.1 The Fluid equations 21

1.6.2 Electron plasma waves 23

1.6.3 Landau damping of electron plasma waves 25

1.7 Particle motion in electromagnetic waves 27

1.8 Propagation of light in a plasma 30

1.8.1 The ponderomotive force in a finite laser pulse 31

1.8.2 The filamentation instability 32

2 Theory of the Raman Instability and Raman Amplifiers **34**

2.1 The Raman instability 34

2.1.1 The matching conditions 35

2.1.2 Derivation of the coupled envelope equations 37

2.1.3 Linear growth rates 41

2.2 The non-linear stage 43

2.2.1 The self similar solution 44
2.2.2 Wave-breaking 48
2.3 Past experimental progress 49

3 Details of Computer Simulations and Plasma Diagnostics 51
3.1 Particle-in-cell computer simulations 51
3.2 Particle-in-cell (PIC) algorithm 54
3.2.1 Interpolation to the grid 54
3.2.2 Field solve 54
3.2.3 Particle push 55
3.2.4 Moving simulation domain 57
3.3 Experimental interferometry diagnostic 58
3.4 Experimental FROG diagnostic 59

4 Optimisation of the Amplifier 62
4.1 Efficiency of the amplifier 63
4.2 Optimisation of the plasma parameters 66
4.2.1 Plasma constraints 66
4.2.2 Ionisation schemes 67
4.2.3 Dynamics of the amplified pulse 69
4.3 Numerical integration of the envelope equations 70
4.4 Particle-in-cell simulation 72
4.5 Parameter scan 76
4.5.1 Discussion of errors 81
4.6 Optimal seed pulse 83
4.6.1 Moving window simulations 83
4.6.2 Fixed window simulations 85
4.7 Effect of competing instabilities 88
4.7.1 Maximum compression ratio and amplification factor 90
4.8 Experimental design 95
4.9 Two-dimensional simulation 96
4.10 Compression of longer pulses 99

5 Solution for Diverging Beams **103**
5.1 Generalisation of the coupled envelope equations 105
5.2 Steady state amplification 110
5.3 Simulation of a diverging beam 113
5.4 Effect of competing instabilities 116

6 Effects of Perturbations From Ideal Conditions **118**
6.1 Required plasma uniformity 118
6.2 Effect of density shifts 120
6.3 Effect of density gradients 122
6.4 Seeding of competing instabilities 124

7 Seeding the Amplifier via Raman Forward Scattering **127**
7.1 Seeding scheme 127
7.2 The Astra laser and set-up 128
7.3 Interferometry results 131
7.4 Measurement of the Raman scatter 133
7.4.1 Energy and profile of the forward scatter 134

8 Experimental Studies of Raman Backscatter **138**
8.1 Experimental objectives 138
8.2 The MTW laser facility 140
8.3 OPAL beam-line and the Raman amplifier chamber 141
8.4 Measurement of the Raman scatter 143
8.5 Particle-in-cell simulation of the full plasma 146

9 Application to X-ray Pulses **150**
9.1 Scaling to shorter wavelengths 150
9.2 Simulation results 154

10 Conclusions and Outlook **159**
10.1 Overall conclusions 159
10.2 Outlook 161

Bibliography **162**

Publications and Communications

[1] Luke Ceurvorst, Naren Ratan, MC Levy, MF Kasim, James Sadler, RHH Scott, RMGM Trines, TW Huang, Marija Skramic, Marija Vranic, et al. "Mitigating the hosing instability in relativistic laser-plasma interactions". In: *New Journal of Physics* 18.5 (2016), p. 053023.

[2] Muhammad Firmansyah Kasim, Naren Ratan, Luke Ceurvorst, James Sadler, Philip N Burrows, Raoul Trines, James Holloway, Matthew Wing, Robert Bingham, and Peter Norreys. "Simulation of density measurements in plasma wakefields using photon acceleration". In: *Physical Review Special Topics-Accelerators and Beams* 18.3 (2015), p. 032801.

[3] Muhammad Firmansyah Kasim, Luke Ceurvorst, Naren Ratan, James Sadler, Nicholas Chen, Alexander Sävert, Raoul Trines, Robert Bingham, Philip N Burrows, Malte C Kaluza, et al. "Quantitative shadowgraphy and proton radiography for large intensity modulations". In: *Physical Review E* 95.2 (2017), p. 023306.

[4] Muhammad Firmansyah Kasim, James Holloway, Luke Ceurvorst, Matthew C Levy, Naren Ratan, James Sadler, Robert Bingham, Philip N Burrows, Raoul Trines, Matthew Wing, et al. "Quantitative single shot and spatially resolved plasma wakefield diagnostics". In: *Physical Review Special Topics-Accelerators and Beams* 18.8 (2015), p. 081302.

[5] N Ratan, NJ Sircombe, L Ceurvorst, J Sadler, MF Kasim, J Holloway, MC Levy, R Trines, R Bingham, and PA Norreys. "Dense plasma heating by crossing relativistic electron beams". In: *Physical Review E* 95.1 (2017), p. 013211.

[6] James D Sadler, Ricky Nathvani, Piotr Oleśkiewicz, Luke A Ceurvorst, Naren Ratan, Muhammad F Kasim, Raoul MGM Trines, Robert Bingham, and Peter A Norreys. "Compression of X-ray free electron laser pulses to attosecond duration". In: *Scientific reports* 5 (2015).

[7] James D Sadler, Raoul MGM Trines, Max Tabak, Dan Haberberger, Dustin H Froula, Andrew S Davies, Sara Bucht, Luís O Silva, E Paulo Alves, Frederico Fiúza, et al. "Optimization of plasma amplifiers". In: *Physical Review E* 95.5 (2017), p. 053211.

[8] James D Sadler, Marcin Sliwa, Thomas Miller, Muhammad F Kasim, Naren Ratan, Luke Ceurvorst, Alex Savin, Ramy Aboushelbaya, Peter A Norreys, Dan Haberberger, et al. "Robustness of raman plasma amplifiers and their potential for attosecond pulse generation". In: *High Energy Density Physics* 23 (2017), pp. 212–216.

[9] JD Sadler, LO Silva, RA Fonseca, K Glize, MF Kasim, A Savin, R Aboushelbaya, MW Mayr, B Spiers, RHW Wang, R Bingham, RMGM Trines, and PA Norreys. "Advantages to a Diverging Raman Amplifier". In: *Communications Physics* (Submitted 2018).

Conference Presentations

- Contributed talk, Central laser facility Christmas meeting, Abingdon, UK (2014, 2015, 2017)
- Poster, IOP plasma physics conference, Oxford, UK, (2017)
- Invited talk, anomalous absorption conference, Old Saybrook, USA, (2016)
- Poster, radiative properties of hot dense matter conference, Santa Barbara, USA, (2016)
- Poster, Omega laser user group meeting, Rochester, USA, (2016, 2017)
- Poster, National ignition facility user group meeting, USA, (2016, 2017)
- Contributed talk, APS plasma physics conference, Savannah, USA, (2015)
- Contributed talk, ignition sessions workshop, Oxford, UK (2015)

Contribution of the Author

All text in this thesis has been written by myself. The background theory and derivations in Chapters 1 to 3 have been reproduced, with

some modification, from the given references. All content from Chapter 4 onwards is original work by myself. However, there was some experimental support from others, as well as advice and helpful theoretical discussions.

All analytical work from chapter 4 onwards was devised and explored by myself only. All computer simulations were conducted and analysed by myself, using a widely available code that was developed and maintained by a separate collaboration. All figures in the thesis were created by myself. The idea for using Raman forward scatter to seed the amplifier was put forward by Peter Norreys. The idea for a diverging beam geometry was mine.

The experiment in Chapter 7 was planned and conducted by myself, the PI Peter Norreys, and the rest of our experimental group. The data analysis and all subsequent work was performed by myself.

The experiment in Chapter 8 was planned and conducted by the PI Dustin Froula and his research group at LLE. I also participated in the laboratory experimental work for 10 weeks during the commissioning and data collection period. I then analysed the resulting data myself.

I have published the work in Chapter 4 in reference [7], the work in Chapter 6 in reference [8] and the work in Chapter 9 in reference [6]. The work in Chapter 5 has been submitted for publication in reference [9] and publication of the work in Chapter 7 is in preparation. The experiment in Chapter 8 is unpublished.

Mathematical Conventions and Definition of the Symbols

Vector quantities will be given in bold. This work will use the Einstein summation convention, where any contracted indices are summed over the spatial or velocity space dimensions. Partial derivatives will be contracted to read $\left(\frac{\partial f}{\partial t}\right)_x = \partial_t f$. Linearisation of the equations will introduce small perturbation quantities denoted by tildes. Properties of the pump, seed and plasma waves will be denoted by the subscripts $0, 1$ and 2 respectively.

c	The speed of light	$\hbar$	Reduced Planck constant
e	Euler's number, fundamental charge	Λ	Number of e-foldings
		f	Particle distribution
n_i	Ion number density	n_{crit}	Pump critical density
m	Electron mass	m_i	Ion mass
Z	Ion charge	$\lambda_{0,1}$	Pump/seed wavelength
$k_{0,1}$	Pump/seed wave-number	$\omega_{0,1}$	Pump/seed frequency
$\Delta\omega_{0,1}$	Pump/seed bandwidth	$T_{0,1}$	Pump/seed duration
T	Temperature	k_B	Boltzmann constant
p	Pressure	$\mathbf{p}$	Momentum
$\mathbf{v}$	Velocity, group velocity	$\mathbf{u}$	Fluid velocity
b	Impact parameter	$\mathbf{A}$	Vector potential
$a_{0,1}$	Dimensionless amplitude	ϕ	Scalar potential, phase
ρ	Charge density	$\mathbf{J}$	Current density
γ	Lorentz factor	n	Free electron number density
ν	Adiabatic index		
F	Focal number, EM tensor	I	Intensity
$P_{0,1}$	Pump/seed power	η	Total efficiency
η_1	Efficiency to leading spike	δ	Polarisation
w	Waist of a Gaussian beam	t_f	Time of best focus
L	Amplifier length	L_R	Rayleigh length
λ_D	Debye length	ω_p	Electron plasma frequency
τ_{ei}	Electron ion collision time	v_p	Phase velocity
$\mathbf{E}$	Electric field	$\mathbf{B}$	Magnetic field
γ_{RBS}	Raman growth rate	γ_{coll}	Collisional damping rate
v_{th}	Thermal electron velocity	L_{max}	Maximum amplifier length
β	Self similar coordinate	$\Phi_{0,1}$	Pump/seed fluence

Chapter 1

Introduction and Theory of Plasmas

1.1 Introduction

Modern high-power laser systems, such as the Vulcan petawatt laser at the Central Laser Facility in the U.K., can be focussed to intensities exceeding $10^{21}\,\mathrm{Wcm}^{-2}$. The electric field of the light is strong enough that atoms in the target are ionised within several laser wavelengths and the resulting free electrons oscillate at speeds close to c. The interplay of the short electromagnetic pulse with collective long range plasma effects leads to a variety of extreme phenomena. These include absorption, scattering and modulation of the light, nuclear reactions, bulk acceleration of the target, generation of huge magnetic fields, electron and proton acceleration, production of gamma rays and production of electron positron pairs.

The unique laboratory environment is suitable for studying processes that otherwise occur only in stars and outer space [83, 48, 7]. Laser pulses can be used both as a driver for creating these conditions, and as a probe to diagnose physical conditions and processes.

There are predictions that even higher intensities of the order $10^{23-24}\,\mathrm{Wcm}^{-2}$ will induce quantum electrodynamics corrections to the interaction physics. Measurements of this could include vacuum phase shifts due to polarisation of the virtual electron positron pairs [54, 55], production of real electron positron pairs and efficient conversion of the laser light to gamma rays [100]. Even far below the critical Schwinger field, the electromagnetic field becomes nonlinear and photon-photon scattering is possible [65, 6, 44]. It may be practically infeasible for light at this intensity to propagate, since any electrons in the beam would initiate a

cascade of pair production and quickly convert the laser light into a macroscopic matter-antimatter plasma and gamma photons.

Measurement of these effects could be possible by colliding laser pulses with a relativistic electron beam [5]. However, improvements to currently available laser pulse powers may be required.

In addition, a number of high power laser applications have been explored, including inertial confinement fusion [63], acceleration of electrons and protons [36, 141], neutron beam production [103] and accelerating small spacecraft to relativistic speeds [80].

Firstly, a hydrogen capsule compressed by an intense radiation drive is capable of sustaining a self heating fusion burn, a clean, compact and abundant form of energy that could sharply cut the cost of electricity [19]. It was predicted that a burn wave can be ignited by a short, high power laser pulse [116]. Due to energy losses in the surrounding coronal plasma, greater pulse powers in the ultra-violet range will be required.

Second, the laser wake-field accelerator concept [78, 82] promises to reduce the size and cost of particle accelerators by 3 orders of magnitude. This is due to the large electric fields that can be sustained in a plasma wave. Many experiments have reached high quality electron bunches with GeV energies. Due to complications with staggering multiple accelerating stages, it is likely that reaching the TeV level will require a single intense laser pulse, beyond current capabilities.

Finally, a theoretical study has indicated that ultra-intense laser plasma interactions will be a powerful gamma ray source [89]. This could have many implications in nuclear physics and industrial sensing applications.

Worldwide, a large number of systems have now reached a power of 1 PW. This includes the Vulcan system (600 J in 600 fs), the National Ignition Facility (1.8 MJ in 2 ns) and the Astra Gemini facility (30 J in 30 fs).

Despite innovative proposals to upgrade the Vulcan laser to 20 PW [47], in recent years conventional laser systems have met an inherent limitation. Damage to solid state optics is often sustained when the incident fluence exceeds approximately $1\,\mathrm{Jcm}^{-2}$, or the intensity is high enough to ionise the surface, typically greater than $10^{12}\,\mathrm{Wcm}^{-2}$. Petawatt systems therefore require metre scale optics, a significant and expensive engineering challenge. Furthermore, the damage limitations become more restrictive for ultra-violet and shorter wavelengths because of photo-ionisation.

Compact laser systems hoping to exceed 1 PW must therefore use plasma based optics.

Laser transitions could be possible in a partially ionised medium, however a simpler method is to start with a long, low power pulse and transfer its energy to a shorter pulse in a plasma. This is known as pulse compression. The most promising mechanism is that of stimulated Raman scattering [85], where two electromagnetic waves can couple via a plasma wave and exchange energy. Predictions indicate that a high efficiency non-linear stage exists [73, 112]. This work will assess the limitations of the scheme and assess the performance across all main plasma and laser parameters.

The simulation results show that there is a parameter window where pre-focussed intensities of $10^{17}\,\mathrm{Wcm}^{-2}$ can be reached. This is a full 5 orders of magnitude greater than conventional laser systems. Furthermore, I show that the scheme is robust to variations in the plasma conditions and beam geometry. Over the last decade, several attempts were made to reach this non-linear stage of Raman amplification and some experiments have diagnosed characteristic features of it [98]. However, none have achieved the full potential of the scheme because there are a wide variety of physical processes that affect the performance and a large number of free parameters. Experimental design therefore requires an extensive body of realistic computer simulations, of which this work is a part. Furthermore, the simulations presented here can predict the performance of some more exotic schemes such as varying the beam geometry, ionisation method and plasma density gradients.

We have also tested one of these schemes experimentally at the UK's Central Laser Facility. The results show it is possible to easily seed the amplifier through a two stage plasma interaction, therefore requiring only a single initial laser pulse. This represents a significant simplification.

1.2 High-power laser systems

Many materials exhibit an atomic energy structure that allows laser operation through a cycle of four levels. The first is the ground state which is heavily populated. External pumping can excite stimulated absorption, promoting electrons to a higher energy level. The electrons then quickly decay to a third level, the upper laser level. This should have a long spontaneous decay lifetime, to easily build a

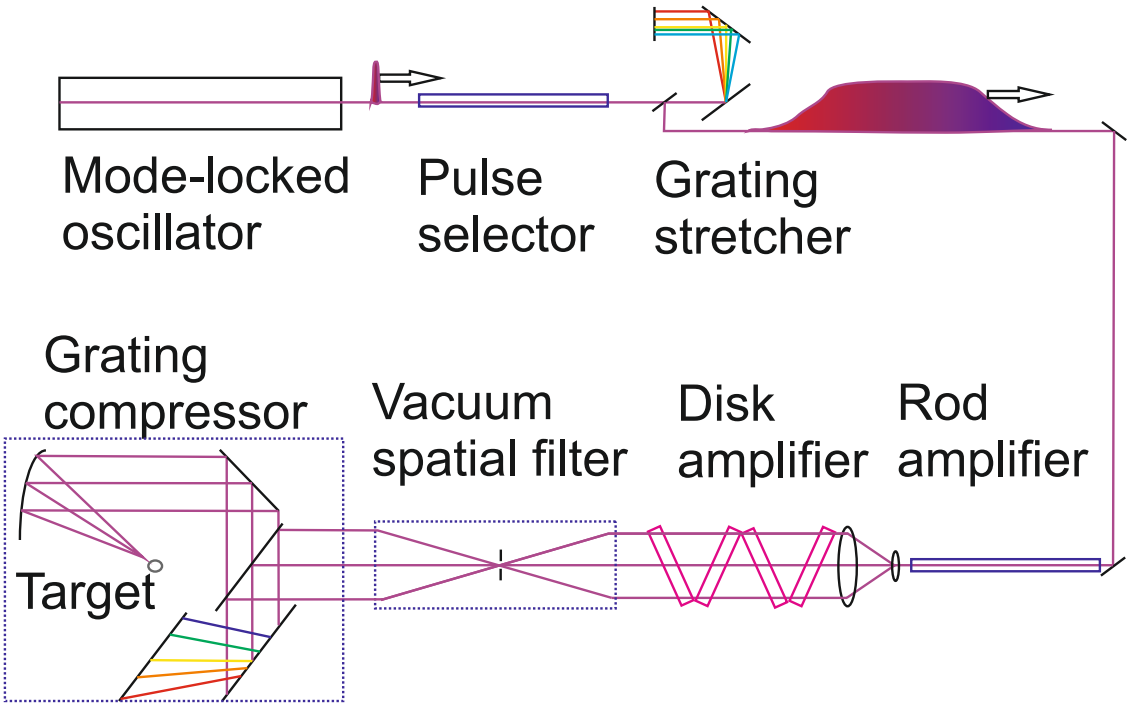


Figure 1.1: Schematic of the main parts of a high-power laser system.

population inversion. This means there are more atoms in the excited upper laser level than the lower laser level. The transition to the lower laser level is typically in the visible or near-infra red for many high-power laser systems.

Once the laser medium has a large population inversion, light can be emitted using some seed radiation. Incoming laser light induces a stimulated decay of the excited atom, releasing a photon with the same direction and polarisation as the incoming beam.

When propagating through a laser amplifier, the light intensity increases and becomes high enough to induce a change in the refractive index. This causes the plane wave to focus and damage the optics, placing an intensity limitation on the amplifier. Since it is expensive and impractical to make the beam much wider, it is preferable to stretch the beam in duration to reduce its intensity. This method is known as chirped pulse amplification (CPA) and has massively increased peak laser powers.

CPA uses a high quality, coherent short pulse seed for the amplifier. This is shown in the schematic in Fig. 1.1. It typically starts at sub-picosecond duration. To reduce the intensity sustained by the amplifiers, a system of reflective gratings induces a wavelength dependent path difference and stretches the short pulse to

several nanoseconds. This gives the pulse a linear gradient in wavelength, known as a chirp. The pulse is often split into several beams to allow more versatile experiments.

The laser medium is doped into a host material, typically a transparent solid, although gaseous amplifiers are also possible. These are cast into a rod shape and pumped from the sides by light from flash-lamps. In some cases, a secondary pulsed laser is required to pump the laser medium, since the upper laser level has a short lifetime.

The long pulse is then amplified to near the damage threshold of the amplifiers and then widened. This reduces the intensity to below the damage threshold, allowing further amplification. However, wider rod amplifiers cannot be effectively pumped across their interior and so the secondary amplifiers must have a disk shape to allow pump light to enter from all sides. Also between these stages is a vacuum spatial filter. These remove any stray amplified light that is not co-linear with the laser light, as well as smoothing spatial phase imperfections introduced by the laser medium. This is achieved by focussing the beam between an aperture, improving the focal spot profile. There is also a Faraday isolator which prevents amplification of backscattered light.

A second set of gratings reverse the chirp and compress the pulse to the transform limited duration. This is the minimum duration allowed by the bandwidth theorem, given the energy spread (bandwidth) of the photons. All laser photons result from the same atomic transition and so have a very narrow energy spread, typically only broadened by Doppler effects or the local environment of each active atom. This bandwidth therefore depends on the host material. Typically the amplifier bandwidth is narrower than the initial oscillator, a process known as gain narrowing.

Due to the issue with the intensity dependent refractive index, even in air, this final stage must use reflective optics in a vacuum chamber. This means the final focussing and the target must also be housed in vacuum.

The Vulcan laser system is an example of a neodymium glass laser. The glass host allows easy manufacturing and scaling of the amplifiers. The active atoms are pumped by flash-lamps with a burst of current from a capacitor bank. However, much of the pump light is spectrally separated from the pumping bands of neodymium, wasting a lot of energy as heat. The amplifiers must then be cooled

before before the arrival of the next pulse. This limits the pulse repetition rate to once every 20 minutes.

This problem may be overcome by changing the host material to a crystal. The electric fields within the crystal spectrally broaden the pump bands and laser transition. In addition, heat can be dissipated quickly due to the high thermal conductivity of crystals. Another popular amplifier uses titanium atoms doped in a sapphire crystal, operating on a broadened laser transition from $700-900\,\mathrm{nm}$. Considering the bandwidth theorem, this wider bandwidth supports shorter duration pulses in the time domain.

The lifetime for spontaneous decay of the upper state in the titanium laser is approximately $3\,\mu s$. This means the whole crystal must be pumped within this time-frame, requiring a high brightness pump source. In practice, a second high-power laser system operating around $527\,\mathrm{nm}$ can be used to pump the titanium sapphire crystal.

In the frequency domain, a pulse with a Gaussian bandwidth may be written $E(\omega) = \frac{E_0}{\Delta\omega}\sqrt{\frac{2\ln 2}{\pi}}\exp\left(-\frac{(\omega-\omega_0)^2 2\ln 2}{\Delta\omega^2} + i\phi(\omega)\right)$, where $\phi(\omega)$ is the spectral phase and $\Delta\omega$ is the full width at half maximum (FWHM) bandwidth. Using $\phi = 0$ means all frequencies have the same phase and so the pulse is transform limited in the time domain, meaning it has the shortest possible duration. The electric field in the time domain is the Fourier transform of the spectrum, giving $E(t) = E_0\exp\left(-\frac{t^2\Delta\omega^2}{8\ln 2}\right)\exp\left(i\omega_0 t\right)$.

The pulse stretcher and compressors effectively apply a phase $\phi(\omega) = \alpha\omega^2$ where α is the chirp parameter. This has the effect of lengthening the pulse in the time domain.

The power in the spectral and time domains are $|E(\omega)|^2$ and $|E(t)|^2$. The full width at half maximum of these quantities are related by $T\Delta\omega > 4\ln(2)$. The transform limited duration, for a pulse with a Gaussian envelope, is then

$$T = 4\ln(2)\frac{1}{\omega}\frac{\omega}{\Delta\omega} \simeq 4\ln(2)\frac{\lambda}{2\pi c}\frac{\lambda}{\Delta\lambda}. \tag{1.1}$$

For lasers such as Vulcan petawatt with a $3\,\mathrm{nm}$ gain bandwidth, this gives the minimum duration $0.5\,\mathrm{ps}$. However, the titanium sapphire lasers have a spectral range from $750-850\,\mathrm{nm}$, giving a minimum duration of just $10\,\mathrm{fs}$. These ultra-short pulse lasers need a pulse energy of just $10-50\,\mathrm{J}$ to reach petawatt powers. This offers a much cheaper route to a petawatt facility, spawning a proliferation of high-power laser facilities worldwide over the last decade.

As discussed previously, the effect of the intensity dependent refractive index can be minimised with reflective optics in a vacuum chamber. The maximum achievable intensity may be evaluated using wave optics. The following derivation shows that it depends on the wavelength and focal angle of the beam.

Starting with a Gaussian transverse envelope beam at the focus, travelling along the z axis, one calculates the far field intensity profile using the Fraunhofer diffraction integral. X and Y are the transverse coordinates in a plane far from the focal plane and x and y are the transverse coordinates at the focus. The distance from focus $R = \sqrt{X^2 + Y^2 + z^2}$ and $k = 2\pi/\lambda$ is the wavenumber. For a Gaussian beam profile at its focus, the electric field is $E(r) \propto \exp(-(x^2 + y^2)/w_0^2)$ where w_0 is the beam waist. The Fraunhofer integral is

$$\frac{1}{i\lambda R} \int \exp\left(-\frac{x^2 + y^2}{w_0^2} + \frac{ikXx}{R} + \frac{ikYy}{R}\right) dxdy, \tag{1.2}$$

$$= \frac{1}{i\lambda R} \int \exp\left(-\frac{x^2}{w_0^2} + \frac{ikXx}{R}\right) dx \int \exp\left(-\frac{y^2}{w_0^2} + ikYyR\right) dy, \tag{1.3}$$

$$= \frac{\pi w_0^2}{i\lambda R} \exp\left(-k^2 w_0^2 (X^2 + Y^2)/4R^2\right). \tag{1.4}$$

By comparing the beam waist in the near field and far field, one finds the relation $w(z) = \lambda z/(\pi w_0)$, valid far from the Rayleigh range. Defining the F number $z/(2w)$ as the distance to focus divided by the full beam width, then $w_0 = 2\lambda F/\pi$ and the focal spot size of the beam electric field is $2w_0 = (4/\pi)\lambda F$. Therefore a smaller focal spot and higher intensity is achieved for a smaller F number, meaning either a wide focussing optic or a short focal length. This may be interpreted through the Heisenberg uncertainty principle, as the greater the spread in transverse momentum of the photons, the more spatially localised they can be at the focus.

A more detailed analysis gives w as a function of the longitudinal distance z from the focus.

$$w^2 = w_0^2 + z^2(w_0/L_R)^2, \quad L_R = \frac{\pi w_0^2}{\lambda} = \frac{4F^2\lambda}{\pi}. \tag{1.5}$$

The Rayleigh length L_R is the distance from the focus where the intensity has dropped by a factor of two. This means that a tight focus, with a small F number, has a high peak intensity over a small volume. For some applications a larger interaction volume is needed so F must be larger.

The power of an ideal Gaussian beam is

$$P = \int I dA = \int_0^\infty \int_0^{2\pi} I_0 e^{-2r^2/w^2} r \, dr d\phi = \pi w^2 I_0/2. \tag{1.6}$$

Therefore the peak intensity is $I_0 = 2P/(\pi w_0^2) = \pi P/(2\lambda^2 F^2)$. For a pulse of power 1 PW, wavelength 1μm and an F number of 3, this yields $10^{22}\,\mathrm{Wcm}^{-2}$.

Even higher intensities may be possible using the method of optical parametric chirped pulse amplification. This is a promising route to lasers exceeding 1 PW power. Some crystalline materials exhibit a non-linear refractive index and birefringence, where the refractive index depends on both intensity and the angle between the light propagation and the crystal axis. If a pump and seed wave with differing frequencies overlap, energy can be transferred from one to the other. Energy and momentum are conserved by the emission of a third wave known as the idler. Using a pulse with a known chirp allows compression of the seed via a set of reflective gratings.

For near-optical lasers, the photon energy is approximately 1 eV, which is less than the ionisation energy of many elements. However, the next section will show that the high electric field is sufficient to ionise the atoms.

1.3 Ionisation

For a laser operating in the visible and near infra-red spectrum, the wavelength greatly exceeds the size of the target atoms. This means one can approximate the effective atomic potential energy as the sum of the nuclear potential plus a uniform oscillating electric field from the electromagnetic wave

$$U(x) = -\frac{e^2}{4\pi\epsilon_0|x|} + exE_0\cos(\omega t). \tag{1.7}$$

This is plotted in Fig. 1.2 for times $\omega t = 2n\pi$, for integer n. This potential has a local maximum value of $U_0 = -2e\sqrt{eE_0/(4\pi\epsilon_0)}$.

In hydrogen, the electron kinetic energy in the ground state is $\alpha^2 mc^2/2$ and the potential energy is $-\alpha^2 mc^2$, where α is the fine structure constant. When the light decreases the effective potential, the electron is able to escape the potential well of the nucleus. If the bound electron escapes then it becomes a free electron and the remaining atom is left as a positively charged ion.

This process, known as barrier suppression field ionisation, occurs when $U_0 = -\alpha^2 mc^2/2$ or $E_0 = \pi\epsilon_0\alpha^4 m^2 c^4/4e^3 \simeq 3\times10^{10}\,\mathrm{Vm}^{-1}$. This is equivalent to an intensity of $1.4\times10^{14}\,\mathrm{Wcm}^{-2}$ for linearly polarised light. This is the field ionisation intensity for hydrogen and is independent of the laser wavelength.

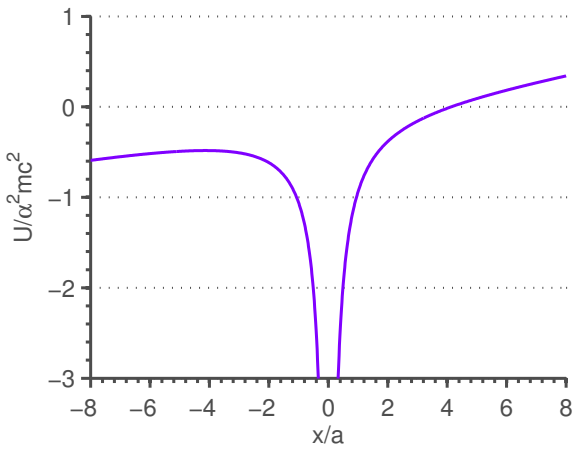


Figure 1.2: The electron potential energy in a hydrogen atom with a $3 \times 10^{10}\,\mathrm{V/m}$ electric field polarised along x. The ground state energy is $-\alpha^2 mc^2/2$, so the electron escapes the atomic potential because the barrier is suppressed. Units are given in terms of the Bohr radius a and the fine structure constant α.

Quantum effects mean that the electron has a finite probability of tunnelling through the potential barrier at intensities lower than $1.4 \times 10^{14}\,\mathrm{Wcm^{-2}}$. In fact, the probability depends exponentially on the potential barrier and so it only lowers the threshold ionisation intensity slightly. This process is known as tunnelling ionisation and it means a hydrogen gas is close to fully ionised at around $10^{14}\,\mathrm{Wcm^{-2}}$.

The field ionisation intensity depends on the ionisation potential, which is unique for each element. This means different elements field ionise at different intensities.

If $\hbar\omega$ is greater than the ionisation energy of the atom then the photon may also directly ionise it. This is known as photo-ionisation. The photons must be in the ultra-violet and x-ray region of the spectrum. Ionising the electrons with a principle quantum number of 1 gives the K absorption edge in the absorbed spectrum of light. As the photon energy is increased still further, the probability of ionisation decreases because the total momentum must be conserved.

It is possible for several photons to couple to the bound electron, meaning photo-ionisation is possible below the ionisation threshold. The probability decreases with the number of required photons, meaning this mechanism is more likely for high frequency light.

Field ionisation depends on the bound electron tunnelling through the lowered potential barrier, which must occur within the time period of the light. This means field ionisation is more likely for lower frequency light. The dimensionless Keldysh

parameter

$$\chi = \frac{mc\omega}{eE}\sqrt{\frac{2V}{mc^2}} \tag{1.8}$$

gives an indication of the ionisation regime. The light has frequency ω, with peak electric field E, incident on an atom with ionisation potential V. Values less than 1 mean tunnel ionisation is much more likely than multi-photon ionisation. For example, Hydrogen in a $1\,\mu$m wavelength infra-red pulse at intensity $10^{15}\,\mathrm{Wcm}^{-2}$ has a Keldysh parameter of 0.2, indicating tunnel ionisation.

1.4 Neutral gases

Under the classical theory of gases, the state variables are related through the ideal gas equation $p = nk_BT$. Here p is the pressure, n is the particle number density, k_B is the Boltzmann constant and T is the temperature. This equation of state is supported experimentally and is widely applicable under the following assumptions. First, the volume occupied by the gas must greatly exceed the total volume of the molecules and the duration of a collision is much less than the time between collisions. Second, there are no long range forces between the molecules, so the molecules only interact with one other molecule at a time through binary collisions. These collisions cause the system to undergo a random walk through the space of microstates. Since, in statistical equilibrium, each microstate is equally likely, the observed macrostate (with its associated observables such as the pressure) will be the one containing the most microstates. This is the macro-state that maximises the system's entropy.

The pressure on the walls of the container is exerted by particle collisions with the wall. Since the occurrence of a collision is a stochastic process, there is one final assumption for the state variables to be constant. The number of molecules in the container must be large, such that the stochastic variations in the observables are negligible with respect to the value of the observable. This also implies that the observed volume is much greater than the size of a molecule. In the limit that the number of molecules tends to infinity, each molecule is an infinitesimal element of a continuous fluid. This allows the use of calculus, rather than a sum over the finite number of molecules. This is known as the fluid approximation and is met for all practical applications since the number of molecules in a macroscopic sample is typically of the order of the Avogadro number $N_A \simeq 6.022 \times 10^{23}$. The

stochastic standard deviation in any measurement of state variables is then a factor $1/\sqrt{N_A} \simeq 10^{-12}$ smaller than the observed variable, a very good approximation to a continuous fluid.

Under this assumption, it allows one to treat the body of molecules not with a list of their individual positions and velocities, but as a continuous probability distribution function. This greatly simplifies the calculation of the dynamics, since instead of a list of the properties of each of $\simeq 10^{23}$ molecules, one just calculates the dynamics of the continuous probability distribution function (CPDF). The distribution of molecules varies as a function of position, due to gradients in the density, and as a function of velocity.

The distribution function is therefore written $f = f(t, \mathbf{x}, \mathbf{v})$. The Cartesian coordinates will be denoted x, y, z and their respective velocity components v_x, v_y, v_z. Although CPDF functions are usually normalised to 1, here I take the convention that $\int f(t, \mathbf{x}, \mathbf{v})\, d^3\mathbf{x} d^3\mathbf{v} = N(t)$, the total number of molecules in the system. I also define $n(t, \mathbf{x}) = \int f(t, \mathbf{x}, \mathbf{v})\, d^3\mathbf{v}$ as the local particle density, implying that $N(t) = \int n(t, \mathbf{x})\, d^3\mathbf{x}$. Since particles of different species are distinguishable, there must be a separate CPDF for each particle species in the sample, and the fluid approximation must hold for each.

The most probable state is equivalent to the state maximising the entropy of the system. It may be shown that the entropy of each molecular species is maximised under the velocity space distribution $f(t, \mathbf{x}, \mathbf{v}) = n(t, \mathbf{x}) f_0(\mathbf{v})$ where $f_0(\mathbf{v}) = (2\pi k_B T/m)^{-3/2} \exp(-\frac{m\mathbf{v}.\mathbf{v}}{2k_B T})$. This is referred to as the Maxwellian distribution. The distribution of the speed $v = \sqrt{\mathbf{v}.\mathbf{v}}$ is this multiplied by the Jacobian for spherical coordinates in velocity space. The angular distribution is uniform so the angular part of the Jacobian is just 4π and the radial part is v^2. This gives $f_0(v) = 4\pi(2\pi k_B T/m)^{-3/2} v^2 \exp(-\frac{mv^2}{2k_B T})$.

The distributions $f_0(\mathbf{v})$ and $f_0(v)$ are shown in Fig. 1.3. The ensemble average of a quantity is denoted through an over-line and is defined by $\overline{g(\mathbf{v})} = \int g(\mathbf{v}) f(\mathbf{v}) d^3\mathbf{v}$. The distribution gives physical meaning to the temperature, since the average kinetic energy of a single particle is $\frac{1}{2}m\overline{v^2} = \frac{3}{2}k_B T$. The temperature should therefore have units of energy.

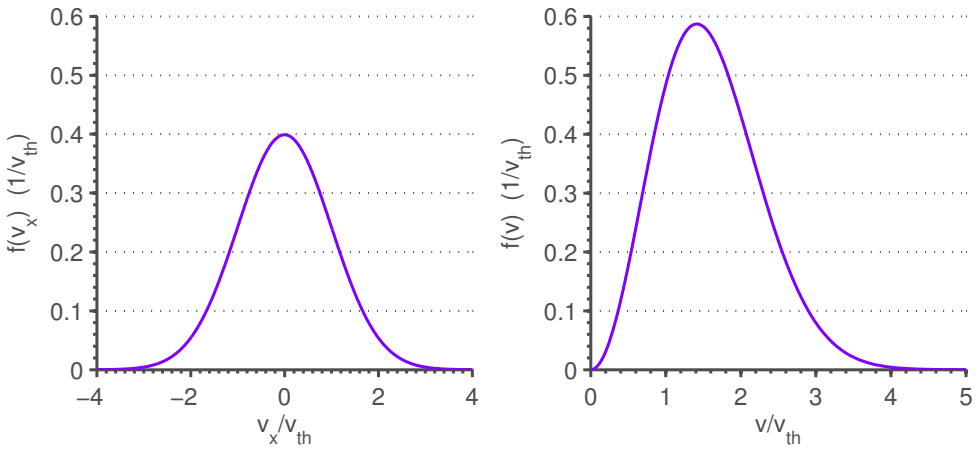


Figure 1.3: Panel (a) shows the most likely velocity distribution function, along one direction. This is known as a Maxwellian distribution. In spherical coordinates, the angular distribution is uniform and the speed distribution is shown in panel (b). The thermal velocity is $v_{th} = \sqrt{k_B T/m}$.

1.5 Plasmas

I now relax the assumption that there are no long range forces between the particles. Clearly binary collisions still play a role, just as in the neutral gas case, but now each particle in the sample is coupled to many particles at once. Of the four fundamental interactions, only gravity and the electromagnetic force are long range. The gravitational attraction between two electrons is smaller than the electric repulsion by a factor $4\pi G m^2 \epsilon_0 / e^2 = 2.4 \times 10^{-43}$. The long range force for free electrons and ions is therefore dominated by the electromagnetic interaction.

The particle distribution is given by f, but the charged particles now generate electromagnetic fields and the fields affect the particle evolution. The electromagnetic field is described by the Maxwell equations:

$$\nabla . \mathbf{E} = \frac{\rho}{\epsilon_0} \tag{1.9}$$

$$\nabla . \mathbf{B} = 0 \tag{1.10}$$

$$\nabla \times \mathbf{E} = -\frac{\partial \mathbf{B}}{\partial t} \tag{1.11}$$

$$\nabla \times \mathbf{B} = \mu_0 \mathbf{J} + \mu_0 \epsilon_0 \frac{\partial \mathbf{E}}{\partial t} \tag{1.12}$$

The electric field $\mathbf{E}(\mathbf{x},t)$ and the magnetic field $\mathbf{B}(\mathbf{x},t)$ are coupled through the four partial differential equations. For a species of charge q, its charge density ρ

and current density $\mathbf{J}$ are defined through the CPDF:

$$\rho = q \int f d^3\mathbf{v} = qn \tag{1.13}$$

$$\mathbf{J} = q \int \mathbf{v} f d^3\mathbf{v} \tag{1.14}$$

This is summed over all species to find the total charge and current density. The electric and magnetic fields can be derived via a scalar and vector potential, written as ϕ and $\mathbf{A}$. Writing $\mathbf{B} = \nabla \times \mathbf{A}$ ensures that $\nabla.\mathbf{B} = 0$. The third Maxwell equation then gives $\nabla \times (\mathbf{E} + \frac{\partial \mathbf{A}}{\partial \mathbf{t}}) = 0$. The general solution is $\mathbf{E} + \frac{\partial \mathbf{A}}{\partial \mathbf{t}} = -\nabla\phi$ for arbitrary ϕ, since the gradient of a scalar field always has zero curl. This means the fields can be found from a scalar and vector potential

$$\mathbf{E} = -\nabla\phi - \frac{\partial \mathbf{A}}{\partial t}, \tag{1.15}$$

$$\mathbf{B} = \nabla \times \mathbf{A}. \tag{1.16}$$

Under vacuum, $\rho = J = 0$. Substituting (1.12) and (1.9) into (1.11) leads to a wave equation for $\mathbf{E}$. The phase velocity is $c = 1/\sqrt{\epsilon_0 \mu_0}$. The plane electromagnetic wave solution has electric field $\mathbf{E} = E_0 \hat{\mathbf{x}} \sin(kz - \omega t)$ and the magnetic field is $\mathbf{B} = (E_0/c)\hat{\mathbf{y}} \sin(kz - \omega t)$. This is an electromagnetic wave that is linearly polarised along the x axis and travels along the z axis.

1.5.1 The Debye length

Consider the effect of placing a positive test charge or electrode inside an otherwise charge balanced plasma at a temperature T. The negative charges in the bulk plasma are attracted to this test charge, whereas the positive ones are repelled. This gives rise to a potential which under equilibrium, cancels that of the test charge. Due to thermal fluctuations, this potential cancellation will not be perfect and it has a characteristic length scale known as the Debye length. The plasma has the effect of screening the field of the test charge.

The potential around an isolated charge is given by the Coulomb potential $\phi = Q/(4\pi\epsilon_0 r)$, where r is the distance from the test charge Q. In a plasma, this affects the equilibrium charge distribution of the electrons and ions.

The Boltzmann distribution gives the probability of a particle penetrating to a potential energy $q\phi$ in a steady state plasma. The local particle density is proportional to this probability, so $n(r) = n_0 \exp\left(-\frac{q\phi}{k_B T}\right)$ where ϕ is the unknown steady

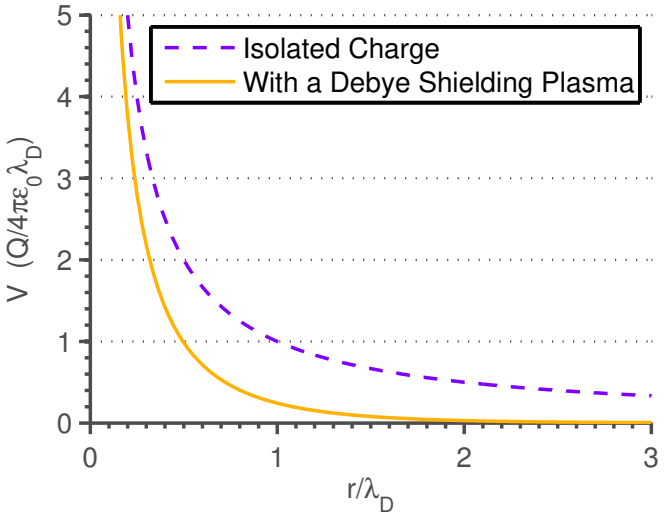


Figure 1.4: The effect of plasma Debye shielding on the potential around a test charge. The realignment of the local charges causes the potential of the test charge to be suppressed above a distance λ_D.

state potential such that $\mathbf{E} = -\nabla\phi$. Far away from the test charge the plasma is unaffected, so that $n_0 = n(r \to \infty)$ is the constant unperturbed number density of free charged particles.

The Boltzmann distribution applies to positive and negative plasma particles of charge q. For the solution in the region where $e\phi/(k_B T) \ll 1$, simply Taylor expand the Boltzmann factors. Assuming that the ions are protons, equations (1.9) and (1.13) then give

$$-\nabla^2\phi = \frac{\rho}{\epsilon_0} = \frac{en_0}{\epsilon_0}\exp\left(-\frac{e\phi}{k_B T}\right) - \frac{en_0}{\epsilon_0}\exp\left(\frac{e\phi}{k_B T}\right) \simeq -2en_0\frac{e\phi}{\epsilon_0 k_B T}. \tag{1.17}$$

Using the Coulomb potential as a trial solution $\phi(r) = f(r)/r$ gives $f''(r) = 2f/\lambda_D^2$, where

$$\lambda_D = \sqrt{\frac{\epsilon_0 k_B T}{n_0 e^2}}. \tag{1.18}$$

The solution is $\phi = \frac{Q}{4\pi\epsilon_0 r}e^{-r\sqrt{2}/\lambda_D}$ and n_0 is the unperturbed free electron density. The potential is therefore the usual Coulomb potential at positions close to the charge, but at radii of λ_D or greater it becomes exponentially suppressed by the shielding, as shown in Fig. 1.4.

This suggests that on scales less than λ_D, the potential is rapidly varying and depends on the precise positions of the nearest particles, but on scales larger than λ_D these variations are suppressed by the Debye screening effect. The fields therefore have two aspects. On scales less than λ_D the fields are rapidly varying and

stochastic. The field is dependent on the precise locations of the $4\pi n\lambda_D^3/3$ particles within the local Debye sphere. On scales greater than λ_D the particles move to screen these stochastic potentials and the remaining smooth fields are just due to collective charge separations and currents over the bulk of the plasma. It is these fields which cause the long range interactions. They are tractable analytically using the fluid approximation, whereas motion on scales less than λ_D must use a particle description.

The Debye length gives the maximum length scale over which charge separation is possible in a steady state plasma. In high-power laser interactions with a gas target $\lambda_D \simeq 100\,\mathrm{nm}$ and in laser solid interactions $\lambda_D \simeq 1\,\mathrm{nm}$. However, rapidly varying plasmas and plasma waves can support charge separation over much longer distances.

The screening on scales of λ_D requires a large number of free particles to accurately move to cancel out the test potential. Therefore smoothing of the stochastic fields via screening requires a large number of particles within a range λ_D and so the Debye number $n\lambda_D^3$ is large. $n\lambda_D^3 \propto n^{-1/2}(k_BT)^{3/2}$ increases with the plasma temperature and decreases with increasing density.

Neutral or weakly ionised gases have a low free electron density, such that the Debye length and Debye number tend to infinity. However, these systems do not exhibit collective behaviour because the system size is typically much smaller than λ_D.

I will now show that a plasma with $n\lambda_D^3 \gg 1$ must also be weakly coupled, meaning the kinetic energy of a typical free electron greatly exceeds its potential energy. This gives $k_BT \gg e\phi$, the same assumption as in the derivation (1.17). Writing ϕ as a Coulomb potential and taking the distance to the nearest ion as the average free electron separation $n^{-1/3}$, one finds that

$$\frac{k_BT}{e\phi} = \frac{4\pi\epsilon_0 k_BT}{e^2 n^{1/3}} = 4\pi n^{2/3}\lambda_D^2 = 4\pi(n\lambda_D^3)^{2/3}. \tag{1.19}$$

The assumptions are therefore equivalent. A high Debye number implies that the average free electron kinetic energy is much higher than the binding energy to its nearest ion, allowing effective motion of the electrons to shield the stochastic fields. Fig. 1.5 shows the Debye number as a function of the plasma density and temperature. Various physical and astronomical plasmas are indicated on the axes. This thesis will only investigate plasmas for which $n\lambda_D^3 > 1$, the weakly coupled ideal plasma section above the orange line.

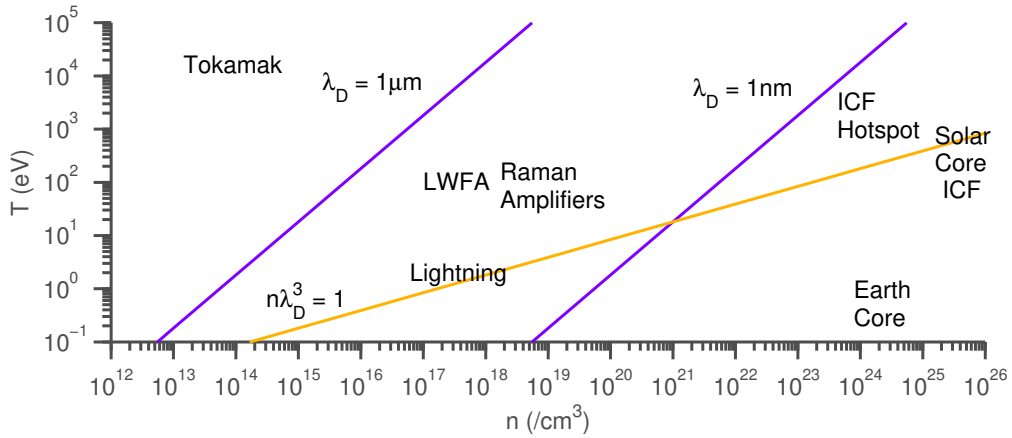


Figure 1.5: Various experimental and astrophysical plasmas are shown in the plane of free electron number density and temperature. Contours of the Deybe length are indicated by the purple lines and the Debye number with the orange line. Tokamaks aim to contain a burning fusion plasma using magnetic fields, whereas inertial confinement fusion (ICF) uses a much denser compressed fuel. Laser wakefield acceleration (LWFA) is a promising route to miniaturise particle accelerators, requiring similar plasma conditions to Raman amplification at visible wavelengths.

Measurable plasmas span over fourteen orders of magnitude of density, from the low values in the ionosphere and magnetic confinement fusion, up to stellar plasmas and laser fusion targets. High-power laser systems can induce weakly and strongly coupled plasma conditions, depending on the specific target and laser pulse.

Following the analysis of Debye shielding, the electric and magnetic fields can be split into a smooth, long range component, hereafter labelled $\mathbf{E}(\mathbf{x}, t)$ and $\mathbf{B}(\mathbf{x}, t)$, and a stochastic, short range component on scales less than λ_D.

These two components will be analysed in the following sections. The short range stochastic component must be treated on a particle level and leads to binary particle collisions that tend to increase the plasma entropy, whereas the long range components are treated with a fluid description. The effect on the particle velocity is decoupled into a stochastic collisional part for interactions on short scales, and a large scale plasma behaviour which supports collective plasma waves.

1.5.2 Particle collisions

On a scale less than λ_D one cannot assume that the potentials of the particles are shielded, so the situation is a many body problem of particles with a Coulomb potential. The influence of the small scale fields is that each particle undergoes a separate Coulomb collision with every other particle within a distance λ_D. The

contributions of any particles at distance greater than λ_D are negligible because their potentials are exponentially suppressed.

To calculate the effect of binary collisions between electrons and ions, assume that the ion is infinitely massive and therefore the Rutherford scattering cross section is applicable. The deflection angle of an electron, in a collision with a single ion, is

$$\theta = 2\tan^{-1}\left(\frac{Ze^2}{4\pi\epsilon_0 m v_0^2 b}\right), \tag{1.20}$$

$$\simeq \frac{Ze^2}{2\pi\epsilon_0 m v_0^2 b}. \tag{1.21}$$

The expansion is valid for small deflections $\theta < 1$. The initial incoming speed of the electron is v_0 and its impact parameter b is the distance of closest approach to the ion, assuming a constant velocity straight line trajectory. Ze is the charge of the ion. The expansion is valid when $b > b_{min} = (Ze^2)/(2\pi\epsilon_0 m v_0^2)$. For $b < b_{min}$ the electron suffers a finite deflection where θ is of order 1 and the expansion is invalid.

The change in velocity has magnitude approximately $v_0\theta$. For collisions with infinitely massive ions, a beam of electrons will retain its average speed, but will become dispersed in angle. This means the velocity variance increases.

Therefore the relevant quantity for measuring the electron ion collision time is the mean square deviation in the velocity. The probability of a collision with impact parameter b in a time dt has a CPDF given by $p.\,db\,dt = n_i v_0.\,d\sigma\,dt = n_i v_0.\,2\pi b.\,db\,dt$ where σ is the cross section and n_i is the ion number density. The square velocity deviation must be integrated over this probability distribution to give the mean square velocity deviation

$$\partial_t\overline{(\delta v)^2} = \int_0^\infty (\delta v)^2 n_i v_0.2\pi b\,db = \int_{b_{min}}^{\lambda_D} (v_0\theta)^2\, n_i v_0.\,2\pi b\,db + \int_0^{b_{min}} (v_0)^2 n_i v_0.\,2\pi b\,db \tag{1.22}$$

$$= \frac{n_i Z^2 e^4}{2\pi\epsilon_0^2 m^2 v_0}\left(\ln(\lambda_D/b_{min}) + \frac{1}{2}\right). \tag{1.23}$$

Here I have estimated the contribution from $b < b_{min}$ by taking $\theta = 1$. In addition, no collisions are considered at distances greater than λ_D because their Coulomb potentials are exponentially suppressed by the Debye shielding. The second integral is finite and found to be smaller than the first if $b_{min} \ll \lambda_D$. It may be shown that $\lambda_D/b_{min} = 3\pi n\lambda_D^3/Z$, so if the Debye number is large then the previous analysis is valid and the second term in equation (1.23) is negligible. This means that in

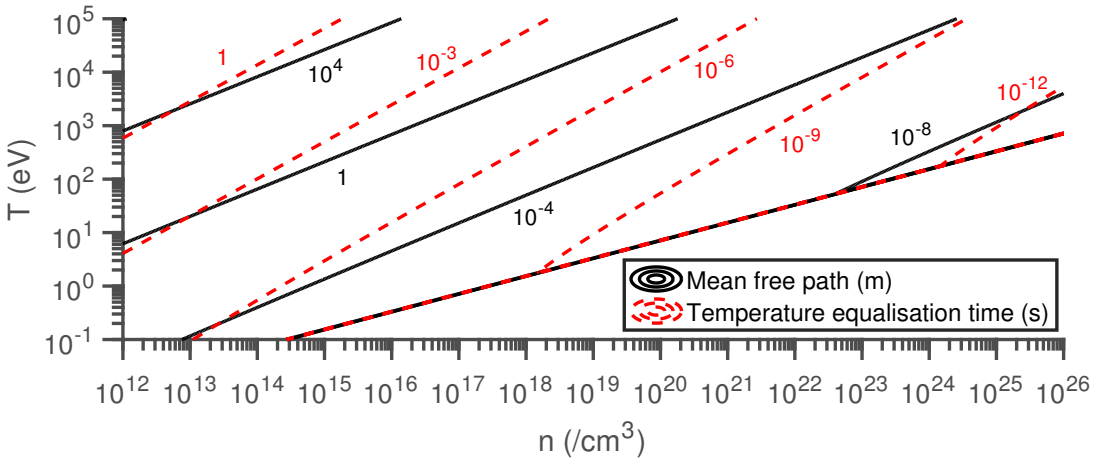


Figure 1.6: The black solid contours give the electron ion mean free path in meters. The red dashed contours give the electron ion temperature equalisation time τ_E in seconds. The analysis is invalid in the strongly coupled plasma region in the bottom right.

a weakly coupled plasma, many successive small collisions with particles satisfying $b_{min} < b < \lambda_D$ are responsible for the velocity spreading. Single large angle deflections with $b < b_{min}$ are relatively less important.

The effect of the deflections will become important when the total average deflection angle is approximately 1 radian, giving $v_0^2 \simeq \overline{(\delta v)^2}$. The electron–ion collision time is then defined by $v_0^2/\tau_{ei} = \partial_t \overline{(\delta v)^2}$, giving

$$\tau_{ei} = \frac{2\pi\epsilon_0^2 m^2 v_0^3}{n_i Z^2 e^4 \ln(3\pi n \lambda_D^3/Z)}. \tag{1.24}$$

The collision timescale increases with the cube of the electron velocity. For a Maxwellian distribution it increases with $T^{3/2}$. The collision time is inversely proportional to the ion density, as is expected on physical grounds. There is also a strong dependence on Z. Indeed, if the charge density $n_i Z$ is kept constant then $\tau_{ei} \propto 1/Z$, suggesting that plasmas of heavy elements such as gold are more collisional.

1.5.3 Stopping power

The previous analysis assumed that the ion is infinitely massive and so the electron retains all of its energy. In reality, the ion has a finite mass and therefore gains energy in the collision. The energy of the faster electron is decreased. This equalises the electron and ion temperatures. The rate of energy loss can be calculated as follows.

	Collision time	Temperature equilibriation time
Electron ion	1	m_i/m
Electron electron	Z	Z
Ion ion	$\left(\frac{T_i}{T_e}\right)^{3/2} \sqrt{\frac{m_i}{m}} \frac{1}{Z^2}$	$\left(\frac{T_i}{T_e}\right)^{3/2} \sqrt{\frac{m_i}{m}} \frac{1}{Z^2}$

Table 1.1: Collision timescales for the electron and ion species in units of τ_{ei}, given by equation (1.24)

The dominant collisions are those with a small deflection angle, so the small angle approximation is valid. Due to conservation of momentum, the change in ion energy is

$$\delta E = \frac{(\delta p)^2}{2m_i} = \frac{(mv_0\theta)^2}{2m_i} = \frac{Z^2e^4}{8\pi^2 m_i \epsilon_0^2 v_0^2 b^2}. \tag{1.25}$$

The energy loss per collision must be integrated over the CPDF to give the average stopping power,

$$\frac{dE}{dt} = \int_{b_{min}}^{\lambda_D} \frac{Z^2e^4}{8\pi^2 m_i \epsilon_0^2 v_0^2 b^2} 2\pi b n_i v_0 \, db = \frac{Z^2e^4 n_i \ln(3\pi n \lambda_D^3/Z)}{4\pi m_i v_0 \epsilon_0^2}. \tag{1.26}$$

The timescale for energy loss and equalisation of the electron and ion temperatures is

$$\tau_E = \frac{1}{2}mv_0^2 \left(\frac{dE}{dT}\right)^{-1} = \frac{2\pi\epsilon_0^2 m m_i v_0^3}{Z^2e^4 n_i \ln(3\pi n \lambda_D^3/Z)} = \tau_{ei}\frac{m_i}{m}. \tag{1.27}$$

Since the ion mass m_i is much greater than the electron mass m, reaching a temperature equilibrium takes much longer than the electron ion collision time. This is because the energy transfer is inefficient due to the large mass difference. Any external energy absorbed by the plasma is often transferred mainly to the electrons due to their lower mass. The timescale for reaching equilibrium with the ions is often much longer than other phenomena in laser plasma interactions, so the electron temperature can be much higher than the ion temperature in many situations of interest.

Table 1.1 gives the relevant timescales. Since $m_i \gg m$, the electrons reach equilibrium first, then the ions reach equilibrium at a separate temperature. Finally the ion and electron temperatures equalise. Fig. 1.6 shows the electron ion mean free path $\lambda_{ei} = \tau_{ei} v_{th}$, and the equilibriation time, as a function of the plasma conditions.

Consideration of Table 1.1 shows that an external beam of electrons will be deflected in angle by both electron and ion collisions, however it is slowed mostly

by electron collisions. This means the beam energy is transferred mostly to electrons and so the electron temperature is typically greater than the ion temperature.

1.6 Plasma evolution at large scales

All that remains to complete the description is an equation for the evolution of f, in terms of the fields.

I use the Einstein summation convention where the indices a, b, c, i, j, k and l can denote each of the Cartesian directions x, y and z. Contracted indices are summed over x, y and z. Derivatives with indices a, b, c denote spatial partial derivatives and derivatives with indices i, j, k, l denote velocity space derivatives. For example, $\partial_a = \frac{\partial}{\partial x_a}$ and $\partial_i = \frac{\partial}{\partial v_i}$. The time, velocity and position coordinates are separate, so that $\partial_t v_j = \left(\frac{\partial v_j}{\partial t}\right)_{\mathbf{x},\mathbf{v}} = 0$ and $\partial_a v_j = \frac{\partial v_j}{\partial x_a} = 0$.

Since particles cannot be created or destroyed, f obeys a six dimensional continuity equation

$$\partial_t f + \partial_a (f v_a) + \partial_i \left(f \frac{dv_i}{dt} \right) = 0. \tag{1.28}$$

In the second term I have written $\frac{d\mathbf{x}}{dt} = \mathbf{v}$.

For the last term in equation (1.28) one must know the forces on a fluid element. As discussed, these forces can be decoupled into particle collisions on small scales and the effects of the smooth fields $\mathbf{E}$ and $\mathbf{B}$ on large scales. For non-relativistic motion, the effect of the large scale fields is given by the usual electromagnetic acceleration $\frac{d\mathbf{v}}{dt} = \frac{q}{m}(\mathbf{E} + \mathbf{v} \times \mathbf{B})$. This gives

$$\partial_j \left(\frac{dv_i}{dt} \right) = \frac{q}{m} \partial_j (E_i + \epsilon_{ikl} v_k B_l) = \frac{q}{m} \epsilon_{ikl} \delta_{jk} B_l = \frac{q}{m} \epsilon_{ijl} B_l. \tag{1.29}$$

Here ϵ_{ijk} is the anti-symmetric tensor and δ_{ij} is the Kronecker delta tensor. They obey the identity $\epsilon_{ijk}\delta_{ij} = 0$. I have also used the identities $\partial_j E_i = 0$ (since $\mathbf{E}$ does not depend on velocity) and $\partial_i v_j = \delta_{ij}$.

The collision term is more complicated and is left separately. Combining equations (1.28) and (1.29) yields

$$\partial_t f + f \partial_a v_a + v_a \partial_a f + f \partial_i \left(\frac{dv_i}{dt} \right) + \left(\frac{dv_i}{dt} \right) \partial_i f = 0, \tag{1.30}$$

$$\partial_t f + v_a \partial_a f + \frac{q}{m} \epsilon_{iil} B_l f + \frac{q}{m} (E_i + \epsilon_{ijk} v_j B_k) \partial_i f = \left(\frac{\partial f}{\partial t} \right)_{coll}, \tag{1.31}$$

$$\frac{\partial f}{\partial t} + \mathbf{v}.\frac{\partial f}{\partial \mathbf{x}} + \frac{q}{m}(\mathbf{E} + \mathbf{v} \times \mathbf{B}).\frac{\partial f}{\partial \mathbf{v}} = \left(\frac{\partial f}{\partial t} \right)_{coll}. \tag{1.32}$$

This is the Vlasov-Fokker-Planck equation and is true for each species separately. The second line used the result (1.29) and the fact that $\partial_a v_a = 0$ since the spatial and velocity coordinates are independent.

I now make the assumption that the phenomena of interest have a timescale shorter than τ_{ei}. This means the plasma is collisionless on the timescale of interest and the collision term in the above equation is neglected. The equation becomes the Vlasov equation. In component and vector forms it is written

$$\partial_t f + v_a \partial_a f + \frac{q}{m}(E_i + \epsilon_{ijk} v_j B_k)\partial_i f = 0, \tag{1.33}$$

$$\frac{\partial f}{\partial t} + \mathbf{v}.\frac{\partial f}{\partial \mathbf{x}} + \frac{q}{m}(\mathbf{E} + \mathbf{v} \times \mathbf{B}).\frac{\partial f}{\partial \mathbf{v}} = 0. \tag{1.34}$$

1.6.1 The Fluid equations

The number density is $n = \int f d^3v$ and the average velocity is $u_j = (1/n) \int v_j f d^3v$. One can integrate the Vlasov equation over velocity space to obtain

$$\partial_t \int f d^3v + \int v_a \partial_a f d^3v + \frac{q}{m} \int (E_i + \epsilon_{ijk} v_j B_k)\partial_i f d^3v = 0, \tag{1.35}$$

$$\partial_t n + \partial_a \int v_a f d^3v - \frac{q}{m} \int (\partial_i E_i + \epsilon_{ijk} B_k \partial_i v_j) f \, d^3v = 0, \tag{1.36}$$

$$\partial_t n + \partial_a (n u_a) = 0. \tag{1.37}$$

The third integral was integrated by parts. The boundary term goes to zero since f tends exponentially to zero at infinity. I have again used $\partial_a v_a = 0$, $\partial_i E_i = 0$ and $\partial_i v_j = \delta_{ij}$. The result is a continuity equation for a fluid of density n. The vector differential and integral forms are

$$\frac{\partial n}{\partial t} + \nabla.(n\mathbf{u}) = 0, \tag{1.38}$$

$$\frac{\partial}{\partial t} \int_V n d^3x = - \int (n\mathbf{u}).d\mathbf{s}, \tag{1.39}$$

where the last integral is over the surface of the volume V. The integral form shows that the rate of change of particle number in a volume element is equal to the flux of particles leaving the volume, integrated over its outer surface. This is the first of the plasma fluid equations and it applies separately for each species.

The next equation can be simplified by defining the pressure tensor $p_{ij} = m \int d^3v (v_i - u_i)(v_j - u_j) f$. The second equation is found by multiplying the Vlasov

equation by v_l and integrating over velocity to get

$$0 = \int v_l \partial_t f + v_l v_a \partial_a f + \frac{q}{m} v_l (E_i + \epsilon_{ijk} v_j B_k) \partial_i f \, d^3 v,$$

$$0 = \partial_t \int v_l f d^3 v + \partial_a \int d^3 v (v_l - u_l + u_l)(v_a - u_a + u_a) f$$
$$- \frac{q}{m} \int d^3 v f \partial_i (v_l E_i + \epsilon_{ijk} v_j B_k v_l),$$

$$0 = \partial_t (n u_l) + \partial_a p_{la}/m + \partial_a \int d^3 v u_l (v_a - u_a) f + \partial_a \int d^3 v u_a (v_l - u_l) f$$
$$+ \partial_a \int d^3 v u_a u_l f - \frac{q}{m} \int d^3 v f (E_i \partial_i v_l + \epsilon_{ijk} B_k (v_j \partial_i v_l + v_l \partial_i v_j)),$$

$$\partial_t (n u_l) + \partial_a p_{la}/m + \partial_a (n u_l u_a) = \frac{qn}{m} (E_l + \epsilon_{ljk} u_j B_k), \tag{1.40}$$

$$\partial_t u_l + u_a \partial_a u_l = -\frac{1}{nm} \partial_a p_{la} + \frac{q}{m} (E_l + \epsilon_{ljk} u_j B_k), \tag{1.41}$$

$$\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u}.\nabla \mathbf{u} = -\frac{1}{nm} \nabla . \underline{\mathbf{p}} + \frac{q}{m} (\mathbf{E} + \mathbf{u} \times \mathbf{B}). \tag{1.42}$$

This is the second plasma fluid equation and gives the development of the fluid velocity under the plasma pressure and electromagnetic forces. Equation (1.38) has been used to reach the last line. To proceed further, an expression for the pressure tensor is needed. However, this depends on higher moments of the Vlasov equation, meaning the system of equations is never closed.

The pressure tensor can be evaluated by assuming an isotropic, Maxwellian plasma with $f \propto n(\mathbf{x}, t) \exp\left(-\frac{m\mathbf{v}.\mathbf{v}}{2k_B T}\right)$. This gives $p_{ij} = p\delta_{ij}$. For example, the z component is

$$p = p_{zz} = m \int d^3 v \, v_z^2 n \exp\left(-\frac{m(v_x^2 + v_y^2 + v_z^2)}{2k_B T}\right) \left(\frac{m}{2\pi k_B T}\right)^{3/2}, \tag{1.43}$$

$$= nm \int dv_z \, v_z^2 \exp\left(-\frac{m v_z^2}{2k_B T}\right) \sqrt{\frac{m}{2\pi k_B T}}, \tag{1.44}$$

$$= nm \sqrt{\frac{m}{2\pi k_B T}} \frac{\sqrt{\pi}}{2} \left(\frac{2k_B T}{m}\right)^{3/2}, \tag{1.45}$$

$$= n k_B T. \tag{1.46}$$

The Vlasov equation has yielded the ideal gas law, assuming the isotropic Maxwellian velocity distribution. This expression for the pressure is known as an equation of state. Clearly $p = p(n, T)$ is a function only of density and temperature. The

equations show how n and $\mathbf{u}$ develop as a function of time under the action of the plasma pressure and the electromagnetic force. These are calculated using the equation of state and the Maxwell equations. The equations apply to each plasma species separately.

1.6.2 Electron plasma waves

I will now use the fluid description to derive an oscillation of the plasma known as an electron plasma wave (or Langmuir wave) and find its dispersion relation [57]. A separation of the plasma electrons and ions creates an electric field and the perturbation then propagates. Classical perturbation theory is used to linearise the fluid equations by treating the perturbation as a small quantity, denoted by a tilde. Now, constrain the problem to one spatial dimension and write the bulk electron velocity as $u = \tilde{u}$, the free electron density as a uniform value plus a perturbation $n = n_0 + \tilde{n}$ and the electromagnetic fields as $E = \tilde{E}$ and $B = \tilde{B}$. The ions are assumed to remain uniform, due to their higher mass. The fluid equations (1.38, 1.42) and the first Maxwell equation (1.9) become

$$\partial_t \tilde{n} + n_0 \partial_x \tilde{u} = 0, \tag{1.47}$$

$$\partial_t \tilde{u} = -\frac{1}{n_0 m}\partial_x p - \frac{e}{m}\tilde{E}, \tag{1.48}$$

$$\partial_x \tilde{E} = -\frac{e\tilde{n}}{\epsilon_0}. \tag{1.49}$$

The magnetic field and $\mathbf{u}.\nabla\mathbf{u}$ terms are quadratic in the small quantities and have been neglected. Since the timescales are shorter than τ_{ei}, the compression waves must be adiabatic and therefore $p \propto n^\nu$ where ν is the adiabatic index. Using the chain rule, the second equation becomes

$$\partial_t \tilde{u} = -\frac{1}{n_0 m}\nu\frac{p}{n_0}\partial_x \tilde{n} - \frac{e}{m}\tilde{E}. \tag{1.50}$$

The mean pressure value is given by $p = n_0 k_B T$. Taking the time derivative of (1.47) and the spatial derivative of (1.50) gives

$$\partial_t^2 \tilde{n} + n_0 \partial_t \partial_x \tilde{u} = 0, \tag{1.51}$$

$$\partial_t \partial_x \tilde{u} + \frac{\nu k_B T}{n_0 m}\partial_x^2 \tilde{n} = -\frac{e}{m}\partial_x \tilde{E} = \frac{e^2 \tilde{n}}{\epsilon_0 m}, \tag{1.52}$$

where I have also used the Maxwell equation (1.49). Finally, combining these equations yields a wave equation. For oscillations proportional to $e^{i(\omega t - kx)}$, the dispersion

relation is:

$$\frac{\nu k_B T}{m}\partial_x^2\tilde{n} - \partial_t^2\tilde{n} = \frac{n_0 e^2}{\epsilon_0 m}\tilde{n}, \tag{1.53}$$

$$\omega^2 = \omega_p^2 + \nu\frac{k_B T}{m}k^2, \tag{1.54}$$

$$\omega^2 = \omega_p^2(1 + \nu k^2\lambda_D^2). \tag{1.55}$$

The waves have a characteristic oscillation frequency of $\omega_p = \sqrt{n_0 e^2/(\epsilon_0 m)}$, known as the plasma frequency. For waves with a wavelength much longer than λ_D, the frequency is essentially ω_p. The phase velocity is approximately ω_p/k. If the plasma is heated so the Debye length approaches the Langmuir wavelength, then the frequency is increased slightly.

The plasma frequency has a physical interpretation as the timescale for a particle at the average thermal velocity to travel a distance λ_D, or $\lambda_D\omega_p = v_{th}$.

The plasma waves are damped by the particle collisions which act to increase the entropy and transfer the energy density of the wave to random thermal motion, on a timescale τ_{ei}. The waves are collisionless and last for many periods so long as

$$1 \ll \omega_p\tau_{ei} = \sqrt{\frac{ne^2}{\epsilon_0 m}}\frac{2\pi\epsilon_0^2 m^2 v_0^3}{nZe^4\ln(3\pi n\lambda_D^3/Z)}, \tag{1.56}$$

$$= \frac{3\pi n\lambda_D^3}{Z\ln(3\pi n\lambda_D^3/Z)}. \tag{1.57}$$

Weakly coupled plasmas with $n\lambda_D^3 \gg 1$ are therefore collisionless for phenomena on timescales ω_p. The neglect of the collisional term in the Vlasov equation is therefore justified. Sustained electron plasma waves can only exist in weakly coupled plasmas; if $n\lambda_D^3 < 1$ then they will be damped within one period. Importantly, the collision timescale is inversely proportional to Z, meaning plasma waves are damped more in plasmas of heavy elements.

In a plane plasma wave, the electric field is purely longitudinal so $\nabla \times \mathbf{E} = 0$. Using the third Maxwell equation (1.11), the waves must have $\mathbf{B} = 0$. The fourth Maxwell equation (1.12) then shows that the electron current exactly cancels the displacement current, so

$$J_i + \epsilon_0\partial_t E_i = 0 \tag{1.58}$$

To recap, $n\lambda_D^3 \gg 1$ implies that the Debye shielding is effective, the plasma is weakly coupled, the fluid approximation is valid on scales larger than λ_D and phenomena on timescales ω_p are collisionless.

1.6.3 Landau damping of electron plasma waves

Since the Maxwellian distribution f_0 is positive everywhere, there are always some plasma electrons with the same velocity ω_p/k as the plasma wave. These electrons become trapped within the wave and become accelerated or decelerated by its electric field. Since $\frac{\partial f_0}{\partial v} < 0$ there are always more electrons accelerated than decelerated and so the wave transfers energy to electrons moving near the phase velocity. This means that all plasma waves, even in collisionless plasmas, are damped [60].

To find the damping rate, one can linearise the one-dimensional Vlasov equation using $E(t,x) = \tilde{E}e^{i(\omega t - kx)}$. The plasma wave has frequency ω and wavenumber k. The imaginary part of ω damps the wave. The distribution function $f(t,x,v) = f_0(v) + \tilde{f}(v)e^{i(\omega t - kx)}$ can be represented with a one-dimensional Maxwellian f_0 at temperature T, plus a small perturbation for the linear plasma wave. The linearised part has been split into its time and space dependence, giving the plasma wave, and the velocity dependent part $\tilde{f}(v)$. The Maxwellian part has $\partial_x f_0 = \partial_t f_0 = 0$. Neglecting the terms that are quadratic in the small quantities, the Vlasov equation becomes

$$0 = \tilde{f}(\partial_t + v\partial_x)e^{i(\omega t - kx)} - \frac{e\tilde{E}}{m}\partial_v f_0, \tag{1.59}$$

$$i\omega\tilde{f} - ikv\tilde{f} = \frac{e\tilde{E}}{m}\partial_v f_0, \tag{1.60}$$

$$\tilde{f} = -i\frac{e\tilde{E}}{m\omega}\frac{\partial_v f_0}{1 - \frac{kv}{\omega}}. \tag{1.61}$$

Equations (1.58) and (1.14) give an expression for the electric field $\tilde{E}$. Substituting equation (1.61) into this cancels out the electric field, leading to

$$\partial_t\tilde{E} = -\frac{\tilde{J}}{\epsilon_0} = \frac{e}{\epsilon_0}\int dv f v, \tag{1.62}$$

$$i\omega\tilde{E} = \frac{e}{\epsilon_0}\int dv \tilde{f} v, \tag{1.63}$$

$$\omega^2 = -\frac{\omega_p^2}{n_0}\int dv \frac{v\partial_v f_0}{1 - \frac{kv}{\omega}}. \tag{1.64}$$

In deriving equation (1.63) I have used that the average velocity of the Maxwellian distribution is zero. The integral has a pole at $v = \omega/k$. The principal

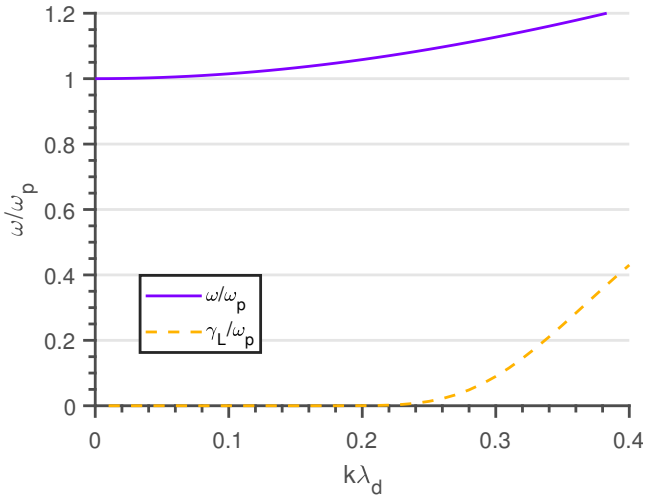


Figure 1.7: The real (solid line) and imaginary (dashed line) parts of the Langmuir plasma wave frequency for a disturbance at late times. Long wavelength Langmuir waves have frequency ω_p and essentially no damping, however if the temperature and therefore λ_D are increased then the wave oscillates at slightly higher frequency and is damped within several cycles.

value of the integral is found by expanding the denominator for small velocities.

$$\omega^2 \simeq -\frac{\omega_p^2}{n}\int dv\left(v+\frac{kv^2}{\omega}+\frac{k^2v^3}{\omega^2}\right)\partial_v f_0, \tag{1.65}$$

$$= \frac{\omega_p^2}{n}\int dv\left(1+2\frac{kv}{\omega}+3\frac{k^2v^2}{\omega^2}\right)f_0, \tag{1.66}$$

$$\omega^2 = \omega_p^2 + 3\omega_p^2\frac{k^2}{\omega^2}\frac{k_BT}{m} \simeq \omega_p^2 + 3k^2\frac{k_BT}{m} = \omega_p^2(1+3k^2\lambda_D^2). \tag{1.67}$$

Equation (1.66) has been derived by integration by parts, where the boundary term is zero since f_0 decays exponentially for high speeds. This dispersion relation agrees with that found in the previous section and shows that $\nu = 3$, corresponding to one degree of freedom for translation in the direction of the wave.

However, the full integral also has an imaginary part which damps the wave. The integration contour is along the real axis but must take a small semi-circle around the pole. The solution at late times will be given by the integral at the slowest decaying pole. The Cauchy theorem gives the imaginary part as the residual at the pole multiplied by $i\pi$. Assuming the damping is much slower than the wave frequency, the real part of the integral is approximately ω_p^2. The full integral is therefore

$$\omega^2 \simeq \omega_p^2 - i\pi\frac{\omega^2\omega_p^2}{k^2n}\frac{\partial f}{\partial v}\Big|_{\omega/k}, \tag{1.68}$$

$$\omega \simeq \omega_p - \frac{i\pi\omega_p^3}{2k^2n}\frac{\partial f}{\partial v}\Big|_{\omega/k}. \tag{1.69}$$

The frequency has been expanded by assuming the imaginary part of ω is much smaller than the real part and taking $\omega \simeq \omega_p$ in the small term. The imaginary part is the damping rate and it is proportional to the derivative of f at the phase velocity of the wave. This agrees with the physical interpretation of Landau damping. For a plasma wave of wavenumber k in a Maxwellian distribution, the damping rate is

$$\gamma_L = \omega_p \sqrt{\frac{\pi}{8}} \frac{1}{k^3 \lambda_D^3} \exp\left(-\frac{1}{2k^2\lambda_D^2}\right). \tag{1.70}$$

The wavelength of a sustained plasma wave must therefore be much greater than the Debye length. Fig. 1.7 shows the damping rate as a function of $k\lambda_D$. It increases when $k\lambda_D \simeq 0.2$, which occurs when the thermal velocity nears the wave phase velocity.

1.7 Particle motion in electromagnetic waves

Now consider the problem of a single isolated charged particle moving under the influence of a plane electromagnetic wave. It is assumed that the particle of charge q is unaffected by the reaction from the radiation it emits itself. Therefore the equation of motion is given exactly by

$$\frac{d(\gamma m \mathbf{v})}{dt} = q(\mathbf{E} + \mathbf{v} \times \mathbf{B}), \tag{1.71}$$

where $\gamma = (1 - v^2/c^2)^{-1/2}$ is the Lorentz factor. This is easily solvable for a non relativistic particle, since in a plane wave the magnitude of the $\mathbf{B}$ field is E/c, so the magnetic term is smaller than the electric term by a factor v/c. The linearly polarised plane wave solution is $\mathbf{E} = E_0 \hat{\mathbf{x}} \sin(\omega t - kx)$, and under a zero order approximation $\gamma = 1$. Integrating the equation of motion gives

$$\frac{d\mathbf{v}}{dt} \simeq \frac{q}{m}\mathbf{E}, \tag{1.72}$$

$$= \frac{q}{m} E_0 \hat{\mathbf{x}} \sin(\omega t - kx), \tag{1.73}$$

$$\mathbf{v} = -\frac{E_0 q}{m\omega} \hat{\mathbf{x}} \cos(\omega t - kx). \tag{1.74}$$

If the particle is an electron, it oscillates transversely with peak velocity $v = a_0 c$ where $a_0 = E_0 e/(m\omega c)$ and E_0 is the maximum electric field during the whole period of the wave. More generally, for all polarisations, the dimensionless amplitude is found from the root mean square of the electric field magnitude,

$$a_0 = \frac{e\sqrt{2\overline{|E(t)|^2}}}{m\omega c}. \tag{1.75}$$

It can be written in terms of the intensity using the Poynting vector $I = \mathbf{E} \times \mathbf{B}/\mu_0$. For all polarisations, the intensity is

$$I = \frac{2\pi^2 \epsilon_0 m^2 c^5}{e^2} \frac{a_0^2}{\lambda^2} = 1.4 \times 10^{18} \frac{a_0^2}{\lambda^2} \mu\mathrm{m}^2\,\mathrm{Wcm}^{-2}.$$

Heavier particles such as protons or ions oscillate at much slower speeds $v = Za_0 cm/m_i$. The electron motion becomes relativistic when $a_0 \simeq 1$ or an intensity of $10^{18}\,\mathrm{Wcm}^{-2}$ for a 1μm wavelength laser. For pulses of this intensity, the magnetic field of the light and relativistic corrections become important [59]. In terms of the electron momentum vector $\mathbf{p}$, the vector potential $\mathbf{A}$ and the Lorentz factor γ, the relativistic equation of motion components are

$$\gamma^2 = 1 + p_i p_i/(m^2c^2), \tag{1.76}$$

$$\frac{d(\gamma mc^2)}{dt} = -ev_j E_j = ev_j \partial_t A_j, \tag{1.77}$$

$$\frac{dp_i}{dt} = -e(E_i + \epsilon_{ijk} v_j B_k) = e(\partial_t A_i - \epsilon_{ijk} v_j \epsilon_{klm} \partial_l A_m). \tag{1.78}$$

The equations are written in the laboratory frame, where the time is t. They can be simplified using the identity $\epsilon_{jki}\epsilon_{ilm} = \delta_{jl}\delta_{km} - \delta_{jm}\delta_{kl}$ and the components of $\mathbf{A}$ for a plane wave. In terms of the ellipticity δ, the components are $A_x = 0$, $A_y = \frac{mca_0}{e}\delta\cos\phi$ and $A_z = \sqrt{1-\delta^2}\frac{mca_0}{e}\sin\phi$. $\delta = 0, 1$ gives a linearly polarised wave and $\delta = \pm 1/\sqrt{2}$ is circularly polarised. $\phi = \omega(t - x/c)$ is the phase for a plane wave travelling in the x direction. Equation (1.78) becomes

$$\frac{dp_i}{dt} = e(\partial_t A_i - \epsilon_{ijk} v_j \epsilon_{klm} \partial_l A_m), \tag{1.79}$$

$$= e(\partial_t A_i - v_j \partial_i A_j + v_j \partial_j A_i) \tag{1.80}$$

$$\frac{dp_x}{dt} = -ev_j \partial_x A_j, \tag{1.81}$$

$$\frac{dp_{y,z}}{dt} = e(\partial_t A_{y,z} + v_j \partial_j A_{y,z}) = e\frac{dA_{y,z}}{dt}. \tag{1.82}$$

On the last line I have used the total derivative of A_y and A_z. The last equation can be integrated to give $p_y = eA_y + \alpha_y$ and $p_z = eA_z + \alpha_z$ where α is an integration constant. The solution with no average transverse momentum is given by $\alpha_y = \alpha_z = 0$. Subtracting c times equation (1.81) from equation (1.77) gives

$$\frac{d(\gamma mc^2 - cp_x)}{dt} = ev_j(\partial_t + c\partial_x)A_j \propto (\partial_t + c\partial_x)\phi = (\partial_t + c\partial_x)(\omega t - \omega x/c) = 0. \tag{1.83}$$

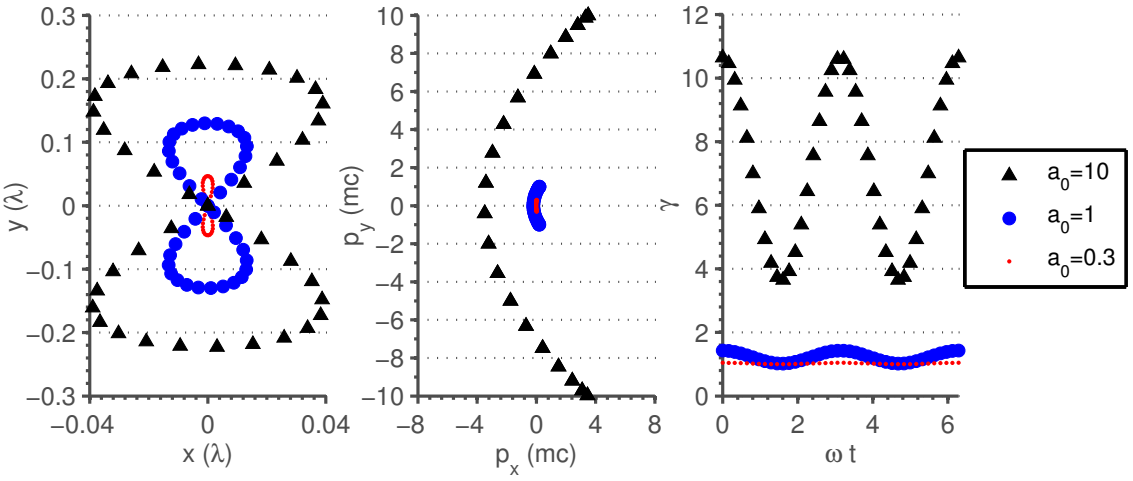


Figure 1.8: The position and momentum of a particle moving in a plane electromagnetic wave propagating along x and linearly polarised along y. The reference frame has zero average motion. The colours show three different values of the dimensionless amplitude $a_0 = eE_0/(mc\omega)$. Each period is shown by 40 points all evenly spaced in time. The panels show (a) The position as a function of time, (b) the momentum as a function of time, (c) the energy as a function of time. For this polarisation, there is no motion in the z direction.

This follows from the chain rule and because $\mathbf{A}$ is proportional to $\cos(\phi)$ or $\sin(\phi)$. Equation (1.83) is integrated to give $p_x = \gamma mc + \alpha_x$. The expression (1.76) is now used to eliminate γ, so that

$$(p_x - \alpha_x)^2 = \gamma^2 m^2 c^2 = m^2 c^2 + p_i p_i = m^2 c^2 + p_x^2 + e^2(A_y^2 + A_z^2), \tag{1.84}$$

$$p_x = \frac{\alpha_x^2 - m^2 c^2 - e^2(A_y^2 + A_z^2)}{2\alpha_x}. \tag{1.85}$$

Here I have also used that $p_{y,z} = eA_{y,z}$. This gives the full general solution in terms of $\mathbf{A}$. For a particle starting from rest, the boundary condition is $p_x = 0$ at $t = 0$. Since $p_x = \gamma mc + \alpha_x$, this gives $\alpha_x = -mc$. The solution is

$$p_x = \frac{mca_0^2}{4}\left(1 + (2\delta^2 - 1)\cos 2\phi\right), \tag{1.86}$$

$$p_y = mca_0\delta\cos\phi, \tag{1.87}$$

$$p_z = mca_0\sqrt{1-\delta^2}\sin\phi. \tag{1.88}$$

If the electron starts at rest, it gains an average longitudinal momentum of $mca_0^2/4$ from the interaction with the wave, irrespective of polarisation. This means a relativistic laser pulse causes an electron to move in the direction of the propagation of the pulse, as well as oscillating transversely.

Using $\gamma mc = p_x - \alpha_x$ gives $\gamma = 1 + \frac{a_0^2}{4}(1 + (2\delta^2 - 1)\cos 2\phi)$, so the average Lorentz factor is $\overline{\gamma} = 1 + a_0^2/4$.

To find the motion in the average rest frame of the particle, one must Lorentz boost to the frame with $\overline{p_x} = 0$. Using the expression for $\mathbf{A}$, equation (1.85) then gives $\alpha_x = -mc\sqrt{1 + a_0^2/2}$. The boost leaves p_y and p_z unchanged, whereas p_x becomes

$$p_x = \frac{mca_0^2}{4\sqrt{1 + \frac{a_0^2}{2}}}(2\delta^2 - 1)\cos 2\phi. \tag{1.89}$$

These can be integrated to find the position as a function of time, which is shown in Fig. 1.8 for several values of a_0. Due to the non-negligible magnetic field of the wave, the particle acquires a longitudinal oscillation and so the orbit is a figure of eight shape. In the ultra-relativistic limit, $y_{max} = \lambda/(\pi\sqrt{2})$ and $x_{max} = \lambda/(8\pi)$. The average Lorentz factor in this reference frame is $\overline{\gamma} = \sqrt{1 + a_0^2/2}$.

1.8 Propagation of light in a plasma

Knowing the response of the plasma electrons to an electromagnetic wave and coulomb collisions, one can calculate the dispersion relation. For a non-relativistic linearly polarised laser pulse ($a_0 < 1$) in a uniform plasma, the linearised equation of motion is $\partial_t u_j = -\frac{e}{m}E_j - u_j/\tau_{ei}$. For a velocity proportional to $e^{i\omega t}$ the solution is

$$\mathbf{u} = -\frac{e\tau_{ei}(1 - i\omega\tau_{ei})\mathbf{E}}{m(1 + \omega^2\tau_{ei}^2)} \simeq \frac{e\mathbf{E}(i\omega\tau_{ei} - 1)}{m\omega^2\tau_{ei}}. \tag{1.90}$$

The current is then $\mathbf{J} = -ne\mathbf{u}$. I have assumed that $\omega\tau_{ei} \gg 1$, giving the behaviour in a weakly collisional plasma with $n\lambda_D^3 > 1$. The ions have an oscillation velocity much smaller than the electrons, so barely contribute to the current.

Using this solution for u, the first fluid equation $\partial_t n = -\partial_i(nu_i) = 0$ shows that in a plane electromagnetic wave with $\partial_y = 0$, the plasma electron density remains uniform. Therefore $\rho = 0$ in the first Maxwell equation (1.9). Taking the curl of the third equation (1.11) and substituting the fourth gives

$$\nabla \times \nabla \times \mathbf{E} = \nabla(\nabla.\mathbf{E}) - \boldsymbol{\nabla^2}\mathbf{E} = -\mu_0\partial_t\mathbf{J} - \epsilon_0\mu_0\partial_t^2\mathbf{E}. \tag{1.91}$$

$$-c^2\boldsymbol{\nabla^2}\mathbf{E} = \frac{\omega_p^2(i\omega\tau_{ei} - 1)}{\omega^2\tau_{ei}}\partial_t\mathbf{E} - \partial_t^2\mathbf{E}. \tag{1.92}$$

Here I have used $\nabla.\mathbf{E} = 0$ and the equation (1.90) to yield a damped wave equation. Substituting the plane wave solution for $\mathbf{E}$ and Taylor expanding the small imaginary part yields

$$\omega^2 = c^2k^2 + \omega_p^2 + i\frac{\omega_p^2}{\omega\tau_{ei}}, \tag{1.93}$$

$$\omega \simeq \sqrt{\omega_p^2 + c^2k^2} + \frac{i\omega_p^2}{2\omega^2\tau_{ei}}. \tag{1.94}$$

The dispersion relation is modified from that in a vacuum, although the vacuum result $\omega = ck$ is recovered when the plasma density goes to zero. Neglecting the damping, the phase and group velocities of the light are

$$v_p = \frac{c}{\sqrt{1 - \frac{\omega_p^2}{\omega^2}}}, \quad v_g = c\sqrt{1 - \frac{\omega_p^2}{\omega^2}}. \tag{1.95}$$

The minimum allowed real frequency is ω_p, showing that an electromagnetic wave must have frequency greater than ω_p to propagate in the plasma. If light propagates into a plasma of increasing density, it eventually encounters a surface where $\omega_p = \omega$, known as the critical surface. The electron density here is the critical density $n_{crit} = \omega^2\epsilon_0 m/e^2$. The wave must be reflected or absorbed at this surface.

Since the plasma frequency depends on the electron mass, it is intensity dependent due to the relativistic mass increase. The plasma frequency is modified to $\omega_p \to \omega_p/\sqrt{\overline{\gamma}}$ where $\overline{\gamma} = \sqrt{1 + a_0^2/2}$ is the average Lorentz factor of the electron in the frame with $\overline{p_x} = 0$.

The small imaginary part of ω gives the collisional damping rate in the plasma. The energy damping rate is

$$\gamma_{coll} = \frac{\omega_p^2}{\omega^2\tau_{ei}} = \frac{\omega_p^3}{\omega^2}\frac{Z\ln(3\pi n\lambda_D^3/Z)}{3\pi n\lambda_D^3}. \tag{1.96}$$

For example, a laser pulse with wavelength $1\,\mu$m in a plasma of density 10^{19}cm^{-3} with a temperature of 100 eV has a damping length of approximately 30 mm.

1.8.1 The ponderomotive force in a finite laser pulse

I have shown that an electron initially at rest is accelerated in the forwards direction by an incoming plane wave, to an average longitudinal momentum $mca_0^2/4$. This is true irrespective of the polarisation. A finite laser pulse, with a gradient in electric field amplitude, therefore accelerates the electron forwards as it overtakes the particle. The movement arises because of the magnetic field of the laser

pulse. In this section, I will show the electron is accelerated by all gradients in the electromagnetic energy density, not just the longitudinal direction.

In an electromagnetic wave, the zero order approximation to the electron motion is the non-relativistic result, equation (1.74). This just gives a sinusoidal oscillation. Using $y(t) = (a_0 c/\omega)\sin\phi$ as an approximation for the electron transverse position, the equation of motion for a pulse with spatially varying amplitude is

$$\frac{dv_y}{dt} = -\frac{e}{m}E_y(\mathbf{x})\sin\phi \simeq -\frac{e}{m}(E_y(0) + y\partial_y E_y)\sin\phi, \tag{1.97}$$

$$= -\frac{e}{m}(E_y(0) + \frac{a_0 c}{\omega}\sin\phi\partial_y E)\sin\phi. \tag{1.98}$$

The first line has been Taylor expanded by assuming the electron oscillates in an almost uniform electric field, that varies only slowly along y. Averaging this acceleration over the full period of the wave from $0 < \phi < 2\pi$ gives

$$\frac{d\overline{v_y}}{dt} = -\frac{ea_0 c}{2m\omega}\partial_y E_y = -\frac{e^2}{4m^2\omega^2}\partial_y(E_y^2) = -\frac{c^2}{4}\partial_y(a_0^2). \tag{1.99}$$

In a plane wave, the electron oscillates about a fixed position. In a finite duration pulse, equation (1.99) shows there is also an average acceleration in the direction of decreasing intensity. This is due to the oscillation taking the electron to regions of differing field strength. Due to other components of the fields that are not part of the plane wave solution, there are also forces in the x and z directions. Including these, the first order approximation to the ponderomotive force is $-mc^2\nabla(a_0^2)/4$.

1.8.2 The filamentation instability

Since the plasma frequency depends on the electron mass, the propagation of light in a plasma is altered if it approaches relativistic intensity. The plasma frequency decreases in a plasma with relativistic electrons, because their effective mass increases to γm. This means the effective critical density increases, an effect known as relativistic transparency.

One of the first noticeable relativistic effects is that a pulse starts to focus in a plasma, if its power is high enough. This is because the phase velocity v_p increases with ω_p, and so the phase velocity in the higher intensity regions of a pulse is less than the lower intensity regions. Since the highest intensity region is usually the centre of the beam, this causes the pulse to focus and the intensity to increase further, leading to a parametric instability. If the pulse power exceeds

a certain critical threshold, this phase velocity difference is enough to exceed the usual divergence of a collimated Gaussian beam. The threshold [57] is

$$P_{fil} = 17 \, \frac{n_{crit}}{n} \, \mathrm{GW}. \tag{1.100}$$

If a beam exceeds this critical power, it focusses into filaments, each containing less than the threshold power. The instability grows from noise, so there must be a range of filament sizes, intensities and positions. In addition, the laser wavefront is altered transversely, meaning the phase and intensity vary as a function of transverse position and the coherence is lost.

There are two other processes that cause filamentation. The ponderomotive force tends to expel electrons from any high intensity region of the laser, meaning the local electron density and plasma frequency is less. This leads to the same focussing and feedback mechanism. High intensity regions also have greater collisional heating of the plasma, causing a thermal expansion of both the electrons and ions. This decreases the local electron density and leads to the same instability. One calculation of the ponderomotive growth rate [81] gives

$$\gamma_{fil} = \frac{\omega_p^2}{8\omega} a_0^2. \tag{1.101}$$

Chapter 2

Theory of the Raman Instability and Raman Amplifiers

2.1 The Raman instability

Plasmas can support many instabilities that grow exponentially. Chapter one showed that the plasma has an oscillation mode at frequency ω_p, so an external driver may resonantly drive a plasma wave under certain conditions. Even neglecting collisions, plasma waves are unstable and will decay via Landau damping. Their energy is transferred to fast electrons propagating at the phase velocity of the wave. Electron-ion collisions were also found to damp the wave, although in plasmas with $n\lambda_D^3 \gg 1$ this effect is minimal. A modified dispersion relation was found for electromagnetic waves in a plasma.

There are other collisionless plasma instabilities which grow exponentially and on a similar timescale to ω_p. These have the effect of modifying the amplitude and wave-front of any light propagating through a plasma. The fastest growing of these is the Raman instability, whereby the light will couple to a plasma wave.

When $\omega = \omega_p$ (meaning $n = n_{crit}$), the light will be resonant with a plasma wave. This occurs at the critical surface, a process known as resonance absorption. The Raman instability is much more ubiquitous because it is resonant at all plasma densities below $n_{crit}/4$. The beating of two separate electromagnetic waves generates an oscillation at their frequency difference. If this frequency difference is ω_p, the beat frequency resonantly excites a plasma wave. If there is a high amplitude monochromatic pump wave in a broad-band radiation field, there will always be some seed radiation at the correct frequency to drive the plasma at ω_p.

Each electromagnetic wave may propagate in any direction [57, 25]. The beating light waves will create a ponderomotive potential that forces the electrons in to

localised regions on a wavelength scale, resonantly driving an electron plasma wave. The plasma wave acts as a moving Bragg grating and scatters the incoming light. This increases the amplitude of the scattered light and therefore the ponderomotive potential, leading to an exponentially growing instability. For the exponential growth to continue, the resonance must be maintained.

The Raman instability has been observed to impede inertial confinement fusion experiments [43, 117, 50], although artificially disrupting the transverse coherence of the laser pulses mitigates this somewhat [91].

2.1.1 The matching conditions

Both the high amplitude pump electromagnetic wave and the infinitesimal scattered wave must obey the plasma dispersion relations. Using subscript 0 for the pump, 1 for the scattered wave and 2 for the plasma wave, this gives $\omega_0^2 = \omega_p^2 + c^2k_0^2$ and $\omega_1^2 = \omega_p^2 + c^2k_1^2$. The maximal exponential growth rate will be when the driven plasma wave is resonant, so its frequency is given by the dispersion relation $\omega_2^2 = \omega_p^2(1 + 3k_2^2\lambda_D^2) \simeq \omega_p^2$.

In the Raman interaction each photon must decay into a plasmon and another photon. To conserve energy and momentum, this requires $\hbar\omega_0 = \hbar\omega_1 + \hbar\omega_p$ and $\hbar\mathbf{k_0} = \hbar\mathbf{k_1} + \hbar\mathbf{k_2}$. This means the fraction of energy transferred to the scattered wave is $\hbar\omega_1/(\hbar\omega_0) = 1 - \omega_p/\omega_0$ [79]. The remaining energy is transferred to the plasma wave, which eventually damps and produces high energy electrons and heating of the plasma.

For direct backscatter, the pump and scattered wave travel in the opposite direction. To allow a closed solution to the equations, approximate the plasma wave frequency as $\omega_2 \simeq \omega_p$. In this situation, there is a system of 5 equations, consisting of three dispersion relations and the energy and momentum matching conditions. The unique solution is [40]

$$k_1 = -\frac{1}{c}\sqrt{\omega_0^2 - 2\omega_0\omega_p}, \tag{2.1}$$

$$k_2 = k_0 + \frac{1}{c}\sqrt{\omega_0^2 - 2\omega_0\omega_p} \simeq (2\omega_0 - \omega_p)/c. \tag{2.2}$$

The scattered light is only down shifted in frequency because the up-shifted peak at $(\omega_0 + \omega_p, k_0 + k_p)$ is very far from the allowed curve for light propagating in a plasma, as shown in Fig. 2.1. The expansion is valid for highly underdense plasmas

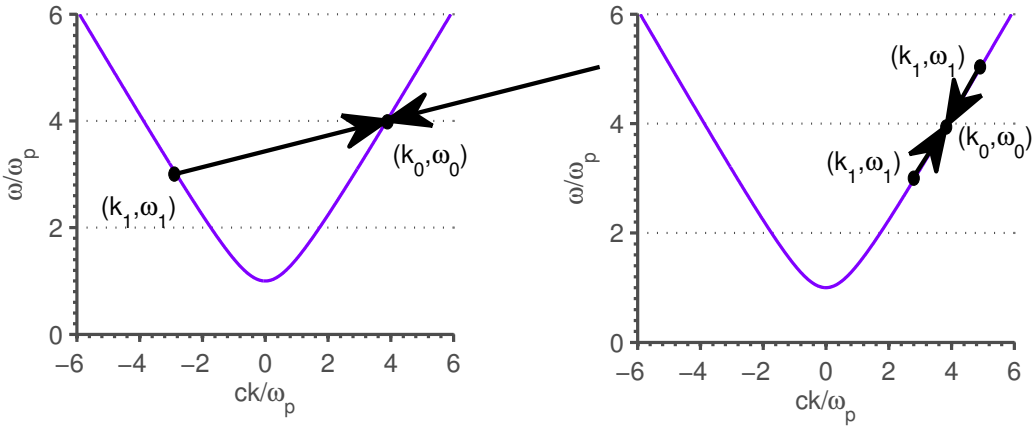


Figure 2.1: The dispersion relation for light (purple) in a plasma and the Raman matching conditions. Panel (a) shows Raman backward scattering for the case $\omega_0/\omega_p = 4$. The scattered light is downshifted in frequency and propagates in the reverse direction. The black arrows represent the k_2 wave-vector of the plasma wave. A resonant plasma wave must have frequency ω_p so the fastest growing mode has a unique frequency and wavenumber for the backscattered light. The up-shifted component is far from the dispersion relation and is not resonant. Panel (b) shows the case for forward scattering, which has two available resonant modes. The scattered light will propagate in the same direction as the pump beam. In general, the matching conditions can be met for scattering at any angle.

where $\omega_p \ll \omega_0$. The solution for k_1 is only real when $\omega_0 > 2\omega_p$, limiting Raman scattering to plasma densities $n < n_{crit}/4$.

The plasma wave-vector is approximately twice that of the pump pulse, meaning the wavelength of the plasma wave will be half that of the incoming light. This agrees with the physical picture, in that the beat wave set up by the pump and scattered light will have a ponderomotive potential proportional to the square of the electric field. This has a wavelength approximately half that of the pump light.

There is also a unique solution when the scattered light is directly forward. In this case both the frequency up shifted and down shifted components are close to the dispersion relation, so both modes are excited. The up shifted solution in this case is

$$k_1 = \frac{1}{c}\sqrt{\omega_0^2 + 2\omega_0\omega_p}, \tag{2.3}$$

$$k_2 = k_0 - \frac{1}{c}\sqrt{\omega_0^2 + 2\omega_0\omega_p} \simeq -\omega_p/c. \tag{2.4}$$

The down shifted solution is

$$k_1 = \frac{1}{c}\sqrt{\omega_0^2 - 2\omega_0\omega_p}, \tag{2.5}$$

$$k_2 = k_0 - \frac{1}{c}\sqrt{\omega_0^2 - 2\omega_0\omega_p} \simeq \omega_p/c. \tag{2.6}$$

2.1.2 Derivation of the coupled envelope equations

To calculate the exponential growth rate of the instability and its development in the non-linear stage, one must find the coupled equations linking the pump, seed and plasma waves. These are most conveniently found in terms of the vector potential. Choice of the Coulomb gauge means that $\nabla.\mathbf{A} = 0$. The light and the plane plasma wave are aligned along the x axis such that $\partial_y = \partial_z = 0$. The Coulomb gauge then reduces to $\partial_x A_x = 0$, a constant potential, taking $A_x = 0$ without loss of generality.

Suppose there is a large amplitude electromagnetic pump wave propagating through the plasma. Equation (1.12) becomes

$$\nabla \times \mathbf{B} = \nabla \times \nabla \times \mathbf{A} = -\nabla^{\mathbf{2}}\mathbf{A} = \mu_0\mathbf{J} - \mu_0\epsilon_0\partial_t\nabla\phi - \mu_0\epsilon_0\partial_t^2\mathbf{A}. \tag{2.7}$$

Since $\partial_y = \partial_z = 0$, the transverse part is

$$(\partial_t^2 - c^2\nabla^2)A_{y,z} = \mu_0 c^2 J_{y,z}. \tag{2.8}$$

Assuming that the electron motion is non-relativistic, equation (1.81) gives the current as $J_{y,z} = -neu_{y,z} = -ne^2A_{y,z}/m$.

One can linearise the equation to be the sum of the large pump wave with subscript 0 and a small scattered wave with subscript 1. This is the sum of two transverse plane electromagnetic waves, so $\mathbf{A_0}$ and $\mathbf{A_1}$ are transverse. Writing the electron density as $n = n_0 + \tilde{n}$, where the tilde denotes small quantities, equation (2.8) becomes

$$(\partial_t^2 - c^2\nabla^2)(\mathbf{A_0} + \mathbf{A_1}) = -\frac{\mu_0 ne^2c^2}{m}\mathbf{A} = -\frac{e^2}{\epsilon_0 m}(n_0 + \tilde{n})(\mathbf{A_0} + \mathbf{A_1}), \tag{2.9}$$

$$= -\omega_p^2(\mathbf{A_0} + \mathbf{A_1}) - \frac{e^2}{\epsilon_0 m}\tilde{n}(\mathbf{A_0} + \mathbf{A_1}), \tag{2.10}$$

$$(\partial_t^2 - c^2\nabla^2 + \omega_p^2)(\mathbf{A_0} + \mathbf{A_1}) = -\omega_p^2\frac{\tilde{n}}{n_0}(\mathbf{A_0} + \mathbf{A_1}). \tag{2.11}$$

The pump wave obeys the plasma dispersion relation and the term that is quadratic in the small quantities can be neglected, giving

$$(\partial_t^2 - c^2\nabla^2 + \omega_p^2)\mathbf{A_1} = -\omega_p^2 \frac{\tilde{n}}{n_0}\mathbf{A_0}. \tag{2.12}$$

In the absence of a pump, the seed wave obeys the normal dispersion relation and is constant amplitude. However the presence of a pump wave and a plasma wave gives a driving term on the right hand side, increasing the energy of the seed. The increase in $\mathbf{A_1}$ is in the same direction as $\mathbf{A_0}$ so the scattered light will gain the same polarisation as the pump. One study showed that even exotic polarisation states can be effectively amplified [84, 128].

To proceed further, one must know the amplitude $\tilde{n}$ of the plasma wave, by linearising the fluid equations. I assume that the ions are static and uniformly distributed. The electron velocity is treated as the large transverse oscillation in the pump beam, plus a small longitudinal component $\tilde{\mathbf{u}}$ for the plasma wave. This means $\mathbf{u} = e\mathbf{A}/m + \tilde{\mathbf{u}}$, where the large transverse motion is $\mathbf{u} = e\mathbf{A}/m$, as given by equation (1.81). This is true only for a non-relativistic pump pulse.

The first fluid equation (1.38) becomes

$$\partial_t(n_0 + \tilde{n}) + \nabla.((n_0 + \tilde{n})(e\mathbf{A}/m + \tilde{\mathbf{u}})) = 0, \tag{2.13}$$

$$\partial_t\tilde{n} + \frac{e}{m}(n_0 + \tilde{n})\nabla.\mathbf{A} + \frac{e}{m}\mathbf{A}.\nabla\tilde{n} + n_0\nabla.\tilde{\mathbf{u}} + \nabla.(\tilde{n}\tilde{\mathbf{u}}) = 0, \tag{2.14}$$

$$\partial_t\tilde{n} + n_0\nabla.\tilde{\mathbf{u}} = 0. \tag{2.15}$$

On the last line I have used the Coulomb gauge and neglected the term that is quadratic in the small quantities. In addition, I have used $\mathbf{A}.\nabla\tilde{n} = 0$ since $\mathbf{A}$ is transverse and the plasma wave is longitudinal. Writing the fields in terms of the potentials and using the linearised form of $\mathbf{u}$ to replace $\mathbf{A}$, the second fluid equation (1.42) becomes

$$\partial_t\mathbf{u} + \mathbf{u}.\nabla\mathbf{u} = -\frac{e}{m}(-\nabla\phi - \partial_t\mathbf{A} + \mathbf{u}\times\nabla\times\mathbf{A}) - \frac{\nabla p}{nm}, \tag{2.16}$$

$$\frac{e}{m}\partial_t\mathbf{A} + \partial_t\tilde{\mathbf{u}} + \mathbf{u}.\nabla\mathbf{u} = \frac{e}{m}(\nabla\phi + \partial_t\mathbf{A}) - \mathbf{u}\times\nabla\times\mathbf{u} + \mathbf{u}\times\nabla\times\tilde{\mathbf{u}} - \frac{\nabla p}{nm}. \tag{2.17}$$

Next I use the vector identity $\mathbf{u}\times\nabla\times\mathbf{u} = \frac{1}{2}\nabla(\mathbf{u}.\mathbf{u}) - \mathbf{u}.\nabla\mathbf{u}$ and the adiabatic equation of state $p \propto n^\nu$. This gives $\nabla p = \nu\frac{p}{n}\nabla n = \nu k_B T\nabla\tilde{n}$ for an ideal gas. The velocity perturbation $\tilde{\mathbf{u}}$ is longitudinal and the plasma wave is assumed to be a plane wave, giving $\nabla\times\tilde{\mathbf{u}} = 0$. This simplifies the equation to

$$\partial_t\tilde{\mathbf{u}} = \frac{e}{m}(\nabla\phi) - \frac{1}{2}\nabla(\mathbf{u}.\mathbf{u}) - \frac{\nu k_B T\nabla\tilde{n}}{n_0 m}. \tag{2.18}$$

The velocity can be eliminated using the linearised form and again dropping the terms quadratic in the small quantities, giving

$$\mathbf{u}.\mathbf{u} = \left(\frac{e}{m}(\mathbf{A_0} + \mathbf{A_1}) + \tilde{\mathbf{u}}\right)^2, \tag{2.19}$$

$$\simeq \frac{e^2}{m^2}(\mathbf{A_0} + \mathbf{A_1})^2 + 2\frac{e}{m}(\mathbf{A_0} + \mathbf{A_1}).\tilde{\mathbf{u}} \simeq \frac{e^2}{m^2}(\mathbf{A_0}^2 + 2\mathbf{A_0}.\mathbf{A_1}) + 2\frac{e}{m}\mathbf{A_0}.\tilde{\mathbf{u}}. \tag{2.20}$$

The pump wave is assumed to be a transverse plane wave and the plasma wave is longitudinal, meaning $\mathbf{A_0}.\tilde{\mathbf{u}} = 0$. Substituting this into equation (2.18) gives

$$\partial_t \tilde{\mathbf{u}} = \frac{e}{m}\nabla\phi - \frac{e^2}{2m^2}\nabla(\mathbf{A_0^2} + 2\mathbf{A_0}.\mathbf{A_1}) - \frac{\nu k_B T \nabla\tilde{n}}{n_0 m}. \tag{2.21}$$

The zero order term $\mathbf{A_0}^2$ leads to a fast oscillation with frequency $2\omega_0$, which will not be resonant with a plasma wave. The scalar potential can be eliminated with Poisson's equation $\nabla^2\phi = \tilde{n}e/\epsilon_0$. Taking the time derivative of equation (2.15), the divergence of equation (2.21) and substituting leads to

$$\left(\partial_t^2 - \nu\frac{k_B T}{m}\nabla^2 + \omega_p^2\right)\tilde{n} = \frac{n_0 e^2}{m^2}\nabla^2(\mathbf{A_0}.\mathbf{A_1}). \tag{2.22}$$

The coupled equations (2.12) and (2.22) show that the pump, seed and plasma waves all obey their respective dispersion relations, but also include a driving term that couples the energy from the pump pulse to the seed and plasma waves. The plasma wave driving term is proportional to $\mathbf{A_0}.\mathbf{A_1}$, meaning the growth rate is maximal if the polarisations are aligned. If the pump and seed beams have orthogonal polarisation then there is no instability.

In the non-linear stage the seed wave amplitude is no longer negligible and the pump becomes depleted. The system can be generalised to the non-linear case by including the equivalent equation for depletion of the pump,

$$(\partial_t^2 - c^2\nabla^2 + \omega_p^2)\mathbf{A_0} = -\omega_p^2\frac{\tilde{n}}{n_0}\mathbf{A_1}. \tag{2.23}$$

Now all that remains is to numerically solve these wave equations. However, analytic solutions exists under certain assumptions. The first is the linear case when the plasma and seed waves are infinitesimal. The second is the highly non-linear case for backscattered light, when the pump wave is fully depleted. These solutions can by found by assuming the frequency and wave-vector matching conditions and using an ansatz for each wave with dimensionless amplitudes a_0, a_1 and n_2. These are $\mathbf{A_{0,1}} = \frac{mc}{2e}\hat{\mathbf{y}}a_{0,1}(x,t)\left[e^{i(k_{0,1}x - \omega_{0,1}t)} + c.c.\right]$ and $\frac{\tilde{n}}{n_0} = \frac{n_2}{2}\left[ie^{i(k_2 x - \omega_2 t)} + c.c.\right]$. This

ensures that the quantities are real and $a_{0,1}$ are the same as defined earlier. The polarisations are linear and aligned. The modulus gives the envelope of the wave as a function of position and time.

I now substitute the ansatz into the first coupled equation (2.12) and use the product rule to evaluate the second derivatives. Since $\omega_0 = \omega_1 + \omega_2$ only one of the four terms will be resonant with an exponent of zero, whereas the others will have a remaining fast oscillation that will average to zero. The remaining terms are

$$\frac{mc}{2e}\left[\partial_t^2 - 2i\omega_1\partial_t - \omega_1^2 - c^2\partial_x^2 - 2ic^2k_1\partial_x + c^2k_0^2 + \omega_p^2\right]a_1 = i\omega_p^2\frac{mc}{4e}e^{i(-k_1-k_2+k_0)x+i(\omega_1+\omega_2-\omega_0)t}n_2a_0. \tag{2.24}$$

The envelope of the laser is given by a_1 and this will vary much more slowly than the carrier wave, assuming the pulse is much longer than one wavelength. Therefore $\partial_t \ll \omega$ and $\partial_x \ll k$. This means the second derivatives can be neglected with respect to the first derivatives. Using the dispersion relations and the matching conditions for all waves, one finds for the perfectly resonant interaction that

$$(-2i\omega_1\partial_t - 2ic^2k_1\partial_x)a_1 = i\frac{\omega_p^2}{2}n_2a_0, \tag{2.25}$$

$$(\partial_t + \frac{c^2k_1}{\omega_1}\partial_x)a_1 = -\frac{\omega_p^2}{4\omega_1}n_2a_0. \tag{2.26}$$

Next, substitute the ansatz into the equation for the plasma wave (2.22). In doing so, one can neglect the non-resonant terms, assume the plasma wave dispersion relation and keep only the lowest order derivatives to obtain

$$\left(\omega_2\partial_t + \frac{\nu k_BT}{m}k_2\partial_x\right)n_2 = -\frac{1}{4}c^2k_2^2a_0a_1. \tag{2.27}$$

To find a simple analytic solution, I now make several assumptions to simplify the equations. These are now summarised along with assumptions made already.

- The plasma wave phase velocity is much greater than thermal electron velocity so that Landau damping is neglected and $\omega_2 \simeq \omega_p$
- Neither of the pulses have relativistic intensity i.e. $|a_0| < 1$.
- There is no damping of any of the waves.
- The interaction is much longer than one wavelength, so that the pulse envelopes vary slowly.

These assumptions cover a wide range of applicability. Writing the group velocities of the pump and seed as v_0 and v_1, the three coupled equations reduce to

$$(\partial_t + v_0\partial_x)a_0 = \frac{1}{4}\frac{\omega_p^2}{\omega_0}n_2 a_1, \tag{2.28}$$

$$(\partial_t + v_1\partial_x)a_1 = -\frac{1}{4}\frac{\omega_p^2}{\omega_1}n_2 a_0, \tag{2.29}$$

$$\partial_t n_2 = -\frac{c^2k_2^2}{4\omega_p}a_0 a_1. \tag{2.30}$$

The three waves all obey a first order wave equation. Each wave has a driving term which is a product of the other two waves. These driving terms also depend on the plasma density. Since the amplitudes $a_{0,1}$ and n_2 can be positive or negative, the amplitude of the seed electric field envelope is therefore $m\omega_1 c|a_1|/e$ and the intensity is proportional to $\omega_1^2|a_1|^2$. The plasma wave amplitude is $\tilde{n}/n_0 = |n_2|$.

2.1.3 Linear growth rates

The growth rate in the linear regime can now found by assuming the seed and plasma wave have infinitesimal amplitude but a finite duration. This is equivalent to the condition that the timescale of growth is much shorter than the seed pulse duration. This means the spatial derivative, for constant time, in equation (2.29) can be neglected. Since their amplitude is infinitesimal, the pump will not be depleted and so a_0 is constant. Taking the time derivative of equation (2.29) and using equation (2.30) gives

$$\partial_t^2 a_1 = -\frac{\omega_p^2}{4\omega_1}a_0\partial_t n_2 = -\frac{\omega_p^2}{4\omega_1}a_0\left(-\frac{c^2k_2^2}{4\omega_p}\right)a_0 a_1 = \gamma_R^2 a_1, \tag{2.31}$$

$$\gamma_R = \frac{1}{4}a_0\frac{\omega_p c|\mathbf{k_2}|}{\sqrt{\omega_1\omega_p}.} \tag{2.32}$$

The scattered light amplitude grows exponentially with growth rate γ_R. The growth rate can be evaluated for Raman scatter in any direction where $\mathbf{k_2} = \mathbf{k_0} - \mathbf{k_1}$ is given by the matching condition. In the special case that the two waves propagate in the same direction, known as Raman forward scatter, equation (2.6) gives $|\mathbf{k_2}| = \omega_p/c$ and the growth rate becomes

$$\gamma_{RFS} = \frac{a_0\omega_p^2}{4\sqrt{\omega_p(\omega_0 - \omega_p)}} \simeq \frac{a_0\omega_p}{4}\sqrt{\frac{\omega_p}{\omega_0}}. \tag{2.33}$$

The simplification is valid for very underdense plasma $\omega_p \ll \omega_0$. In the case that the beams counter align, $ck_2 = ck_0 + \sqrt{\omega_0^2 - 2\omega_0\omega_p}$ is maximal and so the growth rate maximises. This is known as Raman backward scatter, with a growth rate

$$\gamma_{RBS} = \frac{a_0\omega_p}{4}\frac{ck_0 + \sqrt{\omega_0^2 - 2\omega_p\omega_0}}{\sqrt{\omega_p(\omega_0 - \omega_p)}} \simeq \frac{a_0\omega_p}{4}\frac{2\omega_0 - \omega_p}{\sqrt{\omega_p(\omega_0 - \omega_p)}}, \tag{2.34}$$

$$\gamma_{RBS} \simeq \frac{1}{2}a_0\sqrt{\omega_0\omega_p}. \tag{2.35}$$

The expansion is valid for very underdense plasmas. The ratio of forward to backward scattering is $\omega_p/(2\omega_0)$, so forward scatter is always slower growing than backwards. This is somewhat offset by the fact that forward scattered light co-propagates with the pump pulse, whereas backward scattered light quickly exits the short pump pulse. This gives it more time to grow. Forward scattered light is relatively more important in denser plasmas. The growth rate of side-scattered light will be intermediate between the forward and backward rates.

It is important to note that the Raman growth rates are of the order ω_p and increase linearly with the pump pulse amplitude. For example, for a $1\,\mu$m pump pulse at $10^{14}\,\mathrm{Wcm}^{-2}$ in a plasma of density $10^{19}\mathrm{cm}^{-3}$, gives $\gamma_{RBS}^{-1} \simeq 300\,\mathrm{fs}$. This is very quick and means any moderately intense laser crossing an underdense plasma of several millimetres will have many Raman growth times. The backscattered light will continue to grow until it reaches the rear of the pump pulse. The growth rate will be decreased by damping of the plasma wave, such that there is a threshold pump amplitude for the instability to occur.

To investigate backwards scattering, it is often convenient to write the equations in spatial coordinates that move with the seed pulse, since all simulations presented later will use this coordinate system. The coordinate system is arranged so that the seed moves towards negative x and the pump propagates towards positive x. Also inserting the matching conditions and approximating $ck_2 = 2\omega_0 - \omega_p$, the three coupled equations become

$$\left(\partial_t + (v_0 + v_1)\partial_x\right) a_0 = \frac{\omega_p^2}{4\omega_0} n_2 a_1, \tag{2.36}$$

$$\partial_t a_1 = -\frac{\omega_p^2}{4(\omega_0 - \omega_p)} n_2 a_0, \tag{2.37}$$

$$\left(\partial_t + v_1\partial_x\right) n_2 = -\frac{(2\omega_0 - \omega_p)^2}{4\omega_p} a_0 a_1. \tag{2.38}$$

2.2 The non-linear stage

The scattered radiation cannot continue growing exponentially for an indefinite time as this would violate the conservation of energy. Eventually the scattered light (seed pulse) will reach the stage where it starts to deplete the pump pulse and so the growth of the seed slows below exponential [39, 99]. In the limit that the seed is very strong, the pump will be fully depleted. A schematic of this situation is shown in Fig. 2.2.

In this limit, the pump is fully depleted by a short seed pulse and so the duration of the seed is short compared to the timescale of evolution [73, 74, 37, 112, 88]. One can look for these steady state solutions by neglecting the time derivatives with respect to spatial derivatives to find

$$(v_0 + v_1)\partial_x a_0 = \frac{\omega_p^2}{4\omega_0} n_2 a_1, \tag{2.39}$$

$$\partial_t a_1 = -\frac{\omega_p^2}{4(\omega_0 - \omega_p)} n_2 a_0, \tag{2.40}$$

$$v_1 \partial_x n_2 = -\frac{(2\omega_0 - \omega_p)^2}{4\omega_p} a_0 a_1. \tag{2.41}$$

In typical situations, the plasma is very underdense so that $v_0 \simeq v_1 \simeq c$ and $\omega_p \ll \omega_0$. The equations yield the seed pulse requirements for it to fully deplete the pump pulse. This can be estimated by assuming a seed pulse initially with a top-hat shaped temporal profile of duration T_1 and amplitude a_1, such that $\partial_x a_1 = 0$ within the pulse. Taking the spatial derivative of equation (2.39) and substituting equation (2.41) gives

$$0 = c^2 \partial_x^2 a_0 + \frac{\omega_0 \omega_p a_1^2}{8} a_0. \tag{2.42}$$

This is a simple harmonic oscillator equation, so energy oscillates between the three waves. It has the simple solution

$$a_0(x) = a_0(0) \cos\left(\frac{\sqrt{\omega_0 \omega_p} a_1}{2\sqrt{2} c} x \right). \tag{2.43}$$

The pump will be fully depleted at its first zero at $a_0(0)\cos(\pi/2)$. Therefore the seed pulse duration required to achieve full pump depletion is

$$\sqrt{\omega_0 \omega_p} a_1 T_1 = \pi\sqrt{2} \simeq 5. \tag{2.44}$$

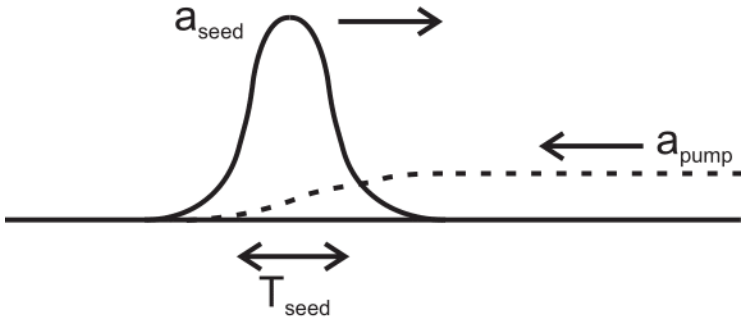


Figure 2.2: A schematic for the non-linear stage of backward Raman scattering. The scattered pulse is short and fully depletes the pump pulse (dashed) as it counter-propagates with it. The amplitude of the scattered light may exceed that of the pump.

This is a very important result and shows that the condition for full pump depletion is a condition on the scattered pulse only. The condition depends only weakly on the plasma density. Fig. 2.3 shows the condition for a $1\,\mu$m wavelength laser at several different plasma densities. It is mainly dependent on the amplitude times the duration of the seed. This is proportional to the square root of the intensity times the duration. A more intense seed will stimulate a greater amplitude plasma wave and deplete the pump more quickly. When the pump is depleted, the energy transfer stops, so any seed light behind this point is not amplified. As the seed intensity increases, its duration therefore shortens.

Squaring the equation (2.44) gives an equivalent condition for the fluence times the duration of the pulse, both shown in Fig. 2.3.

2.2.1 The self similar solution

This section will show that once the pulse grows to meet the condition (2.44), the non-linear stage starts and this equation continues to hold indefinitely. This means that as the amplitude increases in the non-linear stage, the amplified pulse must lie on the line shown in Fig. 2.3a, so the output of any Raman amplifier will lie along this line. This is true until one of the assumptions breaks down [73, 75]. The eventual limitation is that the seed pulse reaches relativistic intensity and the mass increase of the electron means the phase matching between the three waves is lost.

First, take the three coupled equations and again neglect the time derivatives

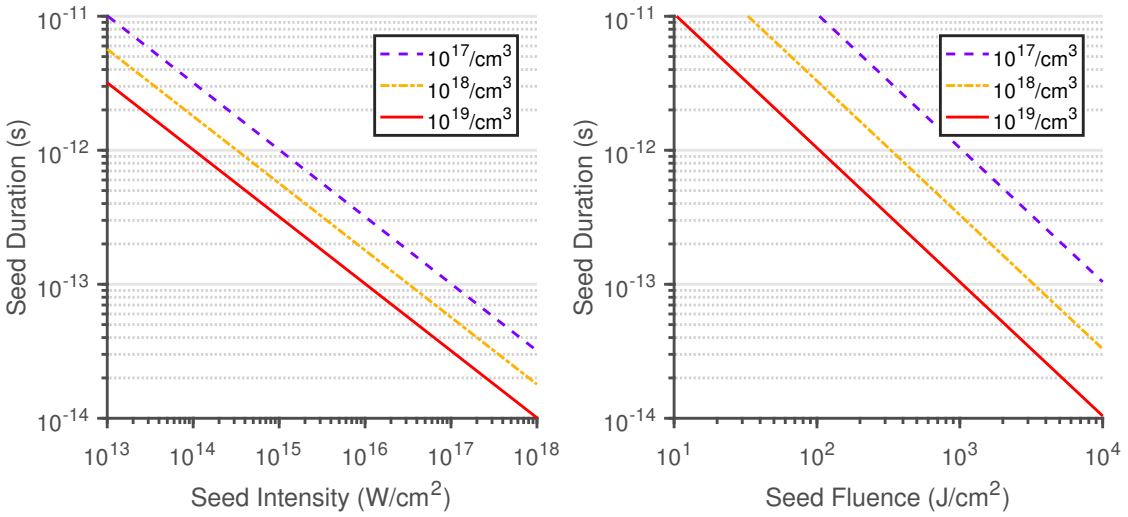


Figure 2.3: The duration that the amplified pulse must have in the non-linear stage, as a function of (panel a) its intensity or (panel b) its fluence. This is shown for a 1 μm wavelength pump in several different plasma densities.

and the dispersion of the laser pulses, such that $v_0 \simeq v_1 \simeq c$. For a function $u(x,t)$, use the trial solutions

$$a_0(x,t) = \alpha_0 \cos(u/2), \tag{2.45}$$

$$a_1(x,t) = \alpha_1 c \partial_x u, \tag{2.46}$$

$$n_2(x,t) = \alpha_2 \sin(u/2). \tag{2.47}$$

The equations (2.39) and (2.41) are satisfied if $\alpha_2 = \alpha_0(2\omega_0/\omega_p)^{3/2}$ and $\alpha_1 = -\sqrt{2/(\omega_0\omega_p)}$. Equation (2.40) then gives

$$c\partial_t\partial_x u = \gamma_{RBS}^2 \sin(u). \tag{2.48}$$

Inspired by the equation (2.44), the trial solutions $u(x,t) = f(\beta)$ should be self-similar in the coordinate xt, or the dimensionless coordinate $\beta = xt\gamma_{RBS}^2/c$. This means the solution for all values of x and t is a function of xt only. Changing variables in the differential equation leads to

$$\beta \frac{d^2 f}{d\beta^2} + \frac{df}{d\beta} = \sin(f). \tag{2.49}$$

In this equation, x and t have been eliminated by the change of variables, so the amplitude of the three waves depends only on a single dimensionless parameter $\beta = \gamma_{RBS}^2 xt/c$. The numerical integration of this equation is shown in Fig. 2.4 for $\beta > 0$. The stationary solution is $f = n\pi$ where n is an integer, however $f = 0$ is an unstable equilibrium corresponding to exactly zero seed pulse. The solution

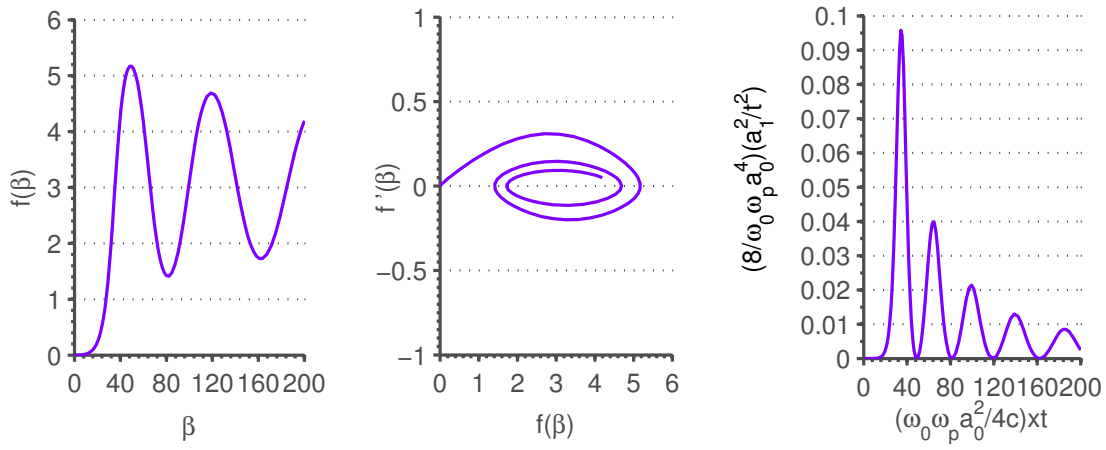


Figure 2.4: The numerical solution of the self similar equation (2.49) converges to $f = \pi$. (a) shows the solution as a function of the self similar coordinate $\beta = xt\gamma_{RBS}^2/c$. (b) shows the convergence to π as a function of $f(\beta)$ and $f'(\beta)$. (c) Gives the shape of the amplified pulse power envelope, as found from this numerical solution. It propagates towards negative x.

therefore converges to $f = \pi$ for any initial seed pulse. To find a_1 in terms of f I use the chain rule to give

$$\partial_x u = \partial_x f = (\partial_\beta f)(\partial_x \beta) = (\partial_\beta f)\gamma_{RBS}^2 t/c. \tag{2.50}$$

Inserting this into the expression (2.46) yields the seed amplitude

$$a_1 = -c\sqrt{\frac{2}{\omega_0\omega_p}}\partial_x f = -\sqrt{\frac{2}{\omega_0\omega_p}}\gamma_{RBS}^2 t\frac{df}{d\beta}. \tag{2.51}$$

Equation (2.51) taken at constant t leads to the relation $-\sqrt{\omega_0\omega_p}\int a_1 dx = c\sqrt{2}\Delta u = c\pi\sqrt{2}$. This shows the integrated amplitude of the amplified pulse is always constant. It is equivalent to the condition plotted in Fig. 2.3. Since the $u = \pi$ stationary point is stable and an attractor, all seed pulses will converge towards this solution for the non-linear stage. Using the convergent value of $u = \pi$ in equation (2.45) shows that the pump becomes fully depleted behind the seed pulse, as expected.

The power density of the amplified pulse, given by the square of equation (2.51), is plotted in Fig. 2.4c. It is not a single spike; it is composed of one main peak along with many trailing sub-peaks. This is a characteristic shape for pulses amplified by the Raman instability and occurs because energy oscillates between the pump, seed and plasma waves.

The analysis shows that a pump pulse may be fully depleted by the interaction. Furthermore, the seed pulse is amplified in the non-linear stage because equation

(2.51) shows its peak amplitude $a_1 \propto t$. Equation (2.44) then means that the seed duration shortens $T_1 \propto 1/t$, the intensity $I_1 \propto t^2$ and the fluence $a_1^2 T_1 \propto t$. This is in accordance with the conservation of energy, since the pump intensity is constant [69, 10].

Although this analysis is very promising for amplification, there are several physical effects which violate the assumptions made and potentially reduce the efficiency and break the coherence of the amplified pulse.

Effects violating the non-linear model

- For high amplitude plasma waves, a significant fraction of the electrons could exceed the phase velocity, causing the wave to break
- Landau damping of the plasma wave
- Collisional damping of the laser pulses and plasma wave
- The frequency matching criterion is not exact
- Light is scattered by an ion plasma wave [127]
- Light may be Raman scattered in other directions and not contribute
- The coherence time of the pump may be less than the duration of the seed, causing loss of resonance.
- Any transverse incoherence of the pump will cause spatially varying gain, leading to filamentation.
- The pump spontaneously back-scatters from thermal plasma noise, depleting it before it reaches the seed.
- The amplified pulse may reach relativistic amplitude, where the increased electron mass will detune the resonance
- Forward Raman scattered light of the seed will propagate with it and modulate the envelope
- The amplified pulse could suffer transverse ponderomotive filamentation
- The amplified pulse could drive a wake-field with its ponderomotive force and lose energy in this way

Several of these processes can be included in the coupled envelope equations and all are found to reduce efficiency.

The main point investigated in this thesis will be the kinetic effects, which act to damp the plasma wave. Since it has a slow phase velocity $\simeq \omega_p c/(2\omega_0)$, the plasma wave quiver velocity $\tilde{u}$ may easily exceed the phase velocity of the wave. If the amplitude of the wave is high enough for this to occur, the wave will cease to be sinusoidal. A significant fraction of electrons will exceed the phase velocity. The electric field of the wave is disrupted by these electrons and so the wave structure is lost.

2.2.2 Wave-breaking

Under certain conditions, the amplitude of the plasma wave becomes so large that electrons oscillate at greater than the phase velocity of the wave. In the non-linear stage, the amplitude of the plasma wave is given by equation (2.47) with $u \to \pi$ and $\alpha_0 = a_0$ is the initial amplitude of the pump. This gives

$$\tilde{n} = a_0 \left(\frac{2\omega_0}{\omega_p}\right)^{3/2} n_0. \tag{2.52}$$

The results of equation (1.49) relate the density and velocity amplitudes, giving $\tilde{u}/v_p = \tilde{n}/n_0$. This shows that the oscillation velocity of the electrons reaches the wave phase velocity precisely when the wave becomes non-linear, i.e. $\tilde{n}$ is an appreciable fraction of n_0. Combining $\tilde{u}/v_p = \tilde{n}/n_0$ with equation (2.52) and assuming the wave breaks when $\tilde{u} = v_p$ gives

$$\tilde{u}/v_p = a_0 \left(\frac{2\omega_0}{\omega_p}\right)^{3/2}, \tag{2.53}$$

$$a_{wb} = 0.35 \left(\frac{n}{n_{crit}}\right)^{0.75}. \tag{2.54}$$

This is the pump amplitude required for the wave to break, assuming the amplification has reached the non-linear stage. When the plasma wave becomes non-linear, so that $\tilde{n} \simeq n_0$, a significant fraction of electrons will exceed the phase velocity and these will not oscillate about a fixed position. The electrons will propagate forwards and won't be trapped within one period of the wave. They will greatly affect the electric field of the plasma wave and reduce the restoring force, disrupting the shape and coherence of the wave.

The discussions up to this point have also assumed that the electron thermal velocity is much less than the plasma wave phase velocity, so a_{wb} applies to cold plasma only. When the thermal velocity approaches the phase velocity, more electrons will exceed the phase velocity and so the threshold is expected to decrease for higher temperatures.

Since the Raman process couples coherent waves, any loss of coherence in the plasma wave will reduce the overall coupling efficiency [34]. Although some Raman scattering may still happen after the plasma wave breaks, the efficiency will be greatly reduced and any scattered light will lose temporal coherence. The effects of wave-breaking on the electron distribution will be simulated in subsequent chapters.

It is important to note that wave-breaking is a separate process to Landau damping. Landau damping depends only on the temperature of the plasma and the phase velocity of the wave, however wave-breaking can happen even in cold plasmas if the pump amplitude is high enough.

As a first order approximation, Raman scattering ceases once the wave breaks. Since the pump amplitude needed to cause this is a_{wb}, only a fraction $(a_{wb}/a_0)^2$ of the pump energy is scattered into the seed pulse. This means increasing the pump amplitude above a_{wb} gives no improvement to the final energy of the amplified pulse and the efficiency must decrease $\propto 1/a_0^2$.

There have been several studies aimed at operating the amplifier above the wave-breaking threshold [121], however other studies have used numerical kinetic simulations to show that efficiency is reduced in this regime [132, 118, 27].

2.3 Past experimental progress

Since the discovery of the high efficiency non-linear stage, there have been several attempts to access it experimentally. Unique signatures of the non-linear stage include broadening of the amplified pulse bandwidth, shortening of the amplified pulse, depletion of the pump pulse and significant energy transfer from the pump pulse to the seed [8]. Verification of the growth rates [105] and high gain have been demonstrated in the linear stage [130, 93].

Open gas jet targets allow versatile experiments, however the plasma is often low density and not uniform. Gas cells and capillary wave-guides offer a more uniform plasma [18], however there is less access for diagnostics [51]. Some experiments

employed a random phase plate to increase the beam waist within a distance L_R of the focus. However, this has a detrimental effect on the pulse coherence.

In the most successful experiment yet, both pulses were reflected back into the target in a double pass design [136, 126, 95, 98, 97]. The efficiency was 6.4%, meaning 6.4% of the pump pulse energy was transferred to the seed. The pump and seed pulse were exactly counter propagated and the plasma was pre-ionised. Progress has also been made towards higher energy pulses and larger plasma volumes [96, 56].

Chapter 3

Details of Computer Simulations and Plasma Diagnostics

This chapter introduces some of the computational and experimental methods used to investigate Raman scattering. All of the methods are well established and effective.

3.1 Particle-in-cell computer simulations

The equations introduced in the first chapter constitute a method to predict the evolution of a plasma system. Specifically, the plasma fluid equations (1.38, 1.42) give the evolution of a fluid, assuming that the pressure tensor is given by some known form. In some cases it can be assumed that the electric and magnetic fields are zero. This fluid method is relatively easy to compute, since only its density and velocity must be numerically integrated over three spatial dimensions. However, an assumption is made about the pressure, effectively averaging over the velocity space of the particles. Treating the plasma as separate ion and electron fluids allows particle separation and electric fields, but still assumes that both distributions are Maxwellian. This condition is violated within the non-linear electron plasma waves in Raman amplifiers.

To model plasma waves, the approach must clearly allow for charges to generate electric and magnetic fields. The two fluid model with inclusion of the fields was solved analytically in Chapter 2. However, this approach will not model the dynamics of the plasma wave and effects such as particle trapping. The physical model can be improved by modelling the full velocity space of the particles.

One must return to the Vlasov equation (1.34) which is completely general for a weakly coupled collisionless plasma. This equation must be integrated for

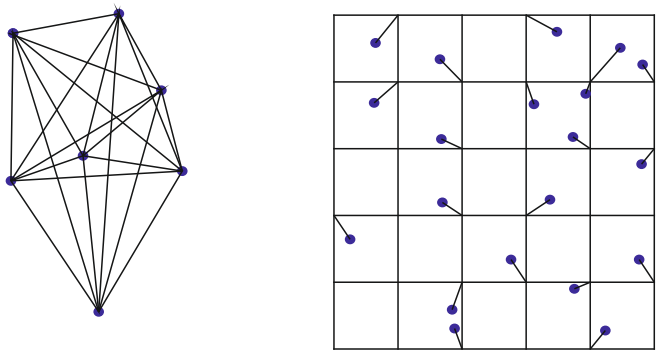


Figure 3.1: Comparison of two particle simulation algorithms. In the first, each particle interacts with all other particles due to the long range forces. In the second, the computational cost is drastically reduced by storing the position of each particle but calculating the forces across a discrete spatial grid. Purple dots represent the simulation particles and the lines represent the communications required.

each plasma species separately. The fields can be calculated through the Maxwell equations (1.9-1.12). This gives a closed system, however there are three spatial dimensions and three velocity dimensions and so any finite difference grid across this space must have a huge number of points. This renders the method extremely lengthy to compute. This method of modelling the full velocity space, as well as the three spatial dimensions, is known as a kinetic method.

To reduce the number of grid dimensions from six to three, it is possible to model the trajectory of a large number of electron and ion particles. This effectively solves the Vlasov equation for a finite number of sample points. If each particle is treated as free, such that quantum mechanical effects are negligible, the resultant dynamics of the many particle system is simply given by Newton's law

$$m_i \partial_t^2 \mathbf{x}_i = \sum_j \mathbf{F_{ij}}, \tag{3.1}$$

where $\mathbf{F}_{ij}$ is the force on particle i as a result of particle j. The electric and magnetic forces must be calculated for each particle and applied to each other particle separately. There are $O(N^2)$ particle pairings where N is the total number of particles, so the computational cost scales as $O(N^2)$. This strong scaling means directly solving the equations of motion is computationally expensive.

However, the scaling may be reduced to $O(N)$ by combining the particle method with a grid across only the spatial dimensions. This allows a large number of

particles to be modelled, typically $N = 10^8$ or more. This is still far short of the true number of particles in even 1 mg of plasma, so each simulation macro-particle typically represents of order 10^{10} real particles. This simplification is demonstrated in Fig. 3.1 which compares the case of interactions between all particles to the case where each particle only communicates with its nearest grid points.

The evolution of a macro-particle is identical to an electron since equation (3.1) depends only on the ratio q/m, which is identical to that of a single electron. Therefore the full number of particles need not be modelled. The ingenious computational saving lies in interpolating the particle quantities to a spatial grid, from which the Maxwell equations yield the fields. In this way, the fields and their effects can be calculated in time of $O(N)$, while keeping the particle kinetic nature and full modelling of the velocity space. The disadvantage is that the spatial resolution is reduced to that of the grid, meaning that the effect of small scale fields will be lost. Since collisionless plasmas are mostly affected by fields on scales $> \lambda_D$, it will be sufficient for the grid resolution to equal λ_D.

The algorithm proceeds as follows. Firstly, the large number of particles are initialised across the grid with a user defined density and velocity distribution. Each macro-particle has its own momentum vector. Since there may be large variations in the plasma density across the simulation box, typically the density of simulation macro-particles is constant, but the number of real electrons that each one represents is varied. This ensures that there are no discontinuities in the macro-particle density. In many situations, a laser pulse or external fields can also be initialised.

The main loop then begins, integrating the trajectory of each particle forwards in time. Since the finite difference integration method is explicit, the duration of each time-step must obey the Courant condition

$$\frac{1}{\Delta t^2} > \sum_i \frac{c^2}{\Delta x_i^2}, \tag{3.2}$$

where the sum is taken over all spatial dimensions. This effectively means that within one time step, a light wave must not be able to propagate more than a single grid point. If the Courant condition is violated, the solution will be numerically unstable and diverge. However, high order interpolation algorithms are not bound by this.

3.2 Particle-in-cell (PIC) algorithm

3.2.1 Interpolation to the grid

In the first stage of the time-step, the particles have their charges and currents deposited on to the grid. Since the grid has a finite resolution, this interpolation will introduce an error. The simplest scheme is to deposit all of the charge and current from a particle on to the nearest grid point, however this leads to discontinuities as the deposition will suddenly switch from one grid point to another as the particle moves a small distance. These discontinuities introduce a large amount of statistical noise. This zero order scheme means the effective weight function of a particle is a rectangle or cuboid.

Higher order interpolation techniques will reduce the noise, since the weight function of the particle is smoother and spread over a greater area. These functions conserve the particle's charge and current. They can be generated via auto-convolutions of the zero order weight function. The penalty for use of a higher order technique is a greater computational cost for each particle, since the weight function is wider and the contribution must be calculated for each grid point.

The relative noise from one grid cell to the next will scale as $1/\sqrt{N_C}$, where N_C is the number of particles per cell. This typically means the total number of particles must be much greater than the number of grid points. Since the grid resolution is of order λ_D, the number of particles $N \gg (L/\lambda_D)^d$, where d is the number of dimensions and L is the spatial scale of the system.

3.2.2 Field solve

Once $\mathbf{J}$ and ρ are known on the grid points, the next stage is to solve the Maxwell equations (1.9-1.12). This is an operation that involves only the grid, and not the particles. Since the number of particles is much greater than the number of grid points, this is typically the stage that is quickest to compute.

For the electrostatic cases, with constant magnetic fields, the Poisson equation can be solved via spectral or finite difference methods. For the general case, the third and fourth Maxwell equations must be integrated in time. The curl of the fields can be evaluated using centred finite differences, then integrated forwards in time using a Runge-Kutta method or similar.

At this stage, the boundary conditions for the fields will be important. The boundaries can be open, meaning that the potential extends to infinity. Another option is periodic boundaries, where the field wraps around to the opposite side of the simulation domain. The potential can also be set to a certain value on the boundary.

The boundary condition at time zero can include a laser pulse initialised in the box. Alternatively, the laser may propagate into the box by initialising the potential on the boundary as a function of time. Some algorithms may also smooth the current and field quantities spatially, to partially remove any high frequency noise resulting from the interpolation.

Lastly, in order to move the particles, the field quantities must be known at the position of the particles. This requires a second interpolation step, similar to the first. The field quantities vary smoothly and important features are resolved by the grid. However, effects such as binary collisions involve fields on scales smaller than λ_D and much smaller than the grid resolution. This is why collisional plasmas are not well modelled by the standard algorithm.

3.2.3 Particle push

Next, the particle positions and momenta must be updated using $\partial_t \mathbf{x} = \mathbf{v}$ and $\partial_t \mathbf{v} = (q/m)(\mathbf{E} + \mathbf{v} \times \mathbf{B})$. The particle momenta and position can be numerically integrated using several methods.

The simplest particle push is to use the Euler method at time-step n such that $\mathbf{x}_{n+1} = \mathbf{x}_n + \mathbf{v}_n \Delta t$. However, this method has truncation error that is linear order with Δt. A more accurate method is the so called leapfrog algorithm that specifies the velocities on a grid at half time-steps $t = (n + 0.5)\Delta t$ for integer n. The time-step is Δt. This yields $\mathbf{x_{n+1}} = \mathbf{x_n} + \mathbf{v_{n+1/2}} \Delta t$. The particle velocities must be initialised at $t = -\Delta t/2$ instead of $t = 0$. This initial step can be done using the Euler method. The leap-frog method gives second order accuracy at no additional computational cost.

The inclusion of magnetic fields $\partial_t \mathbf{v} = (q/m)(\mathbf{E} + \mathbf{v} \times \mathbf{B})$ poses a problem for the leapfrog method, since the centred finite difference equation becomes

$$\frac{\mathbf{v}_{n+1/2} - \mathbf{v}_{n-1/2}}{\Delta t} = \frac{q}{m} \left(\mathbf{E}_n + \mathbf{v}_n \times \mathbf{B_n} \right), \tag{3.3}$$

$$\simeq \frac{q}{m} \left(\mathbf{E}_n + \frac{\mathbf{v}_{n+1/2} + \mathbf{v}_{n-1/2}}{2} \times \mathbf{B_n} \right). \tag{3.4}$$

This yields an implicit matrix equation for $\mathbf{v}_{n+1/2}$, since it cannot be easily re-arranged due to the cross product with $\mathbf{B}$. The re-arranged matrix form gives

$$M_+ \mathbf{v}_{n+1/2} = \frac{q\Delta t}{m}\mathbf{E} + M_- \mathbf{v}_{n-1/2}. \tag{3.5}$$

The components of the matrices $M_\pm$ are

$$(M_\pm)_{ik} = \delta_{ik} \pm \frac{q\Delta t}{2m}\epsilon_{ijk}B_j. \tag{3.6}$$

The updated velocity is then found from equation (3.5) by analytically inverting matrix M_+.

This update of particle position may take it outside of the simulation domain. If so, there are several possibilities, specified by the particle boundary condition. The particle could be reflected by an elastic collision. This conserves energy, but may lead to unrealistic velocity distributions, since the particle will return back to the interaction. The particle could be absorbed and deleted if it hits the boundary, however this will not conserve the number of particles. Another option is the periodic boundaries, however this will again mean a fast particle leaving the box will return back to the interaction region, possibly affecting the results. This can be fixed by returning the particle but changing its energy to one sampled from a Maxwellian distribution. This conserves the particle number and is more similar to experiments, since the interaction will usually be surrounded by a plasma thermal bath.

The final stage is to record any desired quantities for the output of the simulation, such as the field values on the grid or the distribution of particle momenta. These may be spatially or temporally averaged to reduce memory usage. Although it is infeasible to save all of the trajectory data for every particle, it can be done for a sample of the macro-particles.

This whole time step is repeated in a loop for the desired simulation duration T_{sim}. Since $\Delta t \simeq \lambda_D/c$, there will be approximately cT_{sim}/λ_D time-steps. The total computational cost scales as the number of particles multiplied by the number of time-steps, $\propto (L/\lambda_D)^{d+1}$. It is therefore much easier to simulate plasmas with large Debye lengths, with low density and high temperature.

The algorithm does not depend on any specific co-ordinate system for the grid or particles, since the derivatives can be expressed in Cartesian, cylindrical or spherical coordinates. Furthermore, the number of spatial dimensions can be one, two or

three, so long as the particle velocities can still be three dimensional. This can be kept consistent by specifying that the dimensions not simulated are periodic, so that all physical quantities are constant along these other dimensions. All simulations in this work will use the Cartesian co-ordinate system in one or two dimensions.

3.2.4 Moving simulation domain

Many phenomena of interest in plasma physics involve structures that move at large velocities. If the simulation uses a fixed domain in the laboratory frame, it will have to be large enough to contain the full trajectory of the object. Raman amplification is also of this type, since the interaction occurs within the seed pulse which moves at close to c.

The computational cost can be greatly reduced by using a simulation domain that encloses only the interaction volume, then moving along with it as the propagation proceeds. This is possible without a Lorentz boost by simply adding new grid points to the front of the domain and deleting those at the back. New macro-particles must be initialised at the front of the moving window. In addition, modelling Raman amplification requires a long counter-propagating laser pulse to arrive in from the front of the window. This not only requires suitable initialisation of the fields, but also the new macro-particles must be initialised with the velocity given by Fig. 1.8.

Although the moving window simulations greatly decrease the computational expense, they do not properly model the instabilities of the pump pulse as it traverses the plasma. There may also be noise from the initialisation at the front of the window which could artificially seed instabilities such as Raman scattering. It is therefore beneficial to verify the results of moving window simulations with a similar fixed window set-up.

For relativistic PIC codes, another possibility to model moving interactions is to conduct the simulation in a Lorentz boosted reference frame.

All parametric instabilities may be over estimated by PIC codes, since the statistical noise due to the macro-particles provides an artificially large seed to start the exponential growth. This effect may be reduced by using greater numbers of simulation particles.

Another consideration is that a thermal distribution can only be properly modelled if there are enough constituent particles to fill it. For example, particles

exceeding twice the thermal velocity along the x direction constitute only 2% of the total particle number in a Maxwellian distribution. This means Landau damping and kinetic effects will only be modelled well with > 1000 particles per wavelength.

The grid resolution should be fine enough to resolve the smallest spatial scales of interest, which is typically the plasma wave wavelength. This already requires > 30 cells per laser wavelength. However, there is another requirement due to artificial numerical heating. This is an artificial numerical instability which does not conserve energy. Typically if the Debye length is not resolved by the grid, interpolation errors will accumulate and increase the electron temperature until the Debye length is resolved. The timescale for this can be very short, of the order ω_p^{-1}. Once the Debye length is resolved, the artificial heating will slow, but not stop entirely.

The smoothing of the current and electromagnetic fields may also violate the conservation of energy, so it is important to observe the electron velocity distribution to verify that energy is conserved in the simulation. All of these caveats are minimal if they are monitored by observing the simulation data.

Modern particle-in-cell codes have many additional features such as relativistic corrections, a variety of types of laser pulses, optimisation via sub-cycling and processor load balancing, quantum electrodynamics effects and modelling of collisions.

3.3 Experimental interferometry diagnostic

In a high-power laser experiment, it is important to know the shape and density of the plasma generated by the high-power laser, as a function of time. This is especially important for Raman scattering, where the spectrum of the scattered light depends on the plasma density. I used this diagnostic to experimentally measure the plasma profile in Chapters 7 and 8.

The plasma density may be measured with knowledge of its refractive index. Equation (1.95) gives the refractive index of the plasma as

$$\frac{c}{v_p} = \sqrt{1 - \frac{\omega_p^2}{\omega^2}} = \sqrt{1 - \frac{n}{n_{crit}}} \simeq 1 - \frac{n}{2n_{crit}}, \tag{3.7}$$

where the expansion is valid for very underdense plasmas. To measure the refractive index, one must measure the phase difference between a reference wave that does not enter the plasma and one that enters the plasma. This is possible by interfering

the two waves using an interferometer and forming an image of the plasma. In practice, there is an imaging system using lenses and a CCD detector. However, before the image reaches the CCD, it is interfered with another image that passed down the other arm of a Mach-Zehnder or Michelson interferometer. The arms are offset such that the image of the plasma from one arm interferes with the image of a position outside the plasma from the other arm. The offset leads to a path difference between the two images which produce interference fringes on the CCD. The plasma causes a phase delay which bends these fringes. Therefore, by measuring the fringe position, the phase difference due to the plasma can be acquired. This phase difference is

$$\phi = \frac{2\pi}{\lambda} \int (1 - \frac{c}{v_p}) dx, \tag{3.8}$$

$$\simeq \frac{\pi}{\lambda} \int \frac{n}{n_{crit}} \, dx, \tag{3.9}$$

where x is the direction of propagation of the coherent light that forms the image and λ is its wavelength. This direction can be set transversely across the plasma or coincident with the main laser pulse.

If it crosses the plasma transversely and the plasma has axial symmetry, an Abel inversion can be used to extract the electron density. Another technique is to estimate the transverse shape and width of the plasma and perform the integration.

3.4 Experimental FROG diagnostic

Characterisation of a short laser pulse is possible with the FROG laser diagnostic. This was useful in the measurement of the Raman scattered light, described in Chapter 7. Since electronic circuits do not temporally resolve events below approximately 1 ns, optical methods must be used to measure pulses shorter than this. One such method is the auto-correlator, whereby a test pulse is split into two copies and recombined in a crystal with a non-linear refractive index. This has the effect of generating light at double the frequency of the incident light, a process known as frequency doubling. Since the non-linear refractive index depends on intensity, more light will be converted when the incident light intensity is greater.

The FROG diagnostic allows single shot characterisation of a laser pulse because it recombines the two copies at an angle, such that on one side of the crystal the first pulse arrives before the second copy. On the other side of the crystal, this

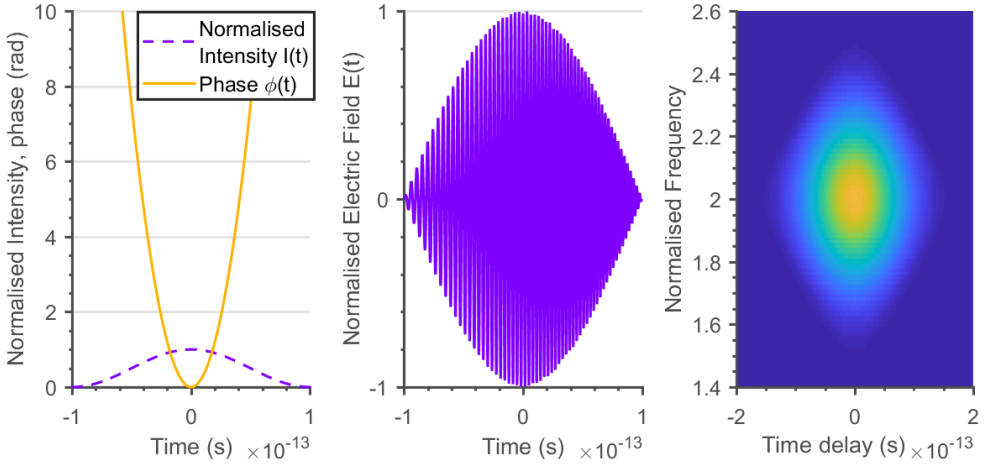


Figure 3.2: Numerical demonstration of the frog diagnostic. (a) The intensity and phase of a chirped test pulse in the time domain. (b) The corresponding electric field as a function of time. (c) The frog trace acquired using equation (3.10)

delay is reversed. Delay time is therefore mapped to the transverse position. The output electric field of the doubled light is proportional to $E(t)E(t-\tau)$, where τ is the delay between the two pulse copies. This is directly proportional to transverse position.

Next, the frequency doubled light is spectrally resolved. Since the spectrometer detector is relatively slow, it will yield a time integrated signal, given by

$$A(\tau,\omega) = \left| \int E(t)E(t-\tau)e^{-i\omega t}\,dt \right|^2 . \tag{3.10}$$

This is the same as the square modulus of a windowed Fourier transform, also known as a spectrogram. However, typically with a windowed Fourier transform (also known as short time Fourier transform) the window function is known. Here, the pulse field itself is what gates the Fourier transform. This is an unknown function, so a simple inversion from $A(\tau,\omega)$ to $E(t)$ is not possible. However, since the spectrogram is a function of two variables constructed from a function of one variable, it is over determined. Every possible $E(t)$ maps to a spectrogram, but only certain spectrograms map to a $E(t)$. The rest are unphysical.

This fact is useful for experimental verification, since if $E(t)$ is recovered, it can be forward transformed to find the spectrogram and this should exactly match with the experimental data. Comparing this retrieved spectrogram to the measured one gives an indication of the measurement validity. The data may be inverted via several algorithms.

This process is shown numerically in Fig. 3.2. Starting with a chirped pulse with a $\sin^2$ intensity envelope and a quadratic phase term $\phi(t) \propto t^2$, the simulated FROG trace is obtained. The pulse duration is 100 fs, with central wavelength 1 μm. The intensity and phase are related to the electric field via $E(t) \propto \sqrt{I(t)}\cos(\omega_0 t + \phi(t))$. This leads to a pulse that is chirped in the time domain and so has a duration greater than its transform limit. The pulse electric field is shown in panel (b). Next, the spectrogram can be found via a discrete Fourier transform for each value of delay. Panel (c) shows the simulated data that would be measured with the FROG device.

The light has approximately double the frequency of the incident pulse and the width gives an approximation to the pulse duration. However, this can be misleading and the algorithm should be used to properly invert the data. There is also a time reversal symmetry with the diagnostic, so it is impossible to know if the chirp is positive or negative.

A second version of the FROG diagnostic is known as GRENOUILLE. This uses the same principle, except the beam is not split. Instead, a Fresnel bi-prism separately focuses the left and right halves of the beam. This provides the two copies of the pulse, incident on the crystal at an angle. Furthermore, the spectrometer is eliminated by using the angular dispersion of the non-linear crystal itself. This significantly simplifies the apparatus.

Chapter 4

Optimisation of the Amplifier

In Chapter 2, the fluid equations showed that a pump pulse counter-propagating with a short seed pulse will beat together in a plasma and produce a ponderomotive pressure that oscillates with the frequency difference of the waves. If this is fixed at ω_p, the ponderomotive pressure will oscillate in the longitudinal direction at frequency ω_p. This will resonantly excite a large amplitude plasma wave which will scatter the pump pulse in the backwards direction.

In the initial linear stage, the pump depletion is negligible, the seed pulse amplitude grows exponentially with growth rate γ_{RBS} and it lengthens in duration. Eventually the seed becomes large enough to appreciably deplete the pump pulse and the process enters the non-linear stage. The intermediate stage is hard to treat, however there is an analytic solution for late times. The energy and amplitude of the amplified pulse grow proportional to t, its intensity grows $\propto t^2$, its duration shortens $\propto 1/t$ and the pump is fully depleted. This solution assumes that $a_0 < a_{wb}$ and damping processes are negligible.

In this chapter, I will verify this high efficiency analytic solution by numerically integrating the envelope equations and compare them to particle-in-cell simulations that capture more of the physics and make fewer assumptions. All of the analytical results will be verified, providing the assumptions hold. The simulations will be used to evaluate the efficiency of the amplifier across the full parameter space, capturing all of the relevant physics.

Next, I turn attention to the seed pulse to find its optimal parameters under experimentally relevant constraints. I will find that the analytical solution breaks down for seeds with a duration longer than γ_{RBS}^{-1}, suggesting that use of a shorter seed pulse will increase performance. Finally, the optimal parameters will be tested in a full scale two-dimensional simulation.

4.1 Efficiency of the amplifier

The coupled envelope equations have only one spatial dimension and the interaction is periodic in the transverse directions, so efficiency can be defined using fluence rather than energy. The energy density of an electromagnetic wave is proportional to $I \propto \omega^2 a_0^2$. The total energy efficiency is defined as the total change in seed pulse fluence Φ_1, divided by the pump pulse fluence Φ_0 that has passed through the seed pulse. The efficiency will therefore change with time and is given by

$$\eta(t) = \frac{(\Phi_1(t) - \Phi_1(0))}{\Phi_0(t)}, \tag{4.1}$$

$$\Phi_1(t) = \frac{1}{c}\int \omega_1^2 a_1^2(t,x)dx, \tag{4.2}$$

$$\Phi_0(t) = \frac{v_0 + v_1}{c}\int_0^t \omega_0^2 a_0^2(t',0)dt'. \tag{4.3}$$

If the pump pulse amplitude is constant then $\Phi_0 = (v_0 + v_1)ta_0(0)^2\omega_0^2/c$, where $a_0(0)$ is the pump amplitude before the interaction. The seed pulse fluence Φ_1 is an integrated quantity and it includes all of the energy in the secondary spikes of the amplified pulse.

Another interesting quantity is the energy transfer efficiency just to the first spike of the amplified pulse. Although it is difficult to define the extent of the first spike, its fluence may be approximated as $\omega_1^2 a_1^2 T_1$ where a_1 is the peak amplitude and T_1 is the duration of the first spike of the amplified pulse. This gives the efficiency to the first spike as

$$\eta_1(t) = \frac{(a_1^2(t)T_1(t) - a_1^2(0)T_1(0))\omega_1^2 c}{a_0^2 t(v_0 + v_1)\omega_0^2}. \tag{4.4}$$

As long as the first spike has a temporal shape typical of a Gaussian laser pulse then η_1 should be less than η, although it is not guaranteed to be. This is because of the rather arbitrary definition of the duration of the first spike as its FWHM, which may not be accurate for some pulse shapes. However, the approximation is valid in most cases of interest, in which the amplified pulse is well behaved and has a well defined main peak.

The efficiency may also be defined through a derivative quantity as the increase in seed fluence divided by the rate that pump fluence reaches the seed,

$$\eta_{diff}(t) = \frac{\Phi_1'(t)}{\Phi_0'(t)}. \tag{4.5}$$

Since this is a derivative quantity, it will give the instantaneous efficiency at any point in the amplifier. However, taking numerical derivatives of simulated data often introduces noise and so I only use the definitions (4.1) and (4.4) to analyse the simulations.

These definitions also have the advantage that the total output energy of the amplifier is easily calculated as $E_{out} = E_{pump}\eta(t = L/c)$, where E_{pump} is the starting pump energy and L is the length of the amplifier. The optimal pump duration should be $2L/c$ since this is the overlap time of two counter-propagating pulses over a length L. In addition, the energy in just the leading spike of the amplifier is $E_{1,out} = E_{pump}\eta_1(t = L/c)$.

The maximal theoretical efficiency may be evaluated from the coupled wave equations. I now show that in the limit of $t \to \infty$, the differential efficiency η_{diff} converges to a constant value that depends only on plasma density. The total efficiency η must also converge to this value.

For a constant amplitude pump, the rate that its fluence reaches the seed is $\Phi_0'(t) = (v_0 + v_1)a_0^2\omega_0^2/c$. Substituting this and equation (4.2) into equation (4.5) gives

$$\eta_{diff}(t) = \frac{\Phi_1'(t)}{\Phi_0'(t)} = \frac{1}{(v_0 + v_1)a_0(0)^2\omega_0^2}\frac{d}{dt}\int \omega_1^2 a_1^2\, dx \tag{4.6}$$

$$= \frac{\omega_1^2}{(v_0 + v_1)a_0(0)^2\omega_0^2}\int 2a_1\partial_t a_1\, dx. \tag{4.7}$$

Now replacing $\partial_t a_1$ with the envelope equations for the non-linear stage (2.40) and $\omega_1 = \omega_0 - \omega_p$, the efficiency becomes

$$\eta_{diff} = \frac{(\omega_0 - \omega_p)^2}{(v_0 + v_1)a_0(0)^2\omega_0^2}\int 2a_1\left(-\frac{\omega_p^2}{4(\omega_0 - \omega_p)}\right) n_2 a_0\, dx \tag{4.8}$$

$$= -\frac{(\omega_0 - \omega_p)\omega_p^2}{2(v_0 + v_1)a_0(0)^2\omega_0^2}\int a_0 a_1 n_2\, dx. \tag{4.9}$$

The instantaneous efficiency is therefore the total overlap integral of the three waves at any given time. The intermediate stage is difficult to treat, however in the advanced non-linear stage equation (4.9) can be simplified using the coupled equation

(2.39) for the non-linear stage to replace $a_1 n_2$. It becomes

$$\eta_{diff}(t) = -\frac{(\omega_0 - \omega_p)\omega_p^2}{2(v_0+v_1)a_0(0)^2\omega_0^2}\int\left(\frac{4\omega_0(v_0+v_1)}{\omega_p^2}\right)a_0\partial_x a_0\,dx \tag{4.10}$$

$$= -\frac{(\omega_0-\omega_p)}{\omega_0 a_0(0)^2}\int\partial_x(a_0^2)\,dx \tag{4.11}$$

$$= -\frac{\omega_0-\omega_p}{\omega_0}\frac{\Delta(a_0^2)}{a_0(0)^2}. \tag{4.12}$$

Here the quantity $\Delta(a_0^2)$ means the change in a_0^2 across the full interaction. The efficiency is proportional to the pump depletion, so the maximum efficiency is when the pump is fully depleted, giving

$$\eta_{max} = 1 - \frac{\omega_p}{\omega_0}. \tag{4.13}$$

This shows that a greater maximal efficiency may be achieved at lower plasma densities, since the plasma wave then has a lower energy density and more energy is transferred to the seed pulse. The quantity η_{max} is plotted in Fig. 4.1 for a pump pulse of wavelength 1 micron.

A more general expression for the efficiency can be found by substituting the trial functions for the self-similar solutions (2.45) into equation (4.9). This leads to

$$\eta = \frac{\omega_1\omega_p^2}{2\omega_0^2}\int_0^\infty\frac{\omega_0}{\omega_p^2}\sin(u)\partial_x u\,dx \tag{4.14}$$

$$= \frac{\omega_1}{2\omega_0}\int_0^{\Delta u}\sin(u)\,du. \tag{4.15}$$

The self similar solution shows that the upper limit on the u integral is just the total integrated amplitude of the seed at any time. Performing the integral gives

$$\eta(t) = \eta_{max}\frac{1-\cos(\Delta u)}{2}, \tag{4.16}$$

$$\Delta u = -\sqrt{\frac{\omega_0\omega_p}{2c^2}}\int a_1\,dx. \tag{4.17}$$

When the integrated amplitude reaches its convergent value of $\Delta u = \pi$, the efficiency converges to η_{max}. As long as the main assumption is still valid, that the duration of the pulse is much less than the timescale of evolution, this expression gives the efficiency for any arbitrary seed pulse.

Interestingly, very energetic seed pulses exceeding the non-linear condition (2.44) will actually have a lower efficiency than η_{max}. This suggests that the seed pulse

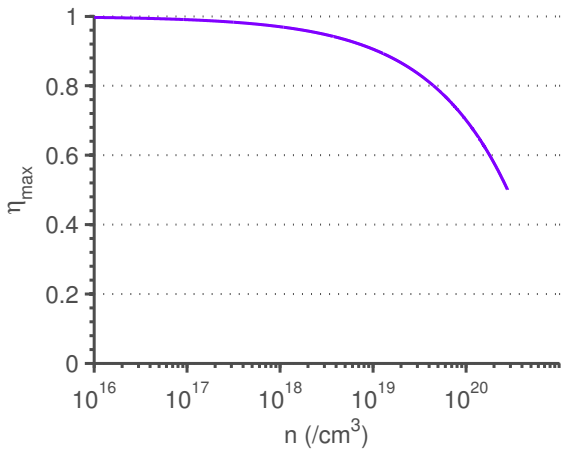


Figure 4.1: The maximum achievable energy transfer efficiency to the amplified pulse, for a pump pulse of wavelength 1 μm. This includes the energy transferred to the whole of the amplified pulse, including any secondary peaks. The model assumes zero plasma temperature and does not account for kinetic or damping effects in the plasma.

must be carefully planned to start with Δu exactly equal to π. Further increasing the integrated amplitude Δu of the initial seed pulse will actually have a detrimental impact on the efficiency.

4.2 Optimisation of the plasma parameters

4.2.1 Plasma constraints

The objective of a plasma amplifier is to maximise the power of the amplified pulse, while maintaining transverse coherence so that it can be focussed. Some applications such as fast-ignition may also require a high energy and/or a certain spectral range for the pulse. Certainly the energy transfer efficiency must also be large to compete with traditional grating compressors, which have efficiency greater than 50%.

I first consider the required plasma conditions. Depending on the plasma temperature and density, the Landau damping rate and collisionality will change [9]. Saturation of the Landau damping may allow slightly greater temperatures than the linear damping rate would suggest [135, 16].

Equation (2.44) shows that an amplified pulse approaching relativistic intensity must have duration $T_1 \simeq \omega_p^{-1}$. For the plasma wave to sustain an interaction of duration ω_p^{-1}, one requires $\gamma_L/\omega_p < 1$ and $\omega_p \tau_{ei} > 1$. Since the phase velocity of the plasma wave for backward scattering is $c\omega_p/(2\omega_0)$, the Landau rate can be calculated

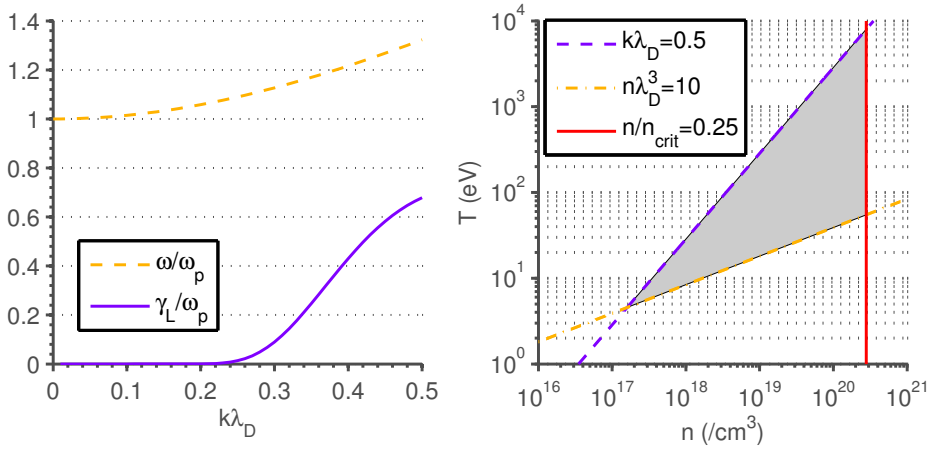


Figure 4.2: (a) The Maxwellian averaged Landau damping rate for the plasma wave in a backwards Raman amplifier. Damping becomes appreciable for $k\lambda_D > 0.3$, where k is the wavenumber of the plasma wave. Also shown is the real part of the plasma wave frequency, showing the thermal correction as the temperature is raised. (b) The plasma conditions for an amplifier with a pump of 1 μm wavelength. Landau and collisional damping on timescales ω_p^{-1} must be avoided, as well as remaining below the quarter critical density. This restricts the plasma conditions to a small region (shaded).

explicitly and is shown in Fig. 4.2. The Landau damping becomes appreciable at the same time as the thermal shift to the plasma wave frequency becomes important. Minimising Landau damping requires that $k_2\lambda_D < 0.5$. Nearing this limit will excessively damp the plasma wave and lower the efficiency.

The condition for a collisionless interaction is $\omega_p\tau_{ei} > 1$, meaning $n\lambda_D^3 > Z\ln(3\pi n\lambda_D^3/Z)$. I take $n\lambda_D^3 > 10$. These conditions are shown in the Fig. 4.2 for a pump pulse of wavelength 1 μm. The equivalent plot for shorter wavelengths will be shown in later chapters. Landau damping favours lower temperature and higher density, but reducing collisional damping requires the opposite.

The final condition for Raman scattering that $n < n_{crit}/4$ confines the interaction to a small region of plasma conditions. The region is narrow and shows a minimum feasible plasma density around $10^{17}\,\mathrm{cm}^{-3}$.

4.2.2 Ionisation schemes

In addition to the plasma variables n and T, the pump pulse amplitude a_0 is the last main variable. Under the standard scheme, the pump pulse will ionise the plasma. This means the pump amplitude must be high enough to form the plasma. For a wavelength of 1 μm in hydrogen plasma, the field ionisation intensity of $1.4\times$

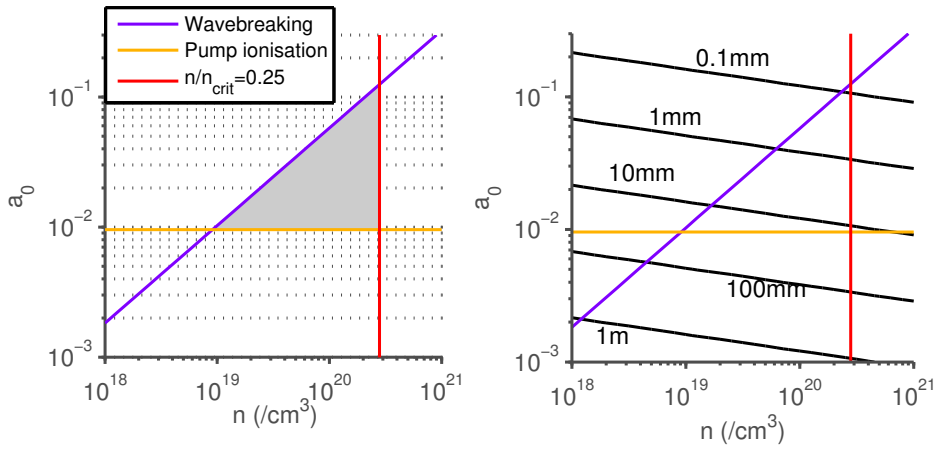


Figure 4.3: (a) The parameter space of plasma density versus pump amplitude, assuming a constant zero temperature, for a 1 μm wavelength. The plasma wave will break and lose coherence above the wave breaking threshold (purple) for a cold plasma (2.54). For warm plasma, this constraint will worsen. In the standard scheme, the pump must ionise the plasma and therefore requires a minimum amplitude (orange). The shaded area indicates the permitted region with high efficiency. (b) The same parameter space and limitations are shown, now with contours for the length L_{max} of the amplifier required to reach relativistic intensity. This distance is measured from the point where the non-linear stage is first reached. The expression for L_{max} also depends on the efficiency, as a simple estimate this is fixed as 50%.

$10^{14}\,\mathrm{Wcm}^{-2}$ corresponds to $a_0 = 0.01$.

Kinetic effects also constrain the pump amplitude. It must not be so high that the wave-breaking limit is exceeded, in which case η decreases. These considerations are shown in the Fig. 4.3, again for a wavelength of 1 μm and assuming zero temperature so that the cold wave-breaking limit (2.54) applies.

Ionisation can also be achieved through an external source such as a heating pre-pulse or a capillary discharge, however this will raise the plasma temperature and the following sections will show that this negatively affects the efficiency.

A third scheme uses the seed pulse to ionise the plasma from the start [68, 11, 12]. If the pump intensity is below the ionisation threshold, this means the plasma will only be ionised once the seed pulse arrives. This has the advantage that the pump will not be subject to plasma instabilities as it traverses the target before meeting the seed pulse. In addition, the lower limit on the pump amplitude is removed.

Considering Fig. 4.3, efficient operation of the pump ionisation scheme will not be possible for plasma densities below $10^{19}\,\mathrm{cm}^{-3}$ for a 1 μm pump pulse. For lower densities, either the seed ionisation scheme or pre-ionisation is required.

4.2.3 Dynamics of the amplified pulse

Using the results of Chapter 2, and considering energy conservation, I now evaluate the seed pulse amplitude. Assuming that the initial seed is small, and that the plasma is very underdense (meaning $v_0 + v_1 \simeq 2c$), equation (4.4) becomes

$$\omega_1^2 a_1^2 T_1 = 2\eta_1 \omega_0^2 a_0^2 t. \tag{4.18}$$

Replacing the seed duration using the main result for the non-linear stage (2.44) yields the seed amplitude as a function of time [123],

$$a_1(t) = \frac{2}{5}\sqrt{\omega_0 \omega_p}\left(\frac{\omega_0}{\omega_1}\right)^2 \eta_1 a_0^2 t. \tag{4.19}$$

The time is measured from the point where the pulse enters the non-linear stage. Also using $\omega_1 = \omega_0 - \omega_p$, the pulse will reach relativistic intensity $a_1 = 1$ after a distance

$$L_{max} = \frac{5c(\omega_0 - \omega_p)^2}{2\eta_1 a_0^2 \omega_0^2 \sqrt{\omega_0 \omega_p}}. \tag{4.20}$$

This places the eventual limitation on the length of the amplifier, since a relativistic seed pulse will change the effective plasma frequency and invalidate the frequency matching condition [76]. This problem and others may be overcome by chirping the pulses [30, 72, 124, 136, 129, 139, 119, 90, 3], however this scheme will not be considered here. The chirp in the pump pulse compensates for the change in the plasma frequency due to the electron mass increase [75].

Contours of L_{max} are shown in Fig. 4.3b for a 1 μm wavelength pump pulse. To stay within the high efficiency region, the length of the amplifier must be between 0.1 mm and 20 mm. The maximum pump duration is then $2L/c = 0.13$ ns. This means any efficient amplifiers using pulses longer than 1 ns must access the region below the yellow line in Fig. 4.3 and the pump is below ionisation intensity. Therefore amplification of any pulse longer than 0.13 ns must use the seed ionisation or pre-ionisation schemes.

Other processes may place a more stringent limit on the total amplification length. Other competing instabilities grow along with the Raman backward scattering. Chief among these are Raman forward scattering of the amplified pulse and ponderomotive filamentation of the amplified pulse. To estimate these limitations, one must first verify the equations (2.44) and (4.19) with detailed simulations.

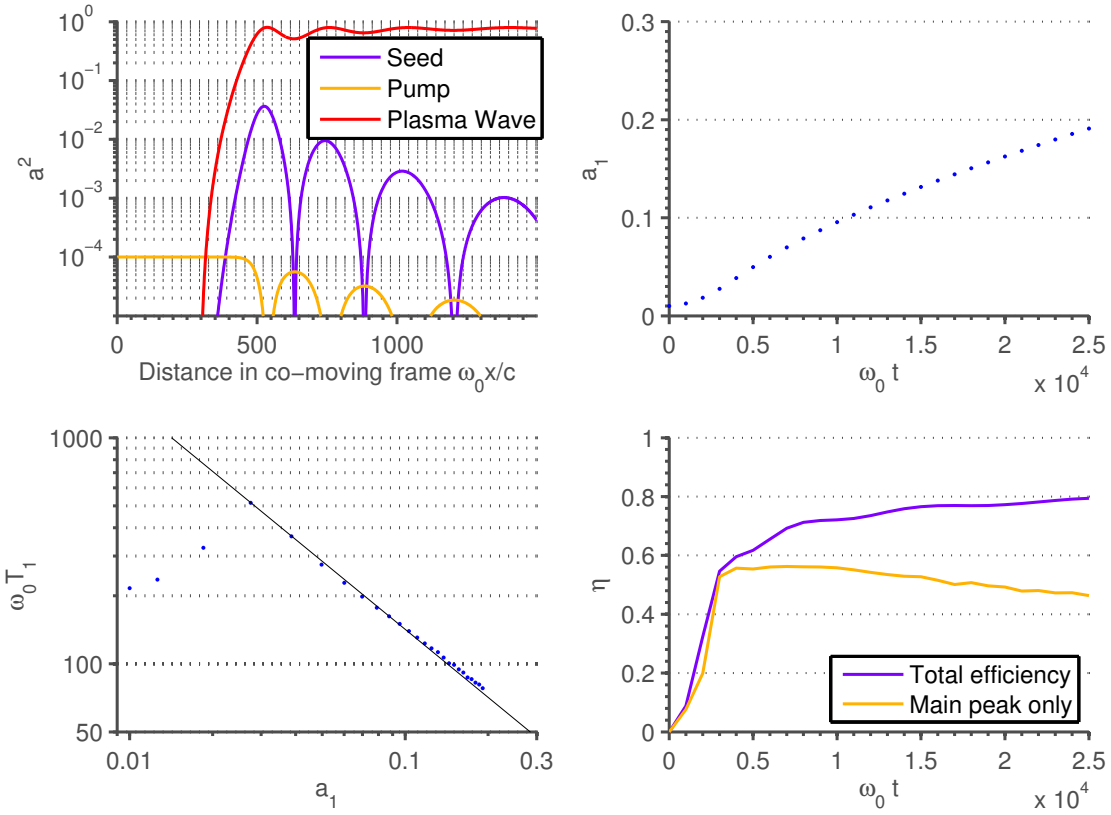


Figure 4.4: Numerical integration of the envelope equations for $n/n_{crit} = 0.01$ and $a_0 = 0.01$, with initial seed amplitude $a_1 = 0.01$ and duration $\omega_0 T_1 = 210$. (a) The longitudinal profile of the pump and seed power densities a_0^2, a_1^2 and the plasma wave amplitude n_2 after a time $\omega_0 t = 25000$. The seed propagates to the left. (b) The seed pulse amplitude as a function of time. It increases linearly with time. (c) The seed pulse amplitude and FWHM duration, over the course of the interaction. Each data point is separated by a time $\omega_0 t = 1000$. The seed amplitude monotonically increases. The seed pulse is attracted to the analytical solution for late times (solid line), given by equation (2.44). (d) The total energy transfer efficiency (η) and the efficiency to just the first spike of the amplified pulse (η_1) as a function of time.

4.3 Numerical integration of the envelope equations

I now numerically solve the envelope equations (2.36-2.38) to verify the analytical solution for late times. The three coupled envelope equations were numerically integrated in time, using a fourth order explicit Runge-Kutta method. Spatial partial derivatives at constant time were taken using a fourth order centred finite difference and a fourth order one-sided finite difference for the four cells nearest the boundaries. The boundary condition was constant a_0 at the front of the box, with a perfectly absorbing boundary condition for all waves at the rear. The box

had length $\omega_0 x/c = 1500$ and was integrated for a time $\omega_0 t = 25000$. The finite difference grid size was $1.5c/\omega_0$ and the time-step was $0.75/\omega_0$. For numerical stability, the time-step multiplied by c was less than the spatial grid resolution.

Conditions just below the wave-breaking threshold were chosen, such that $n_2 < 1$. The seed pulse also started below the condition for the non-linear stage. Fig. 4.4a shows the envelopes of the seed, pump and plasma wave after a propagation time of $\omega_0 t = 25000$. The seed has propagated to the left and the incoming pump pulse is seen in the lower left of the plot, propagating towards the right. The pump is heavily depleted in the interaction and the seed has reached an amplitude much greater than the pump. In addition, the characteristic secondary peaks trail the main spike of the amplified pulse. The plasma wave (red) grows from zero to a high amplitude. It remains at this high amplitude since the model assumes there are no damping mechanisms.

To test the analytical theory, the seed pulse FWHM duration and peak amplitude were taken as a function of time. Fig. 4.4b shows the amplitude as a function of time. There is an initial start up stage, beyond which the amplitude is proportional to time. The final amplitude of $a_1 = 0.19$ is in close agreement to the prediction of equation (4.19) which gives $a_1 = 0.18$.

The analytical solution is seen even more clearly when the amplitude versus seed duration is plotted in panel (c). After the initial linear stage where the pulse lengthens, the amplified pulse satisfies the non-linear criterion $a_1 T_1 \sqrt{\omega_0 \omega_p} = 4.5 \simeq \pi\sqrt{2}$ (solid line). The pulse subsequently shortens as it amplifies.

The efficiency in the numerical integration in shown in Fig. 4.4d. The efficiency starts at zero. The efficiency monotonically increases through the linear stage until it converges to a constant value. This is a characteristic of the non-linear stage.

The theoretical efficiency is $\eta_{max} = 1 - \omega_p/\omega_0 = 0.9$, however this includes infinitely many trailing secondary spikes. Since the box only has a finite size, only some of these trailing spikes are counted and so the measured efficiency of 0.8 is slightly less than η_{max}. However, the total efficiency is seen to be increasing and will converse to η_{max} as $t \to \infty$.

The plot also shows the efficiency to only the first peak η_1. It increases in the linear stage as the main spike grows, but then starts to decrease. This slight decrease is due to the energy transferred to the trailing secondary spikes as these

develop. The secondary peaks developed at the same time as η and η_1 began to diverge.

A high efficiency of $\eta_1 = 50\%$ is achieved. However, the plot shows that for these parameters only around 60% of the total energy of the amplified pulse resides inside the first peak and the rest is of less use since it resides in the undesirable trailing peaks. This fraction may be increased by using a seed pulse with a greater amplitude and duration [73]. The secondary peaks are an unavoidable product of the three wave interaction. The equations are those of an anharmonic oscillator, with the energy oscillating between pump, plasma and seed waves. However, the results of Chapter 5 will show how a different beam geometry can suppress these secondary peaks.

4.4 Particle-in-cell simulation

The envelope equations in the previous section make many assumptions and there are physical effects which are likely to violate these. The particle-in-cell simulation method includes more of the relevant physics and makes only the following assumptions:

- The collision timescale is longer than the interaction so $n\lambda_D^3 > 10$. The simulation is effectively collisionless.
- The plasma is preformed and no ionisation is required.

The particle-in-cell method fully models the velocity space of the particles, meaning all kinetic effects such as wave-breaking and Landau damping are included. In addition, the amplified pulse will forward scatter from noise in the simulation and could filament transversely in 2D or 3D simulations. Therefore all relevant physical effects are included, subject to the finite grid resolution and limited number of simulation particles. This is in contrast to methods making an envelope approximation to the laser fields, such as references [35, 133, 62, 34].

The code Osiris [38] was used to model the interaction using the particle-in-cell method. The simulation had one spatial dimension and three velocity dimensions, with periodic boundary conditions along the second and third spatial dimensions.

The initial conditions were equivalent to those of Fig. 4.4. The simulation used a cell size of $0.008\lambda_0$, with 256 electrons and 256 protons per cell. The simulation

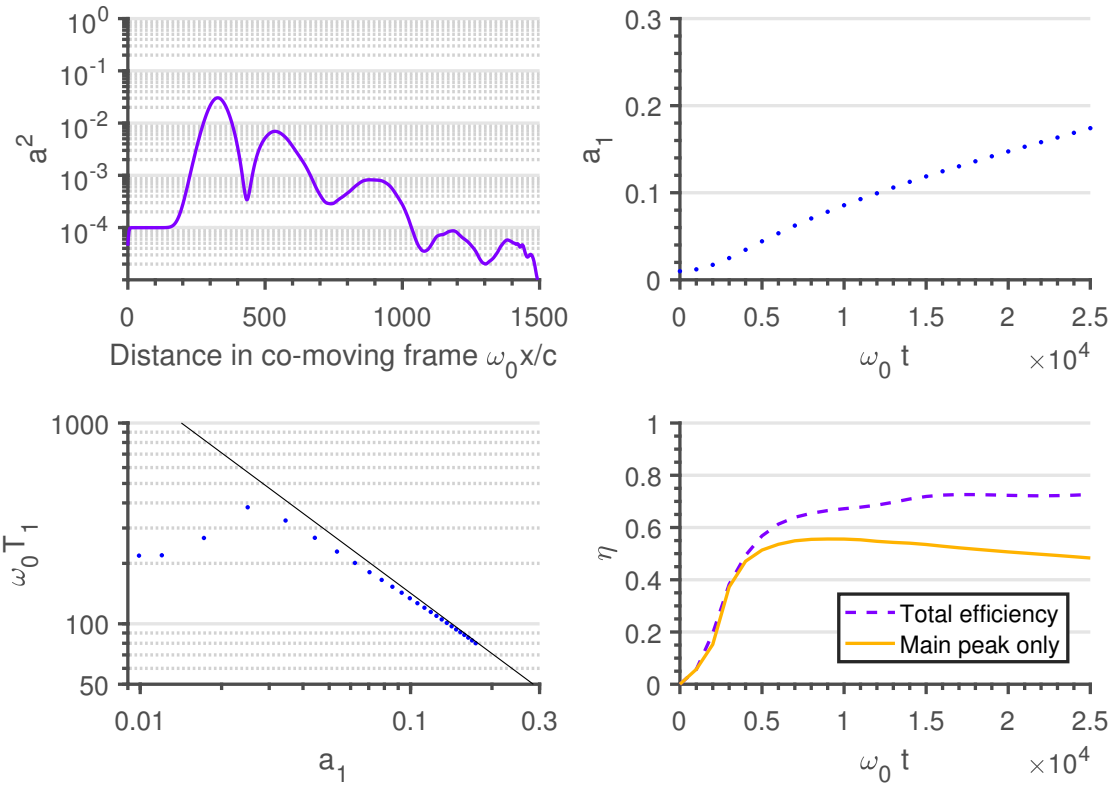


Figure 4.5: One-dimensional particle-in-cell simulation for density $n/n_{crit} = 0.01$ and pump amplitude $a_0 = 0.01$, with initial seed amplitude $a_1 = 0.01$ and duration $\omega_0 T_1 = 210$. (a) The profile of the electromagnetic energy density after a time $\omega_0 t = 25000$, smoothed with a Gaussian filter. The seed propagates to the left. The constant amplitude pump pulse is seen in the lower left. (b) The seed pulse amplitude as a function of time. (c) The seed pulse amplitude and FWHM duration, over the course of the interaction. Each data point is separated by a time $\omega_0 t = 1000$. The seed pulse is attracted to the analytical solution for late times (solid line), given by equation 2.44. (d) The total energy transfer efficiency to the whole of the amplified pulse (including secondary peaks), and to just the main peak of the amplified pulse.

box moved with the seed pulse, with new cells initialised at the front of the box. The particles were initialised with zero temperature and a velocity equivalent to the quiver velocity in the pump pulse. The boundary conditions at the rear were absorbing for the particles and fields, the same as in the numerical integration.

The results are shown in Fig. 4.5. The results are largely similar to the numerical integration shown in Fig. 4.4, proving the validity of the analytical model.

The simulation replicates the non-linear stage of the amplification and the characteristic secondary peaks are seen behind the main peak in Fig. 4.5a. The front peak has a similar final amplitude $a_1 = 0.18$ to the numerical integration. However, the secondary peaks are slightly suppressed.

The condition (2.44) accurately describes the later evolution of the pulse, as seen in Fig. 4.5c. Furthermore, the scalings for a_1 and T_1 agree with the analytical theory. Inspection of the ion velocity space shows that the distribution remains uniform and Maxwellian. Ion waves do not have time to develop within the short timescale of the seed pulse.

The seed pulse amplitude and duration were obtained by taking the total energy density in the electromagnetic field, applying a Gaussian filter to smooth oscillations on the wavelength scale, subtracting the constant energy density of the pump, then taking the peak and FWHM duration of the main spike to yield a_1^2 and T_1. The efficiency was then calculated using equation (4.1).

It is also possible to separate the seed and pump envelopes using Fourier analysis and spectral filtering. However, simply subtracting the pump was judged to be more accurate since the seed amplitude greatly exceeds that of the pump.

Panel (d) shows the efficiency to the first peak, as defined in the previous section. It also shows the total efficiency which is calculated by integrating the full energy density of the seed pulse, including any secondary peaks. For these parameters, the efficiency η is only slightly lower than that of the numerical integration in the previous section, suggesting that the kinetic effects are not too detrimental in this case and the analytical model is valid. The total efficiency monotonically rises to its limiting value, whereas the efficiency to the first peak decreases slightly at later times and converges to a lower value. This shows that some energy is transferred to the secondary peaks and this is unavoidable in the non-linear stage.

If the pump pulse amplitude is raised, the plasma wave amplitude will rise until it breaks, reducing the density perturbation and energy transfer. To demonstrate this, the particle-in-cell simulation was repeated with $a_0 = 0.03$, far above the wave-breaking threshold seen in Fig. 4.3, where the simulated density here is $1.1 \times 10^{19}\,\mathrm{cm}^{-3}$. Fig. 4.6 shows the results in this case. The efficiency has dropped remarkably compared to the case below the wave-breaking threshold. Fig. 4.6a shows that the pump pulse is only depleted by 60%. The plasma wave breaks before the pump is fully depleted, meaning that some energy remains in the pump pulse and the efficiency decreases. The amplified pulse also has more fine scale structure than the previous case. This is due to the amplified pulse Raman scattering in the forwards direction, which propagates with the pulse and grows to appreciable

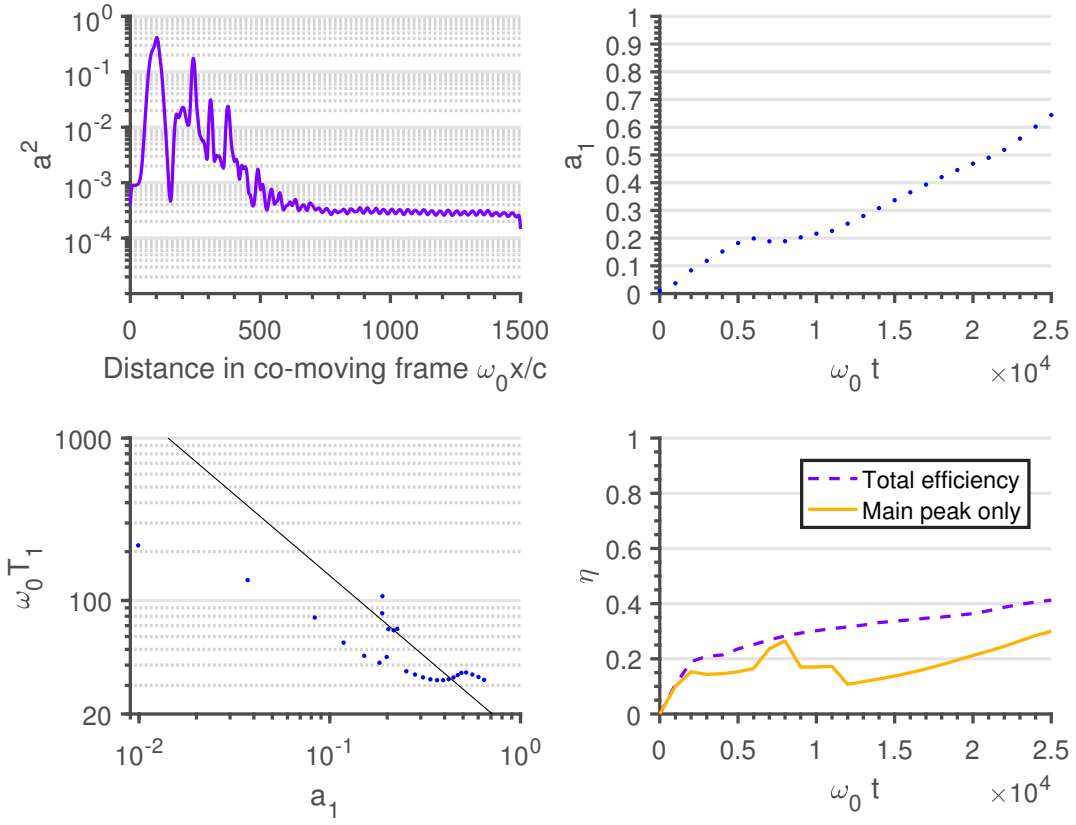


Figure 4.6: One-dimensional particle-in-cell simulation for density $n/n_{crit} = 0.01$ and pump amplitude $a_0 = 0.03$, with initial seed amplitude $a_1 = 0.01$ and duration $\omega_0 T_1 = 210$. (a) The profile of the electromagnetic energy density after a time $\omega_0 t = 25000$, smoothed with a Gaussian filter on a wavelength scale. The seed propagates to the left. (b) The seed pulse amplitude as a function of time. (c) The seed pulse amplitude and FWHM duration, over the course of the interaction. Each data point is separated by a time $\omega_0 t = 1000$. The analytical solution for late times is shown by the solid line and is not applicable in this case with large kinetic effects. (d) The total energy transfer efficiency to the whole of the amplified pulse (including secondary peaks), and to just the highest amplitude main peak of the amplified pulse.

levels. This modulates the main amplified pulse. The noise is incoherent and so the Raman forward scatter is also incoherent, degrading the pulse.

Panel (b) shows that the final peak amplitude becomes much greater than in the previous case, but this is at the expense of losing energy transfer efficiency and increasing the impact of competing instabilities. The development in the amplitude versus duration plane is shown in Fig. 4.6c. The analytical solution is not as applicable in this case, because of the increased kinetic effects [49]. The pulse is kept shorter than the the duration predicted by equation (2.44). This is expected on physical grounds because the plasma wave breaks at some point within the inter-

action and stops the energy transfer before the pump is fully depleted. This must shorten the amplified pulse below the duration expected with no wave-breaking.

In addition, there is competition between two of the amplified spikes. Initially the central part of the seed pulse is amplified, however some of the seed light in front of it grows to a similar level and depletes the pump, meaning the spike behind it grows more slowly. The spike in front reaches similar amplitude to the one behind at around $\omega_0 t = 8000$. This gives a sudden deviation to the pulse duration as it swaps from measuring one spike to the other, seen in panel (c).

At late times, the evolution seems to have a constant duration given by $\omega_0 T_1 \simeq 30$ as a_1 continues to grow. The inverse plasma frequency is $\omega_0/\omega_p = 10$, suggesting that the amplified pulse cannot shorten below this level.

4.5 Parameter scan

The simulated efficiency is less than η_{max} due to kinetic effects in the plasma wave and it will therefore be temperature dependent. This means the maximal efficiency in the non-linear stage depends on the plasma electron density n, the pump pulse amplitude a_0 and the electron temperature T. This is in contrast to the simplistic prediction of $\eta_{max} = 1 - \omega_p/\omega_0$, which only depends on density.

I undertook a full scale three-dimensional parameter scan to identify the efficiency across these three main parameters [107]. Each simulation had different values for the density, electron temperature and pump amplitude. The data were compiled from 1,200 separate particle-in-cell simulations, each with a cell size of $\lambda_0/100$ and 128 electron particles per cell. This gives a sufficient number of particles to accurately model the full Maxwellian distribution, including its tails. To reduce computational expense, the simulations used the moving simulation window as in the previous section, along with the same boundary conditions. The box moved at the seed pulse group velocity.

The efficiency, given by equation (4.1), was extracted after a propagation distance $L_{max}/5$. This is long enough for the efficiency to reach close to its asymptotic value. The seed pulse started with amplitude $a_1 = 0.1$ and a duration $\omega_0 T_1 = 50$, starting close to the non-linear stage of the amplification. The seed temporal profile had a $\cos^2(\pi t/2T_1)$ form, plotted in Fig. 4.7, and the seed pulse frequency was $\omega_0 - \omega_p$. The pump pulse amplitude was constant in each simulation. The plasma was uniform and fully ionised in each simulation and the ions were held static to

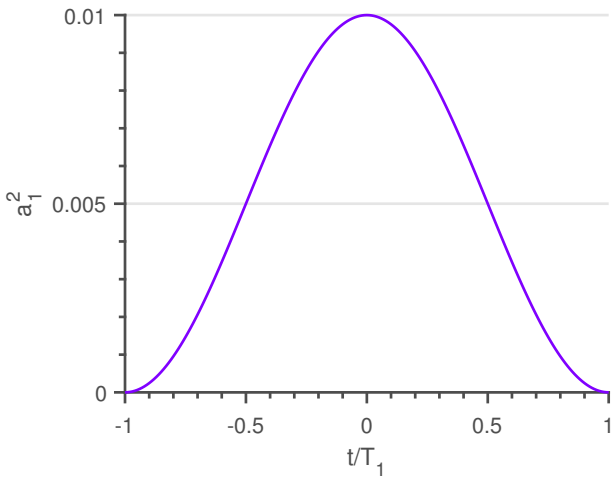


Figure 4.7: Temporal intensity envelope profile of the initial seed pulse used in the particle-in-cell simulations.

reduce computational costs. The laser polarisations were linear and aligned. Noise in the simulations was reduced by using cubic interpolation and current smoothing at each time-step.

Fig. 4.8a shows the convergent efficiency $\eta(n, a_0, T)$ for each simulation across the parameter space. Notice that the simulations in Fig. 4.5 ($n/n_{crit} = 0.01, a_0 = 0.01, T = 0, \eta = 0.75$) and Fig. 4.6 ($n/n_{crit} = 0.01, a_0 = 0.03, T = 0, \eta = 0.4$) agree with this large scale parameter scan. The range of density and amplitude was chosen to cover the estimated optimal range in Figs. 4.2 and 4.3.

In the zero temperature case in Fig. 4.8a, the analytical wave-breaking threshold for cold plasma (2.54) is also plotted. The region below this threshold has high efficiency, as expected. In this region, there is a slight decrease in efficiency at high densities, as predicted by η_{max}. Below the wave-breaking threshold, the efficiency is almost independent of pump amplitude and is slightly lower than η_{max}.

For amplitudes slightly above the wave-breaking threshold, by around 20%, the efficiency is unaffected. This shows that the plasma wave still retains some coherence after it breaks and the Raman scattering continues for some time. As the pump amplitude is increased further, the efficiency rapidly drops. Clearly an efficient Raman amplifier must not greatly exceed the wave-breaking threshold. The $a_{wb} \propto n^{0.75}$ dependence is also validated by the simulations.

The parameter scan was repeated with the electrons initialised as a Maxwellian distribution with temperature 100 eV. Fig. 4.8b shows the total efficiency for this temperature. Compared to the zero temperature case, the region of highest effi-

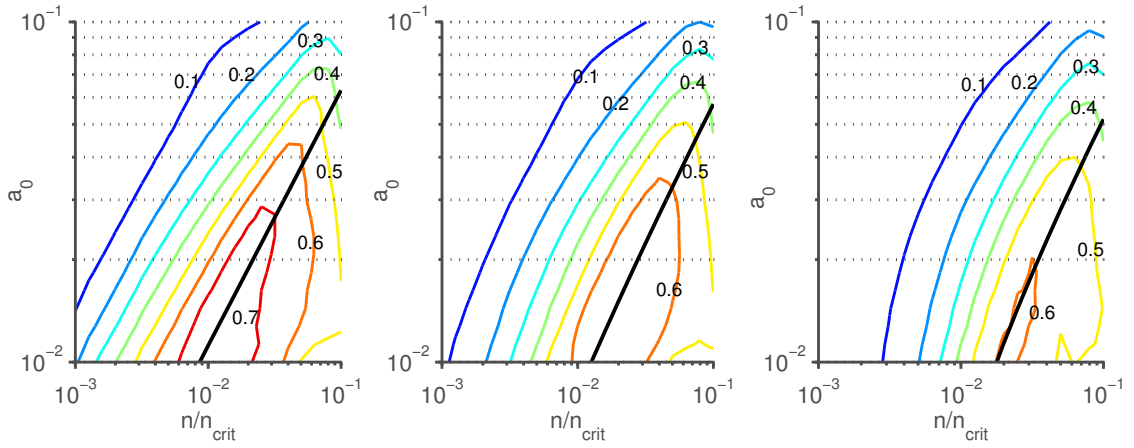


Figure 4.8: Simulated pump to seed energy efficiency of Raman amplifiers, as a function of plasma density n and the pump pulse dimensionless amplitude a_0. The three panels show three different electron temperatures of (a) 0 eV, (b) 100 eV, and (c) 400 eV. The contours give the efficiency, as found from 1200 one-dimensional particle-in-cell simulations. Each simulated a long propagation, such that the efficiency had converged to close to its limiting value. The energy efficiency is the total value, including any secondary peaks, as calculated by equation (4.1). Each simulation had constant density and pump pulse amplitude and started with a seed with $a_1 = 0.1$, $\omega_1 = \omega_0 - \omega_p$ and duration $\omega_0 T_1 = 50$. The solid black lines show the model for the temperature dependent wave-breaking threshold, given by equation (4.22)

ciency is shifted slightly towards higher density and lower pump amplitude. This agrees with the analytical estimates in Fig. 4.2 that as the temperature is raised, the optimal region will narrow and shift towards higher density. It is promising that even for these high electron temperatures typical of experimental pre-formed plasmas [29], efficiency of 60% is still achievable.

The wave-breaking threshold is seen to decrease as the plasma temperature is raised. This is because the electron velocity distribution widens, so more electrons exceed the phase velocity of the wave and become trapped. As a basic model, the sum of the plasma wave quiver velocity $\tilde{u}$ and the electron thermal velocity v_{th} must not exceed the wave phase velocity. The maximal quiver velocity of the wave was found in equation (2.52). The temperature dependent wave-breaking threshold becomes

$$v_p < \tilde{u} + v_{th} = a_0 v_p \left(\frac{2\omega_0}{\omega_p}\right)^{3/2} + \sqrt{\frac{k_B T}{m}}, \tag{4.21}$$

$$a_0 > \left(\frac{\omega_p}{2\omega_0}\right)^{\frac{3}{2}} \left[1 - \frac{2\omega_0}{\omega_p}\sqrt{\frac{k_B T}{mc^2}}\right]. \tag{4.22}$$

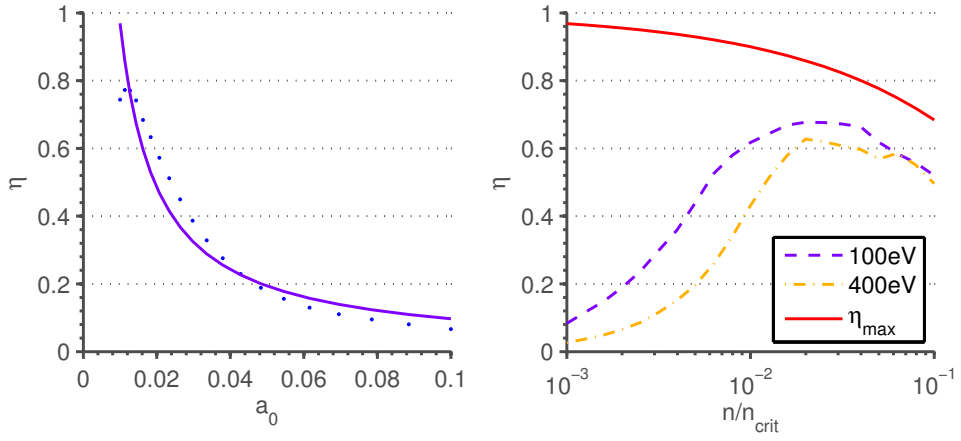


Figure 4.9: (a) Efficiency of the amplifier above the wave-breaking threshold. The plot shows a line-out from the data in Fig. 4.8a, for a constant density $n = n_{crit}/100$ and zero temperature. The blue dots show the simulated total efficiency. The line is a fit of the form $\eta \propto a_0^{-1}$. (b) The maximum available efficiency for a given plasma density. The efficiency η_{max} with no kinetic effects is compared to the simulated efficiency at electron temperatures 100 eV and 400 eV. Kinetic effects decrease η, especially at lower densities.

This prediction is plotted with the solid lines in Fig. 4.8. There is an alternative model given in reference [14] which is significantly more stringent. For each temperature, the model (4.22) more accurately describes the regions with highest simulated efficiency. This discrepancy is likely because the model in reference [14] used a water-bag distribution, rather than the Maxwellian distribution in these simulations.

As the temperature is increased still further to 400 eV in Fig. 4.8c, the optimal region shifts to higher density and lower pump amplitude. The optimal region is narrow and heavily temperature dependent. This suggests that careful control and diagnosis of the electron temperature is important in Raman amplification experiments. In addition, the peak available efficiency decreases for higher temperatures.

For the case at 400 eV, the contours of efficiency begin to curve in the lower left part of the plot. This is due to the presence of Landau damping as well as wave-breaking. Using the wavenumber of the plasma wave, the condition $k\lambda_D < 0.5$ leads to $n/n_{crit} > 16k_BT/(mc^2)$ or $n/n_{crit} > 0.013$ for the 400 eV case. This provides a limitation on the plasma density that restricts the high efficiency region in Fig. 4.8c.

The drop in efficiency above the wave-breaking threshold is shown in Fig. 4.9a, for a fixed density of $n = n_{crit}/100$ and zero temperature. The efficiency drop is well

fitted by the functional form a_0^{-1}, suggesting the loss of efficiency is not as bad as suggested by the simplistic estimate of a_0^{-2}. This is because the plasma wave retains some coherence even after it breaks and some energy transfer continues. A least squares regression of $\log(\eta)$ versus $\log(a_0)$ found the best fit to have a relationship $\eta \propto a_0^{-1.14}$. Since the exponent is greater than -2, this supports the hypothesis that the plasma wave retains some coherence after it breaks and still scatters some of the pump pulse [33].

The detrimental impact of the kinetic processes is clear in Fig. 4.9b, which plots the maximal possible simulated efficiency for a given plasma density. Especially at low plasma densities where the wave phase velocity is low, the kinetic effects decrease the efficiency below that of the prediction η_{max} from the coupled envelope equations [137]. The combination of this, along with decreasing η_{max} at higher density, leads to a maximum in the efficiency of 0.6 at a plasma density $n = 0.02 n_{crit}$.

Since the scope of the parameter scan did not include $a_0 < 0.01$, it may be possible to access lower plasma density by simultaneously lowering the pump pulse amplitude, as demonstrated in the later parts of this chapter.

Fig. 4.10 shows the same parameter space plot but with the efficiency calculated as η_1, the efficiency to the highest intensity peak only. Since some energy will also be transferred to the secondary peaks, this is less than the total efficiency plotted in Fig. 4.8. It shows the same general trends as the total efficiency, decreasing rapidly above the wave-breaking threshold. The model for the temperature dependent wave-breaking accurately predicts the region of highest efficiency.

The wave-breaking also has a greater effect on the secondary peaks. Comparing figures 4.10 and 4.8, η is much higher than η_1 in the region below the wave-breaking threshold. In this region, $\eta_1 \simeq 0.7\eta$, meaning that much of the energy is in the secondary peaks of the peaks of the wave train.

However, at 2-3 times the wave-breaking threshold, comparing Figs. 4.8 and 4.10 shows that $\eta_1 \simeq 0.9\eta$. This means a larger fraction of the amplified pulse energy resides in the main peak. This is because the plasma wave breaks somewhere within the main peak and so the trailing peaks are barely amplified. This has the effect of suppressing any trailing secondary peaks [34].

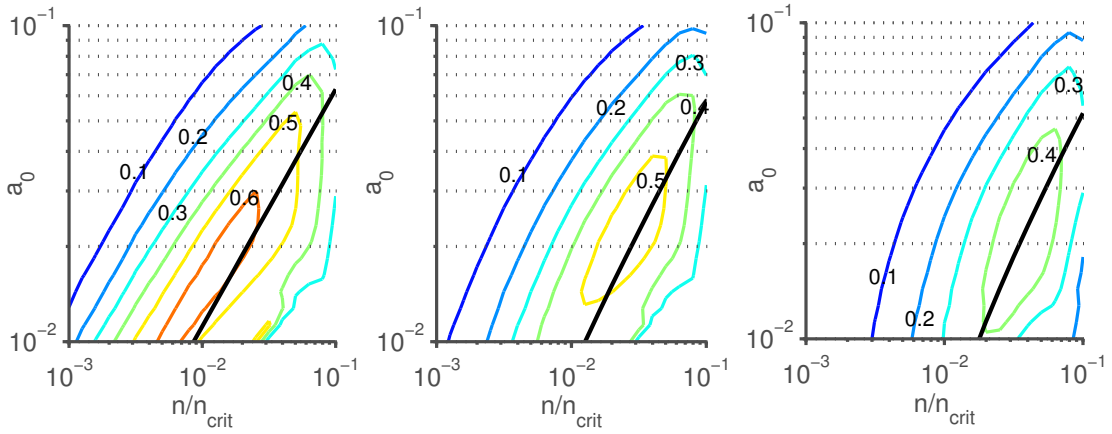


Figure 4.10: Simulated pump to seed energy transfer efficiency of Raman amplifiers, as a function of plasma density n and the pump pulse dimensionless amplitude a_0. The three panels show three different electron temperatures of (a) 0 eV, (b) 100 eV, and (c) 400 eV. The contours give the energy transfer efficiency to the highest amplitude peak only, as found from 1200 one-dimensional particle-in-cell simulations. Each simulated a long propagation, such that the efficiency had converged to close to its limiting value. Each simulation had constant density and pump pulse amplitude and started with a seed with $a_1 = 0.01$ and $\omega_0 T_1 = 50$. The solid black lines shows the model for the temperature dependent wave-breaking threshold, given by equation (4.22).

4.5.1 Discussion of errors

There are several possible sources of error in this parameter scan. First, the propagation distance of $L_{max}/5$ may not be far enough for the efficiency to converge. Also the simulation box had a finite length and so the energy in some of the trailing peaks is not counted.

The effect of numerical instabilities must also be considered. Particle-in-cell simulations are subject to artificial numerical heating of the particles because of the grid interpolation. This is especially prevalent when the cell size exceeds the Debye length. For this reason the simulations in Fig. 4.8b c resolved λ_D. Even in the zero temperature case panel (a), numerical heating was minimal because the new plasma is initialised at the front of the moving window, less than one plasma period in front of the main seed pulse. This is another important advantage of moving window simulations. Conservation of energy in the simulation was verified by inspecting the particle phase space for a repeat of a simulation with no laser pulses.

Particle-in-cell simulations may also over estimate the probability of a particle

becoming trapped in the plasma wave. This could affect the breaking of the plasma wave and change the efficiency. This may be alleviated by increasing the grid resolution, so the parameter scan used a relatively high resolution of $\lambda_0/100$. Since the simulations accurately reproduce the theoretical wave-breaking threshold, this shows the kinetic effects are well modelled, at least in the region immediately where the wave breaks. By the time the plasma wave breaks, the main peak will have passed, so any errors due to finite resolution will only affect the trailing secondary peaks.

All of these errors would tend to decrease the simulated value of η, meaning the simulated values are likely a lower bound. Since η_{max} is a well understood theoretical upper bound, this provides an additional constraint.

I now compare these simulations to a previous experiment. The results of Ren *et al.* [98] used plasma density $n = 1.4 \times 10^{19}\,\mathrm{cm}^{-3}$ and pump intensity $I = 2.3 \times 10^{14}\,\mathrm{Wcm}^{-2}$, giving $a_0 = 0.010$ and $n = 0.008 n_{crit}$. These are similar to the conditions simulated in Fig. 4.5, although that simulation used a cold plasma. There are no indications of experimental temperature measurements, but estimating the temperature of the pre-formed plasma as 100 eV, the parameter scan gives a simulated maximal efficiency of 50%. The experiment measured 6.4% efficiency, the most successful experiment to date.

There are two explanations for this discrepancy. First, as I will show in the next section, the experimental seed pulse was not intense and short enough to reach the high efficiency non-linear stage from the start. The seed pulse was relatively long at 500 fs, and low intensity, with $\Delta u = 0.07$. Although eventually the efficiency should converge to the simulated 50% value, there will have been a large section of the amplifier before the non-linear stage was reached. This reduces the overall efficiency. Secondly, the interferometric measurements of the plasma profile show it did not have the required uniformity, as explored here in Chapter 6.

The next section suggests that use of a shorter, more intense seed pulse would reach the non-linear stage more quickly and increase the performance.

4.6 Optimal seed pulse

4.6.1 Moving window simulations

Since the total efficiency monotonically increases through the interaction, it is important to consider over what timescale it approaches the limiting value given in the parameter scan. This will depend on the properties of the initial seed pulse, such as its duration, peak intensity and contrast.

To study the effects of different seed pulses, I conducted 3 simulations, each with a different seed pulse. I fixed the pump pulse and plasma parameters to $a_0 = 0.01$, $n = 0.01 n_{crit}$ and $T = 0\,\mathrm{eV}$. For these conditions, the parameter scan gave a limiting efficiency of 0.75. The plasma density was uniform in each simulation, with cell size $\lambda_0/60$ and 64 particles per cell.

The temporal intensity envelope of the seed again had the form $a_1^2 \cos^2(\pi t/2T_1)$, where the initial amplitude a_1 and duration T_1 were varied. Seed 1 had $a_1 = 0.0045$ and $T_1 = 50\omega_0^{-1}$. Seed 2 had the same amplitude but a much longer duration of $T_1 = 1000\omega_0^{-1}$. Seed 3 had the same duration as seed 1 but a much higher amplitude of $a_1 = 0.045$. As such, seeds 2 and 3 were close to meeting the condition for the non-linear stage from the start.

To reflect a realistic experimental seed from a chirped pulse amplifier, the initial seed pulse spectral bandwidths were kept constant at $\Delta\omega/\omega = 0.05$ by chirping the initial seeds. The central frequency was $\omega_0 - \omega_p$ in each case.

Fig. 4.11 shows the development in the a_1, T_1 plane of the three seed pulses. Each line represents the trajectory of a separate seed pulse over a propagation time of $\omega_0 t = 25000$. The condition to reach the non-linear stage is shown by the straight solid black line in Fig. 4.11.

The first obvious conclusion is that the analytical solution is valid at late times, since it describes the later evolution of all three seed pulses. This shows that the solution is an attractor.

Seed 1 initially lengthens to reach the non-linear stage, then once it meets equation (2.44) it shortens again, in agreement with reference [53]. It reaches a final amplitude of 0.133.

Seed 2 started with the same amplitude but longer duration than seed 1. It exhibited a different behaviour, where it shortened and split in to several competing spikes. Eventually the leading spike met the non-linear stage criterion, but the

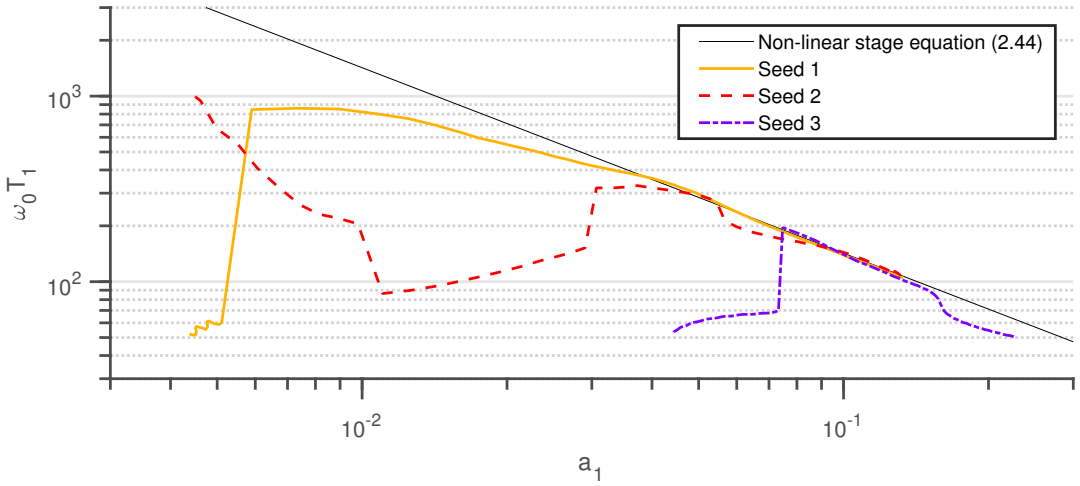


Figure 4.11: The development of 3 separate seed pulses in the a_1, T_1 plane, each shown in a different colour. The seeds were varied in initial duration and amplitude, keeping their bandwidth constant. The plasma was uniform with a density $0.01 n_{crit}$ and zero temperature. The pump pulse amplitude was $a_0 = 0.01$. All seed pulses are attracted to the analytical non-linear solution at late times, given by equation (2.44) and shown by the straight black line. The duration is found from the full width at half maximum of the highest amplitude spike.

effectiveness of the seed was compromised due to the competing peaks. This means that despite starting with twenty times as much energy, seed 2 finished the amplifier with almost identical amplitude to seed 1. In addition, the evolution of the pulse is complex and unpredictable, meaning there would be more shot to shot variation when using a longer seed pulse.

Seed 3 starts with a much higher amplitude than seed 1, but exhibits the same behaviour where it lengthens to meet the non-linear condition, before shortening again and increasing in amplitude. Increasing the initial seed amplitude has resulted in a greater pump depletion and a greater output amplitude of $a_1 = 0.23$. This shows that higher amplitude seed pulses have a greater performance.

The conclusion is that performance can be improved by increasing the seed pulse amplitude at constant duration. Increasing the duration at constant amplitude has little effect on the output pulse, or even worsens the performance.

One of the most interesting features in Fig. 4.11 is that for long duration seeds, the analytical solution shown by the black line starts to break down. For example, seed 2 initially starts near the black line but then deviates and later rejoins with a much shorter duration. This shows that the black line is only an attractor in the lower region of the plot.

This can be traced back to the derivation given in Chapter 2, whereby the self similar solution was found by neglecting the time derivatives of the envelopes. In effect, the assumption is that the seed pulse duration is much less than its timescale of evolution. Estimating the timescale of evolution as the exponential growth time in the linear stage, I find that the self-similar analytical solution is only valid when $T_1 < \gamma_{RBS}^{-1} = \frac{2}{a_0\sqrt{\omega_0\omega_p}}$. For the conditions simulated, this gives $\omega_0 T_1 < 630$. This matches with Fig. 4.11, since the seeds are only attracted to the black line for durations shorter than $\omega_0 T_1 < 500$.

4.6.2 Fixed window simulations

The simulations in the previous section used a moving simulation window. This could introduce an artificial seeding mechanism via the initialisation of new plasma at the front of the box. To eliminate this source of error, several simulations were repeated using a fixed simulation window. This models the evolution of the whole plasma and the full propagation of the pump through the plasma, a more realistic set-up.

A chirped pulse amplification system will provide a seed pulse that has a given fluence and bandwidth, whereas the duration is variable. It is an important question for experimental design to know whether a long, low power seed, or a shorter transform limited seed, is more effective [138, 61]. At constant fluence, the seed can always be lengthened to meet the non-linear condition (2.44). However, this leads to a very long, low intensity seed. I now show that a shorter, more intense seed of the same fluence is more effective, despite starting further from the non-linear condition.

To investigate experimentally relevant constraints, three seeds of equal initial fluence (constant $a_1^2 T_1$) and bandwidth were simulated, again in one dimension. The electrons were initialised with density $n = 0.01 n_{crit}$ and temperature 100 eV, the ions were held fixed and the cell size was $\lambda_0/60$ with 128 electrons per cell. The pump pulse amplitude was constant at $a_0 = 0.02$.

The plasma was initialised with a length of 1 mm across the central region of a box of length 2 mm. The seed pulse was initialised once the pump had crossed the full plasma. The longest seed initially met the non-linear condition (2.44) and the shorter ones were below it.

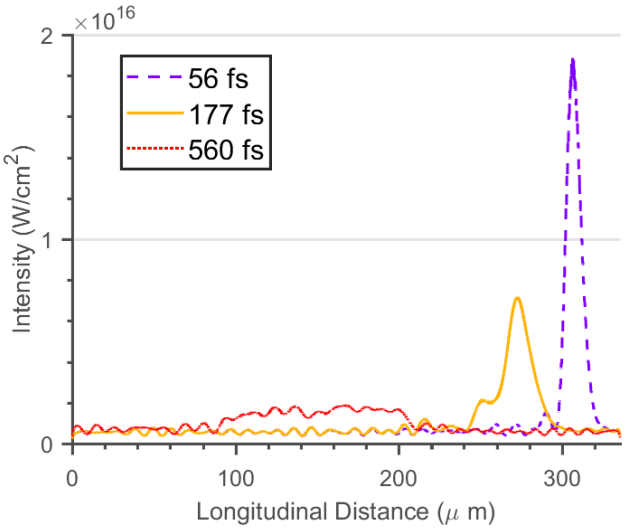


Figure 4.12: Results of three fixed window one-dimensional PIC simulations, each with a different seed pulse. These seeds were varied in duration (given in the legend) at constant energy and bandwidth and propagated for a distance 1 mm through plasma with a uniform density $10^{19}\,\mathrm{cm}^{-3}$ and temperature 100 eV. The counter-propagating pump pulse (also shown in the plot) had constant intensity $5 \times 10^{14}\,\mathrm{Wcm}^{-2}$ and wavelength 1054 nm. The seeds had a chirp to maintain constant bandwidth $\Delta\omega/\omega_1 = 0.03$ in each case, all with initial fluence $130\,\mathrm{Jcm}^{-2}$. Their central wavelength was 1171 nm. Their initial intensities were $23, 7.3$ and $2.3 \times 10^{14}\,\mathrm{Wcm}^{-2}$ respectively. The amplified seeds are shown after they have left the plasma, propagate to the right and their intensity has been smoothed on a wavelength scale.

Fig. 4.12 shows the longitudinal intensity envelope of the final amplified seeds. Importantly, the final amplitude and fluence reached by the shortest seed is much greater than the longest seed, despite starting much further from the analytical solution.

Closer inspection of the simulations found that starting with a short seed will always stimulate a single amplified spike, whereas seeds longer than γ_{RBS}^{-1} tend to split into several amplified spikes, reducing the final peak intensity [124]. For the conditions simulated, $\gamma_{RBS}^{-1} = 180\,\mathrm{fs}$. The seeds shorter than this have led to a single amplified spike and qualitatively looks more like the analytical solution described in Chapter 2.

Splitting up of a longer seed pulse will result in several spikes of similar amplitude, with varying duration and separation. The more useful amplified pulse profile for many applications is that of a single consistent, high amplitude peak, since that will reach the highest intensity at focus. I therefore conclude that a shorter, more intense seed pulse is preferable to a longer pulse that meets the criterion (2.44).

This conclusion is for seed pulses of fixed fluence and bandwidth.

The condition (2.44) for accessing the high efficiency should therefore be supplemented with the condition $T_1 < \gamma_{RBS}^{-1}$. Inserting $T_1 = \gamma_{RBS}^{-1}$ in the equation (2.44) leads to a minimum required seed pulse amplitude to reach the non-linear stage, given by

$$a_1 = \frac{5}{2} a_0. \tag{4.23}$$

Therefore for the self-similar high efficiency behaviour, the seed pulse intensity must exceed the pump pulse intensity.

This also gives a minimum required seed pulse fluence to reach the non-linear stage from the start, given by

$$\Phi_1 = \omega_1^2 a_1^2 T_1 = \frac{25 a_0 \omega_1^2}{2\sqrt{\omega_0 \omega_p}}. \tag{4.24}$$

For a 1μm pump pulse with intensity $10^{14}\,\mathrm{Wcm^{-2}}$ in a plasma of density $10^{19}\,\mathrm{cm^{-3}}$, this gives a minimum seed pulse fluence of $200\,\mathrm{Jcm^{-2}}$ to reach the non-linear stage. The required intensity is $6 \times 10^{14}\,\mathrm{Wcm^{-2}}$. These values are far above damage thresholds of conventional optics, so any amplifier that wishes to reach the non-linear stage from the start must focus the initial seed pulse.

In conclusion, for a given fluence and bandwidth, the shortest seed pulses have higher performance. The seed should certainly be shorter than γ_{RBS}^{-1} to avoid collapsing into several amplified spikes. The simulations showed that shorter seeds reach much higher peak intensity and have more reproducible behaviour. This is because they monotonically lengthen to meet the attractor solution, whereas seeds that are too long must shorten, typically by splitting into several amplified spikes.

For many previous experiments, the initial seed energy was very low and often relatively long seed pulses were used. See, for example, reference [98]. Although this experiment holds the record for the greatest energy transfer efficiency of 6.4% from pump to seed, the seed pulse duration of 500 fs exceeded the Raman growth time.

It would be of significant interest to fully measure the amplified pulse temporal profile in experiments such as this. This could be possible using a FROG diagnostic, described in Chapter 3. Measurement of the characteristic pulse envelope, with the secondary spikes, should also be possible.

4.7 Effect of competing instabilities

The effect of competing plasma instabilities cannot be neglected, since they will dissipate energy and destroy the amplified pulse coherence. I now estimate the growth of the main two competing instabilities, which are ponderomotive filamentation and Raman forward scatter of the amplified pulse.

The forwards scatter will propagate with the amplified pulse and continue to grow and modulate it [77]. Since the forward scatter grows from noise, it will be incoherent and degrade the coherence of the amplified pulse.

To transfer a significant amount of energy, the beam diameter must be wide and the power will greatly exceed the critical power for self focussing. Ponderomotive filamentation is a transverse instability of the amplified pulse that will cause it to lose transverse coherence and spawn high intensity filaments. Both of these instabilities will cause the pulse to lose coherence and so their number of exponentiations must be kept low.

The effect of Raman forward scatter may be calculated as follows. The seed pulse amplitude as a function of time is given by equation (4.19). Since the competing instabilities grow from low amplitude noise, the linear stage exponential growth rates (2.33) are applicable. As the pulse amplifies, its forward scatter component will grow from noise and propagate with it, modulating the pulse and causing its spectrum to grow side-bands. The total number of exponentiations of the Raman forward scatter is

$$\Lambda_{RFS} = \int \gamma_{RFS}\, dt = \int \frac{\omega_p}{4}\sqrt{\frac{\omega_p}{\omega_0}} a_1(t)\, dt \tag{4.25}$$

$$= \int_0^t \frac{\omega_p}{4}\sqrt{\frac{\omega_p}{\omega_0}}\frac{2}{5}\sqrt{\omega_0\omega_p}\eta a_0^2 t'\, dt' = \frac{\eta a_0^2}{20}\omega_p^2 t^2. \tag{4.26}$$

In this equation t is measured from the time that the process enters the non-linear stage. Substituting back in for a_1, one finds the maximum achievable amplitude is

$$a_1 = a_0\sqrt{\frac{80\omega_0\eta\Lambda_{RFS}}{25\omega_p}} \tag{4.27}$$

The same analysis is possible for the case of ponderomotive filamentation. The

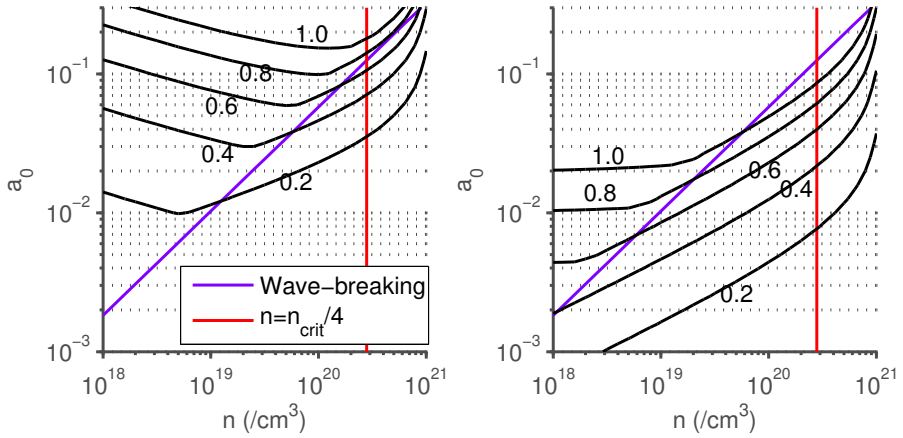


Figure 4.13: Contours of the maximum achievable amplitude a_1 of the amplified pulse before competing instabilities degrade it (black). This is estimated from the theoretical growth rates for the cases of (a) Raman forward scattering of the seed and (b) ponderomotive filamentation of the seed. n is the plasma electron density and a_0 is the pump pulse dimensionless amplitude. The plots assume a pump wavelength of 1 μm. Also shown are the cold plasma wave-breaking threshold (purple) and the quarter critical density (red), repeated from Fig. 4.3.

growth rate is given in equation (1.101). The total number of exponentiations is

$$\Lambda_{fil} = \int \gamma_{fil}\, dt = \int \frac{\omega_p^2}{8\omega_0} a_1^2\, dt \tag{4.28}$$

$$= \int_0^t \frac{\omega_p^2}{8\omega_0}\frac{4}{25}\omega_0\omega_p\eta^2 a_0^4 t'^2\, dt' \tag{4.29}$$

$$= \frac{\omega_p^3}{150}\eta^2 a_0^4 t^3. \tag{4.30}$$

Replacing t with equation (4.19), the maximum achievable amplitude is

$$a_1 = \left(\frac{48\eta a_0^2 \Lambda_{fil}}{5}\right)^{1/3}\left(\frac{\omega_0}{\omega_p}\right)^{\frac{1}{2}}. \tag{4.31}$$

The maximum tolerable growth of the competing instabilities was previously estimated [121] as $\Lambda_{fil} = 10$ and $\Lambda_{RFS} = 10$. The maximum achievable pulse amplitude is plotted on the parameter space diagram in Fig. 4.13. The efficiency $\eta(n, a_0)$ appears in these expressions and it was taken from a fit to the Fig. 4.8a. The contours show the maximum dimensionless amplitude that the amplifier can reach before $\Lambda_{RFS} = 10$ or $\Lambda_{fil} = 10$. It is clear that this increases with the pump pulse amplitude, a somewhat counter-intuitive result. Although the instability growth rates increase for higher a_0, the required length of the amplifier decreases proportional to a_0^{-2} and so the total number of exponentiations is actually less.

This highlights an interesting contradiction with Raman amplifiers. Higher energy transfer efficiency requires high density and low pump amplitude. However, maximising the output intensity before competing instabilities take hold requires the exact opposite conditions.

High efficiency is achieved for more benign conditions when the plasma wave does not break. However, this leads to a slower energy transfer which requires a much longer plasma to reach relativistic intensity. This increases the duration over which competing instabilities can grow and limits the achievable amplitude.

Across much of the parameter space, the length of the amplifier is limited by the competing instabilities, rather than reaching $a_1 = 1$.

Comparing figures 4.13a and 4.13b, the seed pulse is likely to degrade through Raman forward scattering before it degrades transversely. Of course, the Raman scatter growth rate used in the calculation applies to the linear stage and may not apply to relativistic amplified pulses. It also does not include any thermal effects. In addition, any plasma density gradients could de-tune and mitigate the spontaneous scatter while preserving the main pulse [72]. This could alter the plot in Fig. 4.13a to allow greater amplitudes to be reached than the basic estimate given here.

I conclude that greater output pulse intensity should be possible with higher pump intensity and shorter plasmas. However, this is at the cost of lower efficiency, so that a lot more pump pulse energy will be required for only small gains in output pulse intensity. The analysis here also neglects the effect of other plasma instabilities and assumes the simulated efficiency from the previous section is valid.

4.7.1 Maximum compression ratio and amplification factor

The limitation of competing instabilities is clearer if one calculates the maximum compression ratio. This is the ratio T_0/T_1, where T_0 is the pump pulse duration and T_1 is the final seed pulse duration. The two pulses overlap across a distance $cT_0/2$, so this is the optimal plasma length.

In the non-linear stage, a_1 increases proportional to time. The total propagation

time for the seed pulse is $T_0/2$, so the integrated forward Raman growth rate is

$$\Lambda_{RFS} = \int \gamma_{RFS}\, dt = \int_0^{T_0/2} \frac{\omega_p}{4}\sqrt{\frac{\omega_p}{\omega_0}} a_1\, dt \tag{4.32}$$

$$= \frac{\omega_p}{4}\sqrt{\frac{\omega_p}{\omega_0}}\frac{T_0}{4} a_1 = \frac{\omega_p}{4}\sqrt{\frac{\omega_p}{\omega_0}}\frac{T_0}{4}\frac{5}{T_1\sqrt{\omega_0\omega_p}}, \tag{4.33}$$

$$\frac{T_0}{T_1} \simeq 3\Lambda_{RFS}\sqrt{\frac{n_{crit}}{n}}. \tag{4.34}$$

Here I have used the non-linear condition (2.44) and assumed that the non-linear stage is reached from the start. The maximum compression ratio is independent of the pump pulse parameters and depends only on the plasma density.

If there is full overlap of the pump and seed pulses, the energy transfer is $E_1 = \eta_1 E_0$. This leads to the maximum power amplification

$$\frac{P_1}{P_0} = 3\eta_1\Lambda_{RFS}\sqrt{\frac{n_{crit}}{n}}. \tag{4.35}$$

This expression agrees with the Fig. 4.13 in that greater pump pulse intensity leads to greater possible output intensity. However, η_1 decreases above the wave-breaking threshold and so this is not a linear relationship. The previous section found that $\eta \propto a_0^{-1}$ above the wave-breaking threshold, leading to the maximum power output $P_1 \propto \sqrt{P_0}$ for operation above the wave-breaking threshold.

For an example, I take the maximum tolerable number of exponentiations as $\Lambda_{RFS} = 10$. A 25 ps, 1 μm pump pulse in a plasma density $10^{19}\,\mathrm{cm}^{-3}$ has a maximum allowable compression ratio of 320. The maximum power amplification is this multiplied by the efficiency found in the previous section. This yields the maximum power amplification $P_1/P_0 \simeq 150$. In reality, this estimate could be somewhat improved since γ_{RFS} is reduced via de-tuning and thermal effects.

The same growth rate calculation for ponderomotive filamentation leads to

$$\Lambda_{fil} = \int_0^{T_0/2} \gamma_{fil}\, dt = \int_0^{T_0/2} \frac{a_1^2\omega_p^2}{8\omega_0}\, dt \tag{4.36}$$

$$= \frac{\omega_p^2}{8\omega_0}\frac{a_1^2}{3}\frac{T_0}{2} = \frac{\omega_p^2}{8\omega_0}\frac{T_0}{6}\frac{25}{T_1^2\omega_0\omega_p}, \tag{4.37}$$

$$\frac{T_0}{T_1} = \sqrt{2\Lambda_{fil}\omega_0 T_0}\left(\frac{n_{crit}}{n}\right)^{1/4}. \tag{4.38}$$

Again I have used the non-linear condition (2.44). The achievable compression ratio improves for longer pump pulses and sparser plasmas. This is promising for the compression of longer nanosecond pulses such as those of the National Ignition

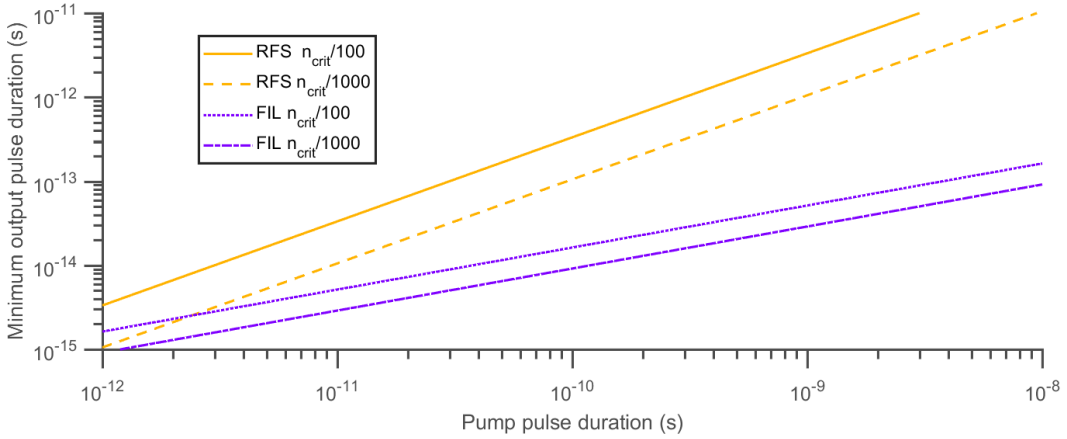


Figure 4.14: Minimum possible amplified pulse duration before competing instabilities degrade the pulse, for a 1 μm pump pulse of varying duration. This is shown for the case of Raman forward scattering and ponderomotive filamentation of the amplified pulse, estimated using the analytical growth rates. The solid lines are for plasma density $n_{crit}/100$ or $10^{19}\,\mathrm{cm}^{-3}$ and the dashed lines are for $n_{crit}/1000$ or $10^{18}\mathrm{cm}^{-3}$.

Facility, as suggested in reference [122]. The minimum achievable output pulse duration is shown in Fig. 4.14 for a pump of wavelength 1 μm.

By far the most stringent limitation is the premature backward scattering of the pump pulse. The pump must traverse across the plasma before it meets the seed pulse and Raman backward scatter will grow from noise. This will deplete the pump pulse and waste energy. The premature backscatter also creates a long duration pre-pulse before the main seed pulse, reducing its pulse contrast. This is usually undesirable for many applications in high energy density physics.

This effect can be estimated using the backward scatter growth rate (2.33), assuming the pump pulse amplitude and wavelength are constant. The number of exponentiations is

$$\Lambda_{RFS} = \int_0^{T_0/2} \gamma_{RFS}\, dt = \int_0^{T_0/2} \frac{a_0\sqrt{\omega_0\omega_p}}{2}\, dt = \frac{a_0 T_0\sqrt{\omega_0\omega_p}}{4}. \tag{4.39}$$

This can be related to the seed pulse duration by combining equations (2.44) and (4.18) to yield

$$a_0^2 T_0 T_1 \omega_0 \omega_p \eta_1 = 25. \tag{4.40}$$

Eliminating the pump amplitude in these equations leads to the final result

$$\frac{T_0}{T_1} = \frac{16}{25}\Lambda_{RFS}^2 \eta_1. \tag{4.41}$$

Taking the maximum tolerable number of exponentiations as 10 means the compression ratio cannot exceed approximately 30. This is much more stringent than the limitations due to seed pulse instabilities.

Fortunately, a high compression ratio can be recovered by several schemes. The resonance can be externally de-tuned with a pump chirp or plasma density gradient, as suggested in reference [72]. The problem may be completely avoided by using the seed ionisation scheme, such that the pump only traverses across neutral gas until it meets the seed [11].

The requirement to maximise the compression ratio as well as the efficiency leads to a unique solution with only one free parameter. To see this, I substitute (2.44) and the Raman forward scatter limitation (4.34) into the energy balance equation (4.18) to yield

$$\frac{25\omega_1^2}{T_1\omega_0\omega_p} = \eta_1\omega_0^2 a_0^2 T_0, \tag{4.42}$$

$$a_0 T_0 = \sqrt{\frac{80\Lambda_{RFS}}{\eta_1}}\frac{\omega_0 - \omega_p}{\omega_0\omega_p}. \tag{4.43}$$

For a given pump duration T_0 the compression ratio will be maximised by using this pump amplitude. If the pump amplitude exceeds this, Raman forward scatter of the seed will degrade it before the end of the amplifier. If the pump amplitude is less than this value, the seed will remain coherent but the compression ratio and power amplification will be less. The optimal performance is to continue the amplification to the point where the forward scatter just starts to degrade the pulse, given by equation (4.43). This requirement, along with the wave-breaking threshold, give two constraints on the parameters.

Fig. 4.15 shows these two requirements together. The efficiency is maximised along the wave-breaking threshold, shown by the purple line. The compression ratio before competing instabilities degrade the pulse is maximised along the blue lines. This is shown for various pump pulse durations. Parameters above this line will lead to unacceptable degradation of the amplified pulse.

This plot shows that for a given pump duration, there is a unique set of parameters that simultaneously maximise the energy efficiency and compression ratio of the Raman amplifier.

The plot assumes $\Lambda_{RFS} = 10$. This was a value extracted from particle-in-cell simulations by Trines *et al.* [121]. Due to unrealistic noise in the simulation, the

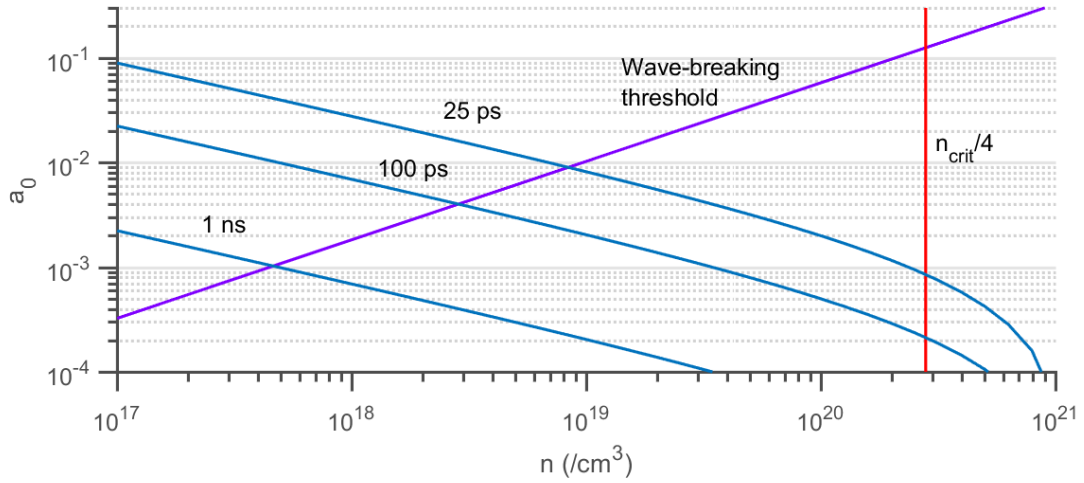


Figure 4.15: Parameter space for Raman amplification, showing the cold plasma wave-breaking threshold in purple, along which the efficiency is maximal. The blue lines show the relation (4.43), for various pump pulse durations, along which the compression ratio is maximal. Exceeding these contours may mean the amplified pulse suffers loss of coherence and envelope degradation due to the competing instabilities. These contours were estimated from the linear Raman forward scatter growth rate, however the true growth rate may be less.

tolerable experimental value will depend on the experimental set-up and may be greater than 10. Note also that the growth rate γ_{RFS} is the maximum theoretical value. Reduction of the forward scatter growth rate through de-tuning or thermal effects will tend to increase the allowable pump amplitude and plasma density.

For example, the optimal parameters for a 25 ps, 1 μm wavelength pump pulse are $a_0 = 0.01$, giving $I_0 = 10^{14}\,\mathrm{W/cm^2}$. The plasma density at the wave-breaking threshold is approximately $n = 10^{19}\,\mathrm{cm^{-3}}$.

For a longer pump pulse of duration 1 ns, the optimal amplitude and density are much lower at approximately $a_0 = 0.001$ and $n = 5 \times 10^{17}\,\mathrm{cm^{-3}}$. Considering Fig. 4.2b, using these low densities will require a low temperature of approximately 20 eV.

In fact, these considerations mean that the maximum feasible pump duration (for wavelength 1 μm) is approximately 1 ns. Longer pulses would require a longer length of plasma, so the plasma density must be lower to mitigate the competing instabilities. However, Fig. 4.2b showed the minimum feasible plasma density is $10^{17}\,\mathrm{cm^{-3}}$.

4.8 Experimental design

Taking all of these considerations into account, there is only one free parameter for the design of the amplifier [69]. Once the pump pulse duration is known, all other parameters are fixed by the considerations in this chapter. The efficiency and compression ratio can be simultaneously maximised using the optimal pump amplitude and plasma density from Fig. 4.15. Once the required pump pulse amplitude is known, this fixes the transverse width of the interaction.

The length of the plasma should cover the overlap distance between the pump and seed, which is $cT_0/2$ where T_0 is the pump pulse duration. The plasma temperature should be kept within the optimal region shown in Fig. 4.2. The seed pulse should be shorter than γ_{RFS}^{-1} to avoid breaking into multiple pulses and ideally meet the condition (2.44) from the start. This will require it to have an intensity greater than that of the pump pulse.

I now present two examples. In the first, the pump pulse has a wavelength of $1.054\,\mu$m, energy 50 J and duration of 25 ps, meaning the plasma length should be 4 mm. Referring to Fig. 4.15, the optimal electron density is around $0.015n_{crit}$ or $1.5\times10^{19}\,\mathrm{cm}^{-3}$ and the optimal pump intensity is around $3\times10^{14}\,\mathrm{Wcm}^{-2}$. Given this pump pulse, these parameters will maximise the efficiency and compression ratio. The electron temperature should be kept below 400 eV, placing a restriction on any external ionisation mechanism. For the seed to be effective and avoid breaking up in to multiple spikes, it should have duration less than $\gamma_{RBS}^{-1} \simeq 230$ fs. To reach the high efficiency non-linear stage from the start, the seed intensity needs to exceed the pump intensity, requiring around 100 mJ in the seed. If this is true, the overall efficiency could be 40%, giving an output energy of 20 J in 20 fs.

For the second example, I show how a greater compression ratio and greater output power may be achieved, albeit with a much larger plasma. I assume the pump pulse has wavelength $1.054\,\mu$m, energy 10 kJ and duration 1 ns, typical of Nd glass laser systems. This would be a more challenging proposition than the first example, since the plasma length must be 15 cm. Fig. 4.15 shows the optimum plasma density is $6\times10^{17}\,\mathrm{cm}^{-3}$ and the optimal pump intensity is $10^{13}\,\mathrm{Wcm}^{-2}$. Since the plasma is large and premature back-scatter of the pump must be avoided, and the temperature must be kept to less than 20 eV, the seed ionisation scheme will be required. The initial seed must be intense enough to ionise the plasma. The seed has wavelength 1081 nm and the efficiency should be 50%, giving a 5 kJ

output in 50 fs. If well focussed, this could give a focal intensity of $10^{24}\,\mathrm{Wcm}^{-2}$, high enough to create electron-positron cascades.

I will now use particle-in-cell simulations to verify these examples.

4.9 Two-dimensional simulation

There are several physical effects that may change the performance when several spatial dimensions are simulated. First of all, the transverse filamentation instability could degrade the pulse [41]. This problem was assessed in the previous section.

The intensities in Raman amplifiers are higher than the damage thresholds of solid state optics, but lower than reached at the target of many current CPA systems. Reaching unprecedented intensities therefore requires focussing of the output of the Raman amplifier. This poses a problem because the unfocused output of the amplifier already exceeds the damage threshold of solid state optics. One obvious solution is to amplify the seed pulse in a pre-focussed geometry, as shown in Fig. 4.16.

A high quality seed pulse is pre-focussed through an underdense plasma, where the pump beam overlaps it and transfers energy. If the coherence of the pulse is retained, it should achieve the same focal spot size as with no plasma interaction. This would allow a high power Raman amplified pulse to focus to an unprecedented intensity of greater than $10^{22}\,\mathrm{Wcm}^{-2}$ [70].

The two-dimensional simulation must therefore verify the high efficiency energy transfer, but also that a good focal quality is possible. The amplified pulse must be simulated along the subsequent propagation through vacuum, right to its focus.

The simulated pump pulse had amplitude $a_0 = 0.016$ and wavelength $1.054\,\mu$m. The plasma density was $1.5\times10^{19}\,\mathrm{cm}^{-3}$ or $0.015n_{crit}$, with fixed temperature 400 eV, placing the interaction near the optimum in Fig. 4.8*c*. The interaction length was 4 mm, approximately half of L_{max}. The transverse width of the pulses was 0.8 mm and the seed pulse focussed 15 mm after the plasma.

These parameters are approximately at the optimum for a 25 ps duration pump pulse, given in Fig. 4.15. The plasma temperature is also within the bounds shown in Fig. 4.2.

These parameters were chosen to reflect the first example given in the previous section and match an upcoming experimental campaign at the University of

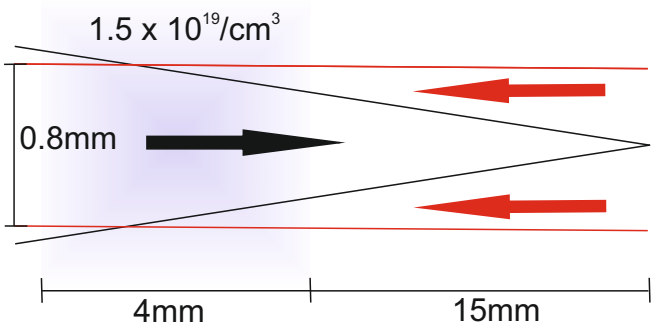


Figure 4.16: Schematic of the simulated Raman amplifier, where the seed pulse (black) is amplified in a focussing geometry by a long pump pulse (red). The plasma had a super-Gaussian spatial profile, with peak density $1.5 \times 10^{19}\,\mathrm{cm}^{-3}$ over a plateau of length 4 mm and ramps of length 1 mm. This gives $n/n_{crit} = 0.015$. The seed and pump were counter-propagating with wavelengths $1.2\,\mu$m and $1.054\,\mu$m respectively. The profiles were Gaussian in space and time, with duration 50 fs and 25 ps. Both had peak intensity $3 \times 10^{14}\,\mathrm{Wcm}^{-2}$, giving $a_0 = 0.016$, and transverse beam size 0.8 mm in the plasma.

Rochester, described in Chapter 8. The simulation used a moving window that followed the pulse.

The simulation box width was 1.3 mm and length was 70 μm, with 30 cells per wave- length, cell width 0.5 μm and 16 electrons per cell. The seed pulse duration of 50 fs is less than the linear Raman growth time of 200 fs, satisfying the conclusions of the seed pulse study. The initial peak intensity of the pump and seed in the plasma were $3 \times 10^{14}\,\mathrm{W/cm}^2$.

The seed pulse was amplified in the plasma and then followed to its position of best focus. Although ponderomotive filamentation was minimal, relativistic self focussing in the plasma moved the position of best focus from 15 mm to 12 mm after the end of the plasma. This suggests the focal position may be strongly affected by the interaction, even if the pulse retains good coherence. Despite the shift, the focal position is likely to be consistent from shot to shot and so it may be accounted for in experimental design.

The resulting intensity profile at the position of best focus is shown in Fig. 4.17a. The final pulse duration of 20 fs agrees with the one-dimensional simulation. The total energy transfer efficiency from pump to seed pulse was 30 %, slightly lower than the 1D parameter scan in Fig. 4.13c which gave 50 %. The 2D efficiency is

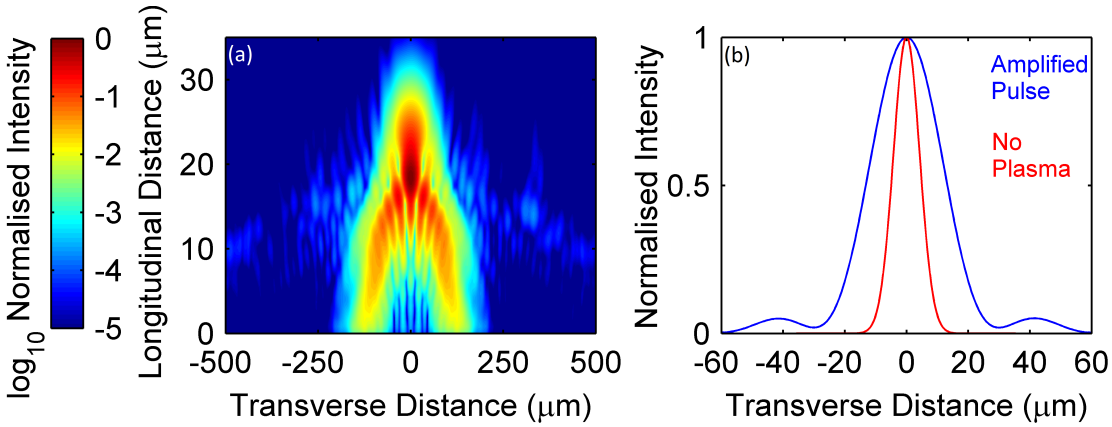


Figure 4.17: Results of the two-dimensional PIC simulation illustrated in the schematic in Fig. 4.16. (a) The amplified pulse intensity profile at the position of best focus. The intensity has been normalised to the peak value and shown on a logarithmic scale. The pulse propagates upwards. (b) A transverse cross section of the pulse at the position of best focus. This is compared to the equivalent profile for the seed pulse focussed with no plasma interaction. Since the amplified pulse is much higher intensity, the graphs have been normalised to compare.

slightly lower because the seed pulse started below the non-linear stage and there were a range of conditions across the width of the interaction.

There is some light surrounding the main high intensity region, with a contrast ratio of around $10-1000$. The contrast ratio at tens of pico-seconds before the pulse is also likely to be around $10-1000$, due to spontaneous Raman backscatter of the pump. This spontaneous scatter across the rest of the plasma was not simulated, since the simulation box only covered the extent of the seed pulse.

Since the pump pulse amplitude in the plasma varies transversely, the pulse contracts to a shorter duration on-axis than the region at the edges. In the focal plane, this means a lot of energy in the wings of the amplified pulse arrives behind the main peak, a less than desirable result. The peak intensity could therefore be improved by using a pump pulse with a flat-top transverse profile, as opposed to the Gaussian profile in the simulation. This would keep the Raman growth rate uniform across the full width of the interaction.

Fig. 4.17b shows a transverse line-out of the amplified pulse, at the position of best focus. It compares this with the equivalent profile at best focus of the seed pulse with no plasma interaction. The focal spot has become slightly widened by the plasma interaction, however is still within a factor of three of the diffraction limited width.

4.10 Compression of longer pulses

It is also important to verify that the scheme works for long pulses such as those available on the National Ignition Facility. There are already some promising experimental results [56]. These pulses have a duration exceeding 1 ns, with an energy up to 10 kJ. If these pulses could be used to drive a larger secondary stage amplifier, the resulting output power could greatly exceed 1 PW.

Considering the findings earlier in this chapter, the efficiency and compression ratio may be optimised by using a lower intensity pump pulse in a lower density plasma [122]. For a 1 ns pump pulse, the length of the plasma should be $cT_0/2 = 15$ cm. To alleviate competing plasma instabilities over such a large plasma volume, the amplifier should use the seed ionisation scheme.

I investigated compression of longer pulses with a further one-dimensional particle-in-cell simulation using the seed ionisation scheme. The set-up was similar to those of the parameter scan in Fig. 4.8. The long counter-propagating pump pulse had duration 1 ns and wavelength 1.054 μm. The plasma density was $6 \times 10^{17}\,\mathrm{cm}^{-3}$, giving $n = 0.0006 n_{crit}$. The pump pulse had linear polarisation and its intensity was $10^{13}\,\mathrm{Wcm}^{-2}$, giving a dimensionless amplitude of $a_0 = 0.003$. These parameters are close to the optimal values given in Fig. 4.15 for a 1 ns duration pump pulse.

The seed pulse had a wavelength of 1081 nm, initial duration 50 fs and peak intensity $3 \times 10^{14}\,\mathrm{Wcm}^{-2}$, placing it close to the non-linear stage. The seed pulse was intense enough to ionise the plasma from the start, however the pump pulse was not. The Hydrogen ionisation front was modelled using a field ionisation algorithm that initialises free electron particles at each time-step, based on the current electric field strength. The electrons are initialised with a velocity along the direction of the polarisation of the laser pulse.

Fig. 4.18 shows the simulated intensity across the space-time diagram. The spatial coordinate is taken in a frame moving at the group velocity of the amplified pulse. The seed pulse is intense enough to ionise the plasma. It initially lengthens in the linear stage, before reaching the non-linear stage after a propagation of 2 cm. The pulse envelope remains smooth and unaffected by Raman forward scattering across the full 15 cm interaction.

Some of the light propagates in the neutral gas ahead of the ionisation front, with a faster group velocity. This does not ionise the plasma or affect the main

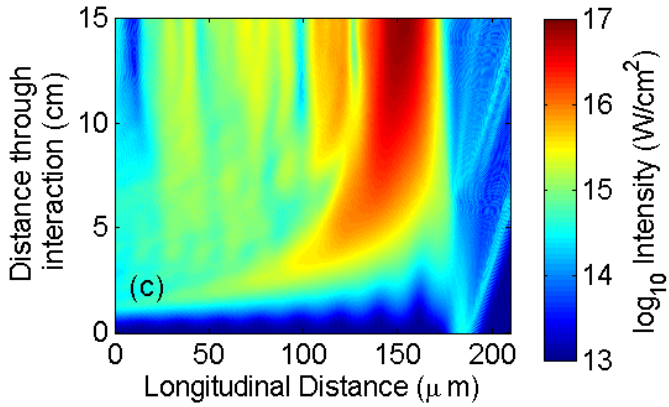


Figure 4.18: One-dimensional PIC simulation, showing effective compression of a nanosecond duration, 10 kJ, 1.054 μm pump pulse using the seed ionisation scheme. The plot shows the intensity across the space time diagram, in spatial coordinates moving at the group velocity of the seed pulse. The plasma had length 15 cm, density 6×10^{17} cm^{-3} and was ionised by the seed pulse, using a field ionisation model. The seed started with energy 15 J and duration 50 fs, with a central wavelength 1081 nm.

interaction, however it will create a pre-pulse and lower the contrast ratio of the final amplified pulse. The characteristic secondary peaks of the amplified pulse grow behind the main spike in the later stages.

The simulated total energy transfer efficiency was 50%. This is a very high efficiency and comparable to many conventional pulse compressors. Furthermore, the seed ionisation scheme would prevent premature Raman scattering of the pump pulse, since it traverses across a neutral gas until it meets the seed.

The seed ionisation scheme also mitigates detrimental kinetic effects. With these parameters, the electron-electron collision time is longer than the duration of the amplified pulse, so the newly ionised electrons move mainly transversely, along the polarisation of the laser. The effective longitudinal temperature is low. This reduces the amount of electron trapping within the plasma wave and gives a greater efficiency.

The seed ionisation scheme provides two major benefits as the premature pump pulse scattering is eliminated and the temperature is kept low, improving the efficiency. Collisional damping of the pump is also avoided, as well as thermal expansion of the plasma, potentially de-tuning the resonance [120].

Extrapolated transversely across a $1 \times 1 \times 15$ cm plasma volume, the simulated pump pulse had energy 10 kJ and power 10 TW. The seed pulse had energy 15 J and power 0.3 PW. A seed pulse such as this could be developed from a smaller

scale Raman pre-amplifier, as demonstrated in the previous section. The wide bandwidth developed by the amplified pulse will give some freedom to use a lower plasma density for the secondary amplifier and still maintain resonance.

The energy transfer efficiency of 50% means the final seed pulse energy was 5 kJ, within a duration of 60 fs. The amplified pulse power is close to 100 PW, a value far in excess of current capabilities.

A two-dimensional simulation is not currently feasible, however transverse ponderomotive filamentation is unlikely to be an issue. The integrated growth rate of the filamentation instability throughout the simulation is $\Lambda_{FIL} \simeq 3$.

The long pump pulse may also experience detrimental instabilities as it traverses the neutral gas. This is due to the non-linear intensity dependent refractive index of the gas. The on-axis phase difference due to this effect can be estimated through the B integral for the pump pulse. This is given by $k_0 \int_0^L n(I) - n(0) dx$, where $n(I)$ is the refractive index including its intensity dependence. For a 15 cm thick Hydrogen target, I estimated this as 0.02, a value well within typical tolerances.

The results suggest that a frontier level pulse power could be reached via staging the amplifiers. Firstly, a millimetre scale amplifier at higher density would generate a Joule level, fully coherent seed pulse. Next, this pulse could seed a much larger 15 cm long amplifier at a lower density. This larger amplifier must use the seed ionisation scheme to enable the use of much lower plasma density. Both of these stages may use pump pulses similar to those at currently operational facilities.

The simulations and analytical predictions in this chapter show that a Raman amplifier can develop a pulse close to relativistic intensity. Particle-in-cell simulations have found the efficiency of the amplifier as a function of the three main parameters. The pulse amplitude and duration were found as a function of time, as well as the maximal power amplification before competing instabilities take hold. These considerations mean that for a given pump pulse duration, there is a unique intensity and plasma density that maximise the efficiency and compression ratio. The optimal seed is shorter than the linear Raman growth time. The scheme is effective in a two-dimensional focussing geometry. Providing the plasma is ionised by the seed, it may be extended to 15 cm long and use nanosecond duration pump pulses.

In the next chapter, I calculate how the beam geometry may be manipulated to combine both of these stages into a single, uniform density amplifier. My new

geometry drastically reduces the required seed pulse energy to millijoule level and the secondary peaks of the amplified pulse are eliminated.

The results in this chapter have been published in reference [107]. All results are my original research, apart from two sections. Some of the analytical results on the parameter space have been previously published in reference [9]. The section on the optimal seed pulse was performed in collaboration with Raoul M. G. M. Trines, who also performed similar simulations [123].

Chapter 5

Solution for Diverging Beams

Short, intense seed pulses have the greatest performance. However, it is difficult to generate powerful pulses at the shifted wavelength that is required to maintain resonance. The condition (2.44) for the high efficiency behaviour is independent of the transverse dimension, so starting the amplifier at the focus of the seed pulse could be an easy way to reach the high intensity required to access the non-linear stage [1].

In three dimensions, a coherent Gaussian laser pulse can be amplified as it converges outside of its Rayleigh range, within its Rayleigh range, or as it diverges outside of the Rayleigh range. The Rayleigh range is the region within a distance L_R of the focus. In the original proposition for Raman amplifiers [73], the beam was amplified as it converged to focus. However, in their mathematical derivation, repeated in Chapter 2, it was assumed to be collimated. In this chapter, I generalise those coupled envelope equations for the Raman amplifier to include a three dimensional geometry. I have submitted this work for publication [109].

Here, for the first time, I propose the exact opposite geometry to reference [73]; the seed should be amplified as it diverges. The rationale is that as it expands, this will moderate the intensity while the power and energy increase, as suggested in reference [23]. Furthermore, near its focus, only a low power seed pulse is required to access the non-linear stage. This drastically reduces requirements on the seed. The new geometry replicates that of typical high-power laser systems, which start with a low power seed and a narrow diameter, then the amplifiers become wider as the power increases [66].

A schematic is shown in Fig. 5.1. The seed pulse is amplified as it diverges away from its focus. The overlap with the pump beam can be maximised by using the same F number and focal position for both beams. Since the transverse spot size

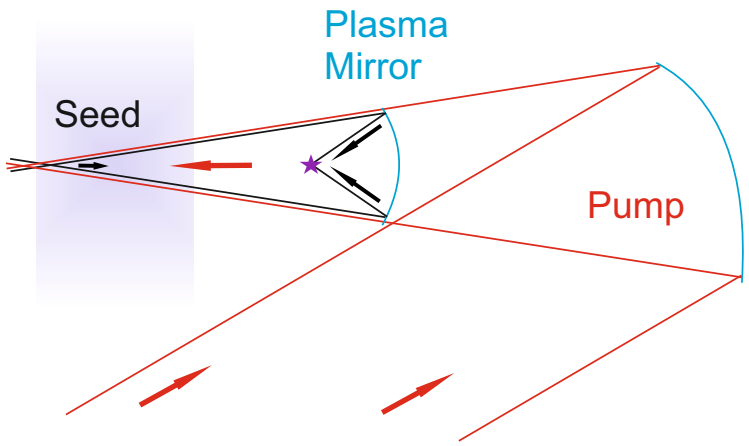


Figure 5.1: Schematic of the new proposed Raman amplifier beam geometry. Amplifying an expanding beam will keep its intensity low as the power increases, avoiding unwanted instabilities in the plasma. The seed (black) propagates to the right and overlaps with the counter-propagating pump (red) inside the plasma (purple). The amplified pulse can then be focussed by a spherical plasma mirror to achieve unprecedented intensity at the focus, marked by the star. The long pump pulse will not ionise the plasma mirror and may propagate through the back. Alternatively, the pump could arrive from a different angle and propagate past the side of the plasma mirror.

is small near the focus, the seed pulse will have high intensity and easily meet the condition (2.44). This contrasts with the case of collimated beams, which would require a much wider beam waist and correspondingly higher seed power.

The seed diverges as it is amplified, maximising the volume of the interaction and therefore the energy transfer. Once the short pulse reaches the end of the amplifier, the amplified pulse is still diverging and it must be focussed via external methods. This is the main disadvantage of the scheme. This poses a problem, since a large system may produce a pulse that greatly exceeds the damage threshold of standard optics. A promising technique is to use a curved plasma mirror to focus the amplified pulse [24, 20].

The pump pulse is approximately 1/1000 the intensity of the amplified pulse and so it will not ionise the mirror and may propagate through the back of it. This method will allow the two pulses to counter-propagate, overlap and have the same focal position, an important aspect of the scheme.

Although reference [73] suggests use of a non-collimated beam, the derivation presented there and in Chapter 2 applies strictly to collimated beams. I now gener-

alise the coupled envelope equations in Chapter 2 to include converging or diverging beam geometries and recover the results for collimated beams as a special case.

5.1 Generalisation of the coupled envelope equations

In a collimated amplifier, the pulse amplitude rises proportional to the propagation time [121]. However, if a beam is expanding away from a focus, this will cause its beam waist to increase and its intensity to fall. For a Gaussian beam, the intensity on the propagation axis is $I_1 = I_{max}(w_0/w(z))^2$ where I_{max} is the maximum intensity at focus and the beam waist $w(z) = w_0\sqrt{1 + z^2\lambda^2/\pi^2 w_0^4}$. For positions far outside the Rayleigh range, this can be approximated as $w(z) = z\lambda/(\pi w_0)$. This leads to $I_1 = I_{max}(\pi w_0^2/z\lambda)^2 \propto 1/z^2$.

The z position of the peak intensity of the pulse is then $z \propto t - t_f$ where t_f is the time of best focus of the pulse. The maximum on axis intensity therefore decreases as $I_1 \propto 1/(t-t_f)^2$ and the amplitude decreases as $a_1 \propto 1/z \propto 1/(t-t_f)$. Assuming a Gaussian laser pulse is outside its Rayleigh range, this means the dimensionless amplitude on the propagation axis evolves as

$$a_1 \propto \frac{1}{t - t_f}, \tag{5.1}$$

$$\frac{da_1}{dt} = -\frac{a_1}{t - t_f}, \tag{5.2}$$

where t_f is the time of maximum intensity of the pulse. At times $t < t_f$ the short seed pulse has not yet reached the position of best focus. Later times $t > t_f$ correspond to a diverging beam. As before, the propagation time t is measured from the start of the non-linear stage and is always positive. Addition of this new geometrical term changes the coupled envelope equations (2.36) to

$$(\partial_t + 2c\partial_x)\, a_0 = \frac{\omega_p^2}{4\omega_0} n_2 a_1, \tag{5.3}$$

$$\partial_t a_1 = -\frac{\omega_p^2}{4\omega_0} n_2 a_0 - \frac{a_1}{t - t_f}, \tag{5.4}$$

$$(\partial_t + c\partial_x)\, n_2 = -\frac{\omega_0^2}{\omega_p} a_0 a_1. \tag{5.5}$$

The presence of this new term changes the self-similar solution for the non-linear stage. The geometrical term for the plasma wave has been neglected, since its phase velocity is much less than c. In addition, the seed duration is short and so

the unperturbed pump amplitude will hardly change within the seed pulse volume. This means the pump geometrical term is also neglected, although the incoming pump amplitude $a_0(t,0)$ is now treated as a variable. The effect of a wave-guide on Raman amplification has previously been explored in reference [18].

To find the new solution, one can proceed in a similar manner to the solution for a collimated amplifier in Chapter 2 [73], by treating the time derivatives in equations (5.3) and (5.5) as negligible. This will find the evolution in the limit that the pulse duration is much less than its timescale of evolution.

For a function $u(t,x)$, I again take the trial solutions $a_0 = \alpha_0(t)\cos(u/2)$, $a_1 = \alpha_1 c\partial_x u$ and $n_2 = \alpha_2 \sin(u/2)$. The variable $\alpha_0(t) = a_0(u=0)$ is the pump amplitude at the start of the short seed pulse. In general, since the pump power may change and it is focussing, this may be a function of time. The pump amplitude $\alpha_0(t)$ is assumed to vary much more slowly than the duration of the seed. As before, equations (5.3) and (5.5) are consistent if

$$\alpha_1 = -\sqrt{\frac{2}{\omega_0\omega_p}}, \tag{5.6}$$

$$\alpha_2 = \alpha_0\left(\frac{2\omega_0}{\omega_p}\right)^{\frac{3}{2}}. \tag{5.7}$$

First of all, I find the solution when $a_0(t,0)$ is constant. In this case, equation (5.4) becomes

$$c\alpha_1\partial_t\partial_x u = -\frac{\omega_p^2}{4\omega_0}\alpha_0^2\left(\frac{2\omega_0}{\omega_p}\right)^{\frac{3}{2}}\sin(u/2)\cos(u/2) - \frac{\alpha_1 c}{t-t_f}\partial_x u, \tag{5.8}$$

$$\left(\partial_t + \frac{1}{t-t_f}\right)\partial_x u = \frac{\omega_0\omega_p}{4c}\alpha_0^2\sin(u). \tag{5.9}$$

In the limit that $|t_f| \gg t$ the focus is very far away compared to the scale of the interaction and one should recover the case of a collimated beam discussed in Chapter 2. Taking the same approach as before, I look for self similar solutions of a function $f(\beta) = u(t,x)$ in the dimensionless co-ordinate $\beta = \omega_0\omega_p\alpha_0^2 xt/4c$. Changing variables to β, and assuming the incoming pump amplitude α_0 is constant, the equation becomes

$$\beta\frac{d^2f}{d\beta^2} + \frac{df}{d\beta} + \frac{t}{t-t_f}\frac{df}{d\beta} = \sin(f), \tag{5.10}$$

$$\beta\frac{d^2f}{d\beta^2} + \left(\frac{2t-t_f}{t-t_f}\right)\frac{df}{d\beta} = \sin(f). \tag{5.11}$$

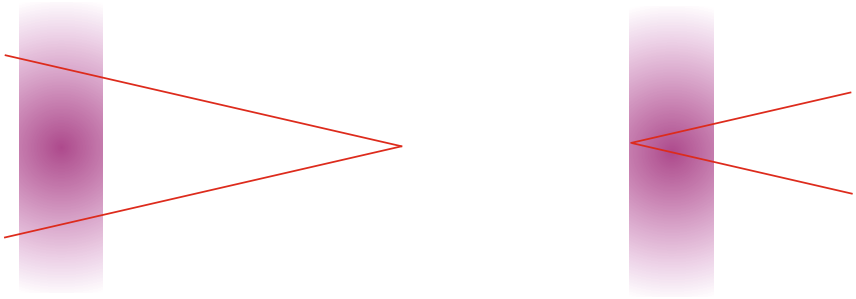


Figure 5.2: The two limits of the beam geometry that exhibit self-similar amplification of the pulse. The extent of the amplified beam is shown in red and the plasma is shown in pink. The seed propagates towards the right. (a) The extent of the non-linear stage in the plasma is far away from the focal position. This is the previous solution for a collimated amplifier. (b) The pulse enters the non-linear stage at its focus and expands as it amplifies. This is the new self similar solution and it has a different amplified pulse shape.

The variable t still appears in this equation and so the evolution in three dimensions is, in general, not self similar. However, in the limit that $|t_f| \gg t$, the beam is collimated, recovering the original equation (2.49) found by Malkin *et al.* [73].

There is also self-similar behaviour in the limit that $|t_f| \ll t$. This is the situation when the pulse enters the non-linear stage at or very near its focus, such that the focal position falls within the extent of the amplifier. Note that this limit will only apply to an expanding amplified beam rather than a converging beam, since it implies the non-linear stage must start at the focus, so that $t_f \to 0$.

The development of the seed pulse is self similar as it expands outside of its Rayleigh range. This is the situation illustrated in Fig. 5.1. The beam geometries for the two limiting cases with the self similar behaviour are also drawn in Fig. 5.2. Any intermediate case will still have amplification, although it won't be self-similar (not solely a function of β).

In this new limit for a diverging beam, equation (5.11) becomes

$$\beta \frac{d^2 f}{d\beta^2} + 2\frac{df}{d\beta} = \sin(f). \tag{5.12}$$

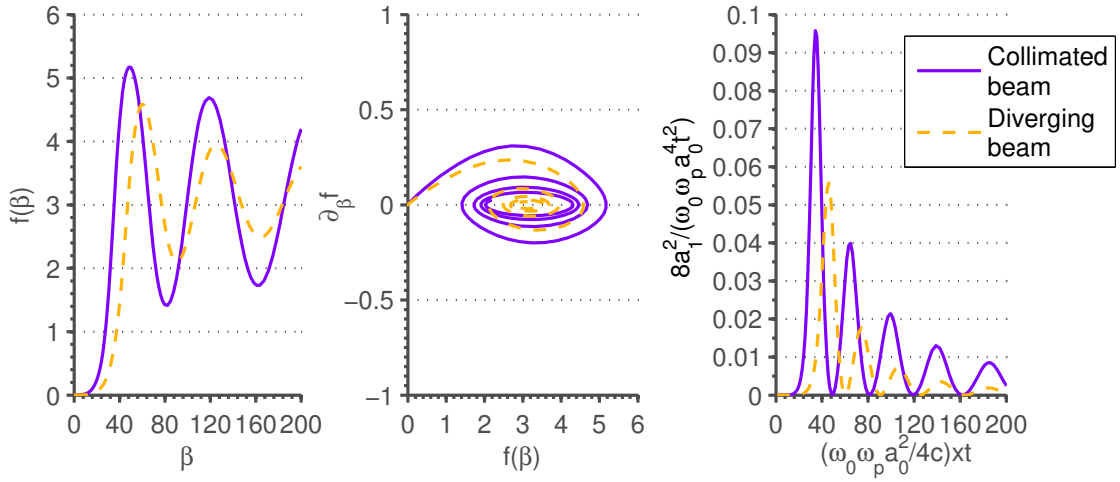


Figure 5.3: Numerical integration of the self-similar equations for the limit of a collimated beam (purple) or a diverging beam (orange). (a) f as a function of the dimensionless self-similar coordinate β. It tends towards π for both cases. (b) The development in the f, f' plane. (c) The shape of the amplified pulse power envelope, as found from these solutions.

This equation is comparable to that for a collimated beam, except for the multiplier of 2 for the first order derivative. The evolution of a diverging beam therefore has the same characteristics and scalings to the standard collimated beam, however the pulse shapes are slightly different. This is illustrated in Fig. 5.3, which numerically integrates equation (5.12) and compares to the collimated case discussed in Chapter 2. This solution assumes the pump pulse amplitude α_0 is constant at all times.

Both self-similar equations have the form of a damped anharmonic oscillator. However, for the expanding beam the damping term is greater and so the secondary trailing peaks are slightly suppressed compared to the standard collimated case. In addition, the peak amplitude is lower than the collimated case, but the pulse duration is slightly longer.

The relation (2.44) still holds for the expanding beam. This can be seen from the trial solution $a_1 = \alpha_1 c\partial_x u$, which shows that the integrated amplitude of the pulse is just $u(\beta \to \infty) = \pi$. This relation is the same as for the collimated case. Furthermore, the solution is still attractive since f tends to π for all non zero initial conditions.

It is also possible to find the seed pulse amplitude as a function of time, assuming the non-linear stage starts at the focus. From equation (4.19) in Chapter 4, the

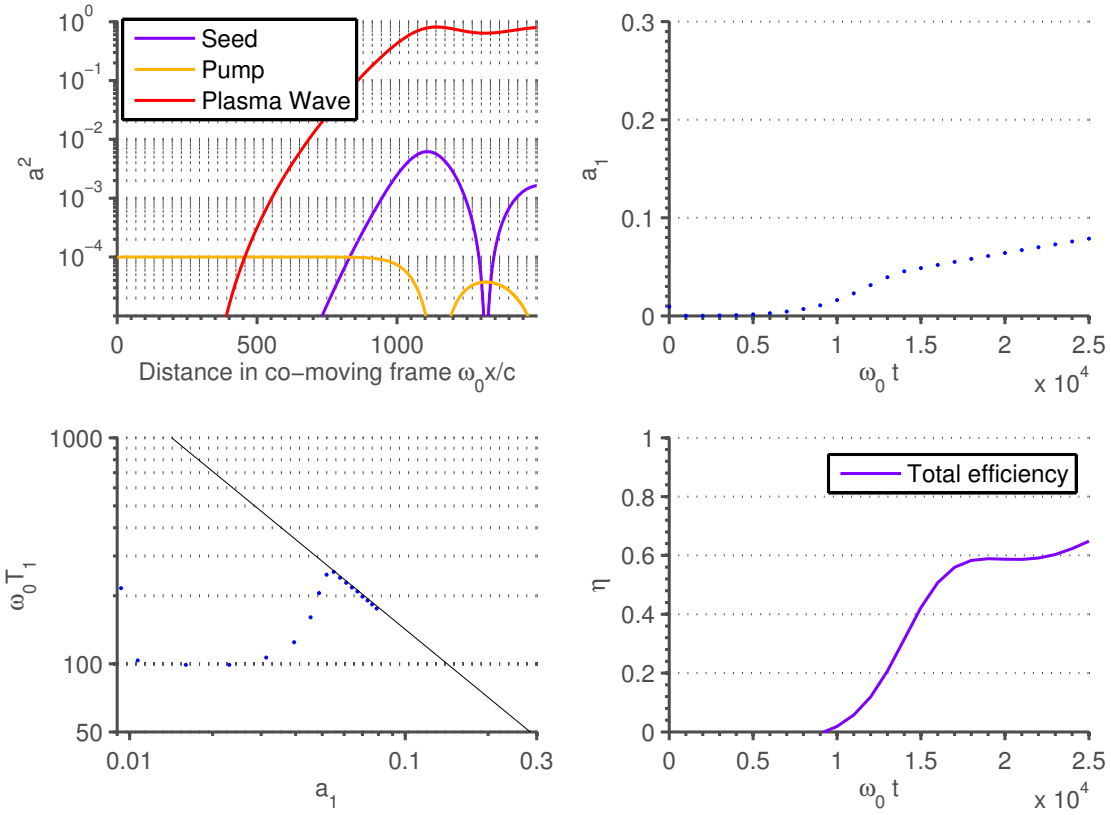


Figure 5.4: Numerical integration of the coupled equations (5.3-5.5) for a Raman amplifier with an expanding beam. The pump pulse amplitude was constant at $a_0 = 0.01$ and the density was $n = 0.01 n_{crit}$. The seed started at a distance $10c/\omega_0$ past the focus, with an amplitude $a_1 = 0.01$ and a duration $\omega_0 T_1 = 210$. Compare to Fig. 4.4 which found the evolution of the pulse for the same conditions in a collimated amplifier. (a) The seed and pump power densities and the plasma wave amplitude at a time $\omega_0 t = 25000$. (b) The amplitude of the amplified pulse as a function of time. (c) The development of the amplitude versus duration of the amplified pulse. (d) The total energy transfer efficiency from pump to seed.

amplitude for a collimated seed beam was

$$a_1(t) = \frac{2}{5}\sqrt{\omega_0 \omega_p}\frac{\omega_0^2}{\omega_1^2}\eta_1 a_0^2 t. \tag{5.13}$$

In the case of an expanding beam, the total pump energy is that contained by a cone of length ct and radius w. This means the energy balance equation becomes

$$2\eta_1 a_0^2 \omega_0^2 (\frac{1}{3} t \pi w^2) = a_1^2 \omega_1^2 T_1 \pi w^2. \tag{5.14}$$

Using the non-linear condition (2.44) gives

$$a_1(t) = \frac{2}{15}\sqrt{\omega_0 \omega_p}\frac{\omega_0^2}{\omega_1^2}\eta_1 a_0^2 t. \tag{5.15}$$

The pulse amplitude is one third of the value for a collimated amplifier after the same length, and it rises linearly.

Fig. 5.4 numerically integrates the coupled envelope equations for the same conditions as the collimated amplifier in Fig. 4.4. In this new case, the pulse is amplified as it expands, with the integration starting close to the focus. The time of focus was $t_f = -10/\omega_0$, so the beam is now expanding. Comparing the pulse after $\omega_0 t = 25000$ to Fig. 4.4, the amplitude is around 1/3 of the previous value. This agrees with the derivation above.

Fig. 5.4b shows that the amplitude initially drops as the pulse expands, however after some time it recovers due to Raman scattering, reaching the non-linear stage. Panel (c) shows that the main non-linear relation (2.44) shown in black, still holds for the later stages.

The result in panel (d) shows that very high efficiency can also be reached with the expanding geometry, although there appears to be a longer start up time in the linear stage. Interestingly, the results are independent of the F number of the beam, so long as the limit $t \gg |t_f|$ applies and the incoming pump intensity is constant for all times.

Importantly, these results show that the high efficiency behaviour is still achievable with non-collimated beams, as originally envisaged in reference [73]. In the case of a diverging beam where the non-linear stage starts at the focus, the evolution of the seed pulse is also self-similar, but with a different pulse shape to the collimated case.

The seed amplitude grows linearly with time $a_1 \propto t$, its intensity $I \propto t^2$, its duration $\propto 1/t$, its power $\propto t^4$ and its energy $\propto t^3$. This is because the energy is proportional to the volume of the cone that the pulse has propagated through, which increases $\propto t^3$.

5.2 Steady state amplification

The previous section assumed that the pump pulse intensity is constant throughout the amplifier. For a collimated pump beam, this means it has constant power. However, for a converging pump beam as shown in Fig. 5.1, there is greater beam overlap if the pump beam is focussed to the same position and F number as the seed. In this case, to maintain the pump intensity constant, the pump power must increase proportional to $(t - t_f)^2$.

I now show that if instead, the pump pulse power increases linearly with $t - t_f$, the seed pulse amplitude will remain constant as it amplifies. Furthermore, the characteristic secondary trailing peaks will be completely eradicated. This is a major benefit, since all of the energy transferred will reside in the main spike of the amplified pulse.

Returning to equation (5.9), before the assumption of constant α_0 was made, there is a static solution when $\partial_t = 0$. In this instance, the equation becomes

$$\partial_x u = \frac{\omega_0 \omega_p}{4c}(t - t_f)\alpha_0^2 \sin(u). \tag{5.16}$$

This means that if the incoming pump pulse amplitude decreases as $\alpha_0 \propto 1/\sqrt{t - t_f}$, the time coordinate will be eliminated and u is constant in time. In this case, u is given by

$$\frac{du}{dx} = \frac{\omega_0 \kappa}{c} \sin(u), \tag{5.17}$$

$$\kappa = \frac{\omega_p (t - t_f)}{4} \alpha_0^2. \tag{5.18}$$

Where I have defined the constant κ, which is related to the rate of increase in the pump pulse power. Since $\alpha_0 \propto 1/\sqrt{t}$, the pump pulse power is

$$P \propto \omega_0^2 \alpha_0^2 w^2 \propto \omega_0^2 \frac{1}{t} t^2 \propto t. \tag{5.19}$$

This means in the diverging seed geometry, if the pump pulse power increases linearly, the seed pulse will have constant intensity and duration as its energy and power increase.

The solution to equation (5.17) is

$$\tan(u/2) = e^{\kappa \omega_0 (x - x_0)/c}, \tag{5.20}$$

where x_0 is an integration constant. The solution $u = 0$ is an unstable equilibrium, so that any initial seed will always tend to the stable equilibrium $u = \pi$. The solution is independent of time, so that the seed pulse has constant amplitude as it amplifies. The shape of the seed pulse envelope can be found by differentiating (5.20) to get

$$\frac{1}{2} \sec^2(u/2) \frac{du}{dx} = \frac{\kappa \omega_0}{c} e^{\kappa \omega_0 (x - x_0)/c}, \tag{5.21}$$

$$a_1(x) = \sqrt{\frac{2\omega_0}{\omega_p}} \frac{\kappa}{\cosh(\kappa \omega_0 (x - x_0)/c)}. \tag{5.22}$$

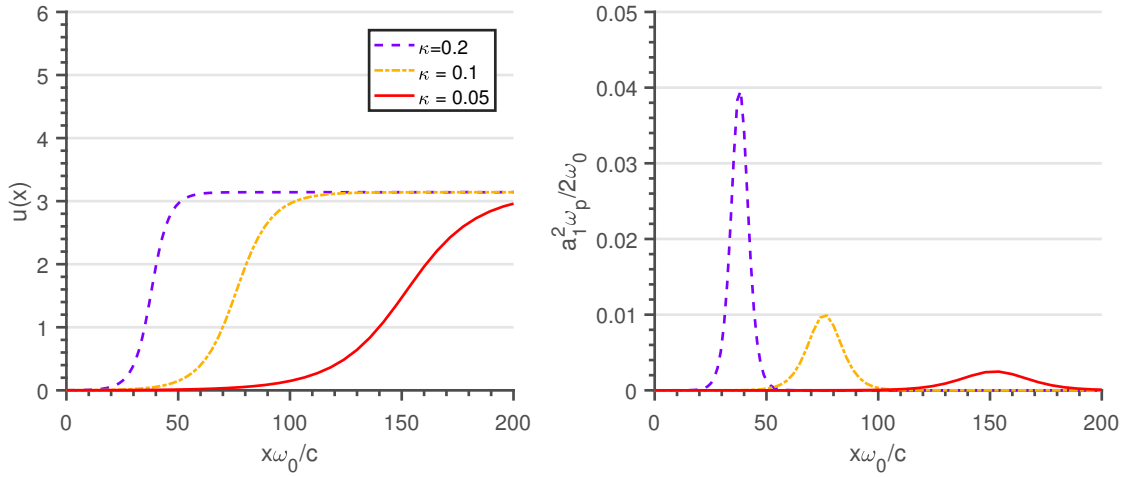


Figure 5.5: Solution to the equation (5.17), giving the amplified pulse power envelope for amplification of a diverging beam with $a_0 \propto t^{-1/2}$. (a) the numerical solution of the equation (5.17) for different values of κ. The boundary condition is $u(0) = 0.001$, corresponding to a small initial seed. (b) The power envelope of the amplified pulse for these same values of κ. The peak intensity is proportional to κ^2.

The pulse consists of a single amplified spike that monotonically depletes the pump pulse. This is illustrated in Fig. 5.5 which shows the solution for numerous values of κ. A higher gradient in pump pulse power means κ is greater and the stationary solution is a shorter, more intense spike. The peak amplitude is $a_1 = \kappa\sqrt{2\omega_0/\omega_p}$ and the duration of the power envelope is $T_1 = 1.76/(\omega_0\kappa)$. This means the amplitude duration product becomes

$$a_1 T_1 \sqrt{\omega_0 \omega_p} = 2.4. \tag{5.23}$$

This means that compared to the collimated case, the output of the amplifier will require a lower energy for a given output power. The constant κ and therefore the duration and amplitude of the output pulse are entirely set by the gradient of the pump pulse power and the F number of the interaction. In terms of these quantities, the parameter κ is

$$\kappa = \frac{2e^2\lambda_0}{\pi^2 m^2 c^6 \epsilon_0} F^2 \sqrt{\frac{n}{n_{crit}}} \frac{dP_0}{dt}. \tag{5.24}$$

For example, a $1\,\mu$m pump pulse increasing in power from 0 to 10 TW over 1 ns with plasma density $1.1 \times 10^{18}\,\mathrm{cm}^{-3}$ and length 15 cm, with a f/5 geometry, gives $\kappa = 0.008$. This means the amplified pulse has constant duration $T_1 = 120$ fs and constant amplitude $a_1 = 0.06$.

5.3 Simulation of a diverging beam

It is currently infeasible to fully investigate the process with 3D simulations, however the diverging beam scheme can also be investigated with 2D Cartesian particle-in-cell simulations.

In a 2D simulation, if the fields and plasma are periodic in the third dimension with Cartesian coordinates, the laser pulse effectively comes to a line focus. As a pulse diverges away from a line focus, its intensity drops $I \propto 1/(t - t_f)$ and the amplitude drops as

$$a_1 \propto \frac{1}{\sqrt{t - t_f}}, \tag{5.25}$$

$$\partial_t a_1 = -\frac{1}{2}\frac{a_1}{t - t_f}. \tag{5.26}$$

The derivation then proceeds as before, except with $\kappa \to 2\kappa$.

This means for the steady state amplification, the pump pulse amplitude must decrease proportional to $1/\sqrt{t - t_f}$. This is in agreement with the three-dimensional case in the previous section. To recover this behaviour in 2D, the pump power must be constant, rather than increasing linearly as in the 3D focussing case.

The situation with constant pump power may be easily simulated in 2D, for a seed pulse starting at its focus. I choose conditions close to the optimal in Chapter 4, with uniform plasma of density $n = 0.01 n_{crit}$ and $T = 100\,\mathrm{eV}$. The simulation box had a width of 1.3 mm and length 68 μm, moving along with the seed pulse at its group velocity. There were 60 cells per wavelength, each with width 0.5 μm and 16 electrons per cell. The ions were held static to reduce computational costs.

The pump pulse had constant power, such that the amplitude decreased proportional to $1/\sqrt{t - t_f}$ as the interaction moved away from the focus. Its wavelength was 1.054 nm and the F number was 5, giving $\kappa = 0.05$.

The seed was initialised at its focus as a transform limited Gaussian pulse, with duration 50 fs and wavelength 1.18 μm. The peak seed intensity was $3 \times 10^{16}\,\mathrm{Wcm^{-2}}$. This meant the seed pulse satisfied the non-linear condition (2.44) from the start, despite having energy of just 0.3 mJ. The beam waist at focus was 3.8 μm. The focal position was the same as for the pump.

The seed pulse amplifies as its FWHM expands from 5 μm to 570 μm. Its evolution across three time-steps is shown in Fig. 5.6. In the first time-step, shown in panel (a), the seed propagates towards the right and the long counter-propagating

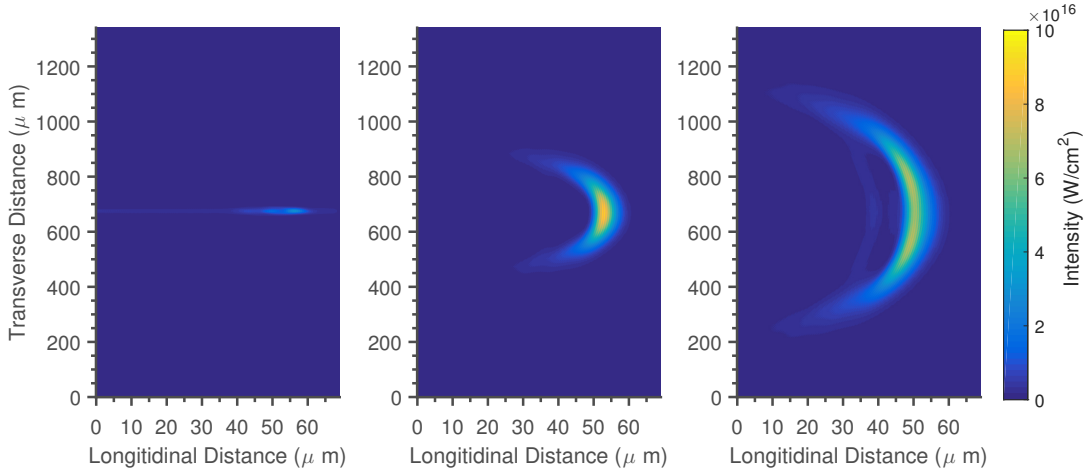


Figure 5.6: Particle-in-cell simulation of a Raman amplifier with a diverging beam. The pump pulse amplitude drops proportional to $1/\sqrt{t-t_f}$, such that the seed pulse intensity is constant as it amplifies. The plasma density was $10^{19}\,\mathrm{cm}^{-3}$ and the pump and seed were both focussed at $t = 0$ with a f/5 geometry. The transverse profiles were Gaussian. The figure shows the simulated intensity at three time-steps of (a) 0 ps, (b) 8 ps and (c) 16 ps. The colour scale is identical for all plots, showing that the peak intensity remains approximately constant as the power and energy increase.

pump pulse is also visible. The seed pulse amplifies in the later time-steps, shown in panels (b) and (c). The simulated intensity has been spatially smoothed on a wavelength scale. The first obvious development is the bowed shape of the amplified pulse, due to its divergence. There are no signs of any filamentation and the pulse envelope remains smooth. The pump is no longer visible on the colour scale because its intensity drops as $I_0 \propto 1/(t-t_f)$. A small secondary peak forms behind the main peak at the front.

This secondary peak can be seen more clearly in Fig. 5.7a, which shows a line-out on axis from the last time-step, after 16 ps. The secondary peaks are reduced in amplitude compared to the usual collimated case presented in Chapter 4. Compare, for example, to the simulation results of Fig. 4.5, in which the secondary peak is more than half the intensity of the first peak. With the new geometry, the secondary peak is 1/7 of the main peak intensity. This difference is despite use of the same plasma density and the same propagation distance. Although the secondary peaks are not completely eradicated in the PIC simulation, they are heavily suppressed, in agreement with the analytical results presented in the previous section.

In the final time-step, the amplitude is $a_1 = 0.24$, in agreement with the analytical prediction of $\kappa\sqrt{2\omega_0/\omega_p} \simeq 0.22$. The FWHM pulse duration is 17 fs, similar

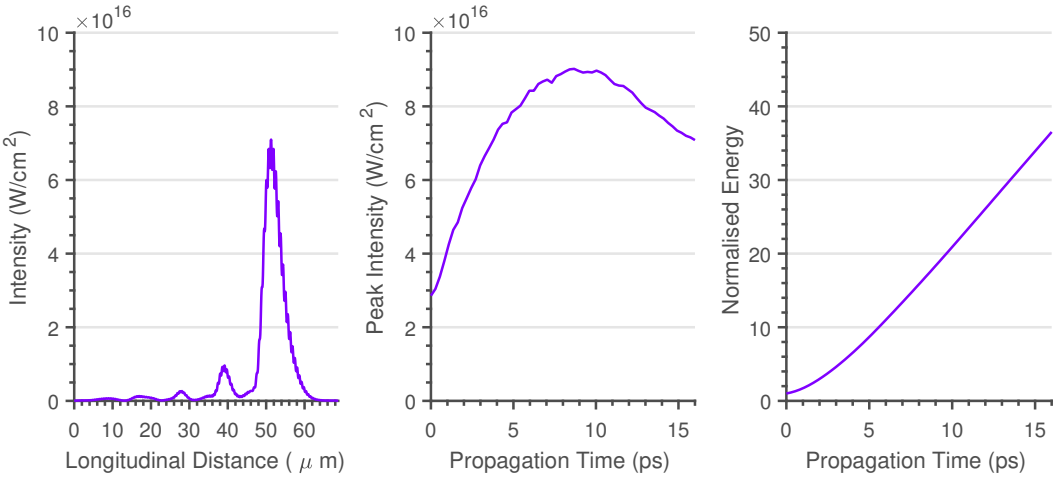


Figure 5.7: Results of the simulation of a diverging seed pulse, also shown in Fig. 5.6. (a) Line-out of the intensity on the axis of propagation, after propagation for 16 ps. This has been smoothed on a wavelength scale. The pulse propagates to the right. (b) The development of the peak intensity of the amplified pulse across the 16 ps propagation. (c) The development of the total energy of the amplified pulse, normalised to its initial value.

to the analytical prediction of 19.6 fs. This gives the duration amplitude product as $a_1 T_1 \sqrt{\omega_0 \omega_p} \simeq 2.3$.

This product is half that of the collimated beam scheme, proving that the simulated amplification accessed the new non-linear solution. In addition, the less useful secondary peaks are suppressed and all of the energy is concentrated in the main spike. This means that for an output pulse of identical power, a much lower total energy is required in the diverging beam scheme than the collimated scheme, possibly by more than a factor of two.

Fig. 5.7b also shows the evolution of the peak intensity of the amplified pulse. After an initial start-up stage where it reaches the stationary solution, the intensity remains constant or even drops slightly at later times. This is in contrast to the standard collimated scheme where the intensity continues to rise. The drop in amplitude at late times could be due to decreasing efficiency as the pump pulse intensity decreases. The peak intensity remains approximately constant throughout the interaction, in agreement with the analytical theory.

Fig. 5.7c shows the evolution of the pulse energy over the same timescale. Since the simulation was 2D, this was obtained by summing over the field energy density across two dimensions, giving a quantity with units J/cm. Although the peak intensity is approximately constant over time, the energy grows by a large factor

of nearly 40. This validates the results of the steady state amplification scheme, derived in the previous section.

The main advantage of a diverging seed is that only a low energy is required to initially meet the condition (2.44) at the focus. This drastically reduces requirements on the seed pulse. Furthermore, the intensity at focus will be high enough to use the seed ionisation scheme, eliminating premature scattering of the pump pulse and greatly improving the amplified pulse contrast.

5.4 Effect of competing instabilities

The eventual limitation with the standard collimated scheme is that the beam waist is constant, so for the power to grow so must the intensity. Since the growth rates of detrimental instabilities increase with intensity [81], this means the pulse will eventually be degraded by instabilities. The maximum allowable compression ratio for the standard scheme was estimated in Fig. 4.13.

In the new scheme, the pulse intensity is constant but the power and energy increase as it diverges. The amplitude a_1 can be fixed at any desired value by changing the pump pulse or beam geometry. I now repeat the calculation of the maximum compression ratio for the case of an expanding beam with constant a_1 and T_1.

$$\Lambda_{RFS} = \gamma_{RFS}\frac{T_0}{2} = \frac{\omega_p a_1}{4}\sqrt{\frac{\omega_p}{\omega_0}}\frac{T_0}{2} \tag{5.27}$$

$$= \frac{5\omega_p}{16}\sqrt{\frac{\omega_p}{\omega_0}}\frac{T_0}{T_1\sqrt{\omega_0\omega_p}}, \tag{5.28}$$

$$\frac{T_0}{T_1} \simeq 3\Lambda_{RFS}\sqrt{\frac{n_{crit}}{n}}. \tag{5.29}$$

Here I have used the growth rate (2.33) and the non-linear stage relation (5.23). This shows the maximum compression ratio for the steady state amplifier is still limited to the same value as the collimated amplifier (compare to equation (4.34)). The maximum compression ratio is independent of κ and the pump pulse parameters.

The same calculation for the ponderomotive filamentation growth rate yields

$$\Lambda_{fil} = \gamma_{fil}\frac{T_0}{2} = \frac{a_1^2\omega_p^2 T_0}{16\omega_0} \tag{5.30}$$

$$= \frac{\omega_p^2 T_0}{16\omega_0}\frac{25}{4\omega_0\omega_p T_1^2}, \tag{5.31}$$

$$\frac{T_0}{T_1} = \sqrt{2\Lambda_{fil}\omega_0 T_0}\left(\frac{n_{crit}}{n}\right)^{1/4}. \tag{5.32}$$

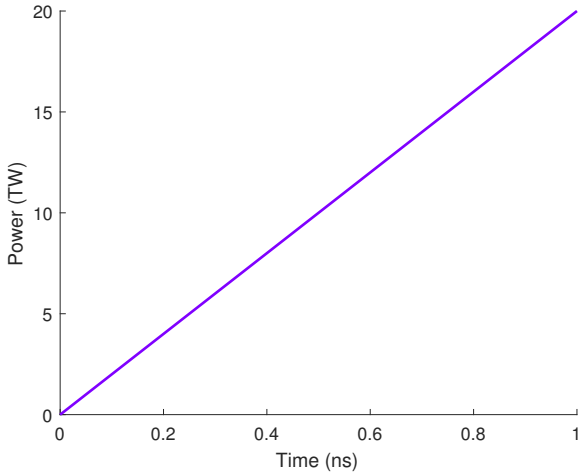


Figure 5.8: The required temporal power envelope of the pump pulse. This will lead to an amplified pulse with no trailing secondary peaks.

The maximum compression ratio before filamentation occurs is also the same as for the collimated amplifier.

A further advantage of the new geometry is the accessibility of the seed ionisation scheme [11, 13] with a low energy seed. For example, with a f/5 focus, a 50 fs seed pulse is capable of ionising a hydrogen gas with less than 1 mJ of initial energy.

The previous chapter suggested a combination of two amplifier stages of differing plasma density. However, the results of this chapter show that by exploiting the beam geometry it is possible to combine these into a single stage of uniform density, greatly reducing the complexity and better maintaining the resonance.

To summarise, my new scheme uses a diverging geometry and pulse shaping so that the pump pulse power increases $\propto t - t_f$, as shown in Fig. 5.8. Pulse shaping is routinely achieved for longer duration pulses on facilities such as the National Ignition Facility. The non-linear stage is reached straight from the start, since the seed has a high amplitude at its focus. The intensity and duration of the amplified pulse are then constant, but the power and energy increase $\propto (t - t_f)^2$ as it diverges away from the focus. The envelope consists of a single amplified peak, which reduces the wastage of energy into the usual trailing secondary peaks.

The new scheme presented in this chapter is completely my own work and has been submitted for publication [109]. The original idea was mine. The mathematical derivation, simulations and analysis were conducted solely by myself.

Chapter 6

Effects of Perturbations From Ideal Conditions

So far, the analysis has proceeded from the idealised coupled envelope equations for the Raman process. These equations make many assumptions about the physics that could overestimate the performance. Indeed, the large scale kinetic simulation study revealed that particle trapping and phase mixing often destroy the delicate plasma wave and reduce performance, especially at lower plasma density.

However, even these simulations assumed highly idealised conditions with perfectly uniform plasma and perfectly coherent, well defined laser pulses. In many experimental scenarios this will not be true and it is an important consideration for any experimental design to know the permitted tolerances. This could influence the uniformity of the gas target design, laser pulse quality and laser pulse spectrum.

6.1 Required plasma uniformity

Since Raman scattering is a resonant process, the highest growth will be when perfect resonance is maintained. For two coherent laser pulses, this will correspond to a given plasma density. However, as with any physical resonance, there will be some finite width of the resonance and so some level of tolerance in the allowable plasma density. This will apply separately at each position in time and space, so the process will also have some tolerance to plasma density gradients [120, 72, 52, 114].

The tolerance is easily estimated in the non-linear stage of Raman amplification [77]. Since the amplified pulse duration is given by equation (2.44), the bandwidth of the amplified pulse must have a minimum width given by the bandwidth theorem

(1.1). As the pulse amplifies, it shortens in duration and so must gain bandwidth in the spectral domain.

Using the solution for the non-linear stage (2.44) to replace T_1, the FWHM pulse bandwidth is

$$\Delta\omega_1 = \frac{4\ln(2)}{T_1} = 4\ln(2)\frac{a_1\sqrt{\omega_0\omega_p}}{5}. \tag{6.1}$$

This will apply in the non-linear stage only. The density tolerance δn can be defined through the loss in efficiency as the density is changed. I take δn to be the change in density when efficiency drops to half its value at perfect resonance. Since $n \propto \omega_p^2$, this becomes

$$\frac{\delta n}{n} = 2\frac{\delta\omega_p}{\omega_p} = -\frac{\Delta\omega_1}{\omega_p}. \tag{6.2}$$

Here I have used the frequency matching condition $\omega_0 = \omega_1 + \omega_p$ and assumed that the pump is monochromatic so that $d\omega_0 = 0$. There is also a factor of $1/2$ since $\Delta\omega_1$ is the FWHM, not the half width. Substituting in equation (6.1) gives

$$\frac{\delta n}{n} = \frac{4}{5}\ln(2)a_1\sqrt{\frac{\omega_0}{\omega_p}}. \tag{6.3}$$

The tolerance to density variations increases with the amplitude of the amplified pulse. This means the pulse will be more susceptible to density variations near the start of the amplifier than at the end. In addition, the tolerance is only weakly dependent on the resonant plasma density.

For example, a pulse with amplitude $a_1 = 0.1$ in the non-linear stage, with the density of $n = 0.01n_{crit}$ simulated in Chapter 4, will have a density tolerance of 17%. This is a large value and easily achievable experimentally.

However, the process may have a much lower tolerance in the early stages of the amplifier. In the linear stage of the instability, if there is a de-tuning given by $\delta\omega = \omega_1 + \omega_p - \omega_0$, the coupled equations show the wave de-tunes in a time $\delta\omega^{-1}$. For efficient amplification, this must be less than the growth rate and so the bandwidth of the instability is similar to the growth rate [21, 31]. Taking them equal, the bandwidth in the linear stage is

$$\delta\omega \simeq \gamma_{RBS} = \frac{1}{2}a_0\sqrt{\omega_p\omega_0}. \tag{6.4}$$

Equation (6.2) then gives the tolerance to density variations in the linear stage as

$$\frac{\delta n}{n} = -\frac{\delta\omega}{\omega_p} = \frac{a_0}{2}\sqrt{\frac{\omega_0}{\omega_p}}. \tag{6.5}$$

Across much of the high efficiency part of the parameter space, this is approximately 1-2%. Achieving this level of plasma uniformity is difficult with conventional targets. Any small signal that amplifies in the linear stage will grow at rate γ_{RBS} until it becomes detuned by a density change of around 1%.

This means there is a significant advantage to starting with an intense seed pulse [108, 107]. Although some authors have argued that the exponential growth in the linear stage means a low power, low bandwidth, long duration seed pulse will be effective, in reality the required plasma uniformity is difficult and the seed will be de-tuned before the non-linear stage. This will drastically lengthen the linear stage of the amplifier, reducing the overall efficiency. Interferometric diagnosis of the density profile in experiments such as reference [95] suggest the required plasma uniformity is not achieved. This means the growth of the seed through the linear stage will be much slower than exponential.

The low tolerance is another disadvantage to using a pre-heating pulse. It will sharply raise the plasma temperature, reducing the efficiency, but also cause thermal expansion of the plasma. This imposes density gradients that reduce the growth of the pulse in the linear stage. The seed ionisation scheme suffers neither of these problems.

The most effective route to avoid this problem is to start with a seed pulse intense enough to deplete the pump, meaning the large density tolerance of $\simeq 16\%$ is more applicable.

6.2 Effect of density shifts

Low amplitude seed pulses may have a greater tolerance to density variations if they have a wide bandwidth. Encountering different plasma densities at different times will cause separate parts of the seed bandwidth to be amplified in the linear stage. To show this, I performed a series of particle-in-cell simulations to explore the efficiency of Raman amplifiers for different levels of de-tuning. Particle trapping may also de-tune the interaction [131].

First of all, I investigated the effect of a constant shift in density δn, away from the resonance. The first set of simulations investigated the non-linear stage, using a pump pulse amplitude of $a_0 = 0.01$, a seed pulse amplitude of $a_1 = 0.1$ and a seed duration of $\omega_0 T_1 = 150$. This means the pulse met the non-linear condition from the start. The plasma temperature was 50 eV.

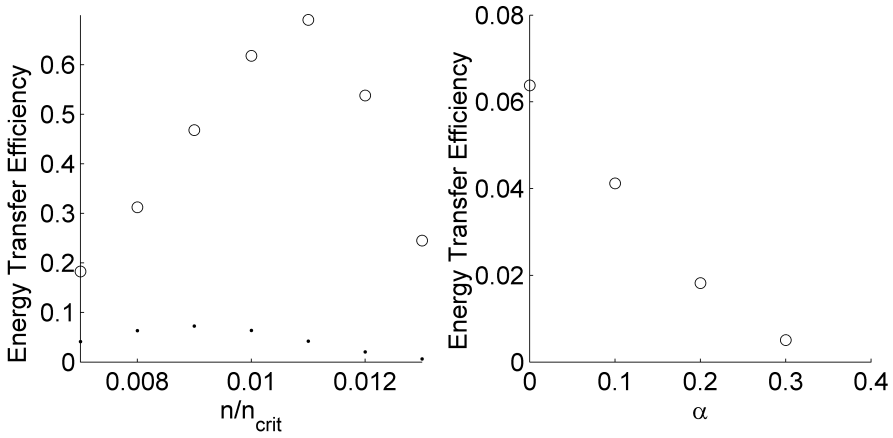


Figure 6.1: Effects of density variations, investigated with 1D particle-in-cell simulations. (a) The effects of a uniform shift to the density. The seed frequency was $\omega_1 = 0.9\omega_0$ and the plasma temperature was 50 eV. The pump had constant amplitude $a_0 = 0.01$. The total efficiency is given after a $300\lambda_0$ propagation. Open circles correspond to a high amplitude seed pulse $a_1 = 0.1$ in the non-linear stage and the points correspond to a low amplitude seed with $a_1 = 0.01$. (b) The effect of a plasma density gradient. The efficiency is shown versus the relative density change α across the $300\,\lambda_0$ propagation, for the lower amplitude seed.

These simulations are shown with open circles in Fig. 6.1a. The plot shows the total energy transfer efficiency after a 300 wavelength propagation. All parameters were kept constant except for the plasma density. The seed pulse frequency was $\omega_1 = 0.9\omega_0$, meaning the expected resonant plasma density is $0.01n_{crit}$.

The simulations show a broad resonance. The peak efficiency is found for a density slightly higher than the expected value of $n = 0.01n_{crit}$. This may be due to increased kinetic effects at lower densities causing a drop in efficiency. In addition, the Raman scatter growth rate increases with density and so there is an advantage to using a slightly higher density than the linear stage resonance.

However, the simulations were repeated in the linear stage, and this trend was reversed. The solid points in Fig. 6.1a show the evolution of an identical seed but with a peak amplitude of 0.01. As such, this seed did not meet the requirement (2.44) for the non-linear stage. This means the achievable instantaneous efficiency is much less.

For this lower amplitude seed, the peak efficiency is at a plasma density slightly lower than expected. This is due to the temperature dependent Bohm-Gross increase to the plasma wave frequency, meaning the resonance is slightly shifted.

Including this effect, the predicted resonance is at

$$\omega_0 - \omega_1 = \omega_p \sqrt{1 + 3k_2^2 \lambda_D^2}. \tag{6.6}$$

The matching conditions give $k_2 \simeq 2\omega_0/c$ and $\lambda_D^2 = k_B T/(m\omega_p^2)$. This leads to

$$\omega_0 - \omega_1 = \omega_p \sqrt{1 + 12\frac{k_B T}{mc^2}\frac{n_{crit}}{n}}, \tag{6.7}$$

$$\frac{\omega_1}{\omega_0} = 1 - \sqrt{\frac{n}{n_{crit}} + 12\frac{k_B T}{mc^2}}. \tag{6.8}$$

For the conditions simulated, with $T = 100\,\mathrm{eV}$, this gives a resonant plasma density of $0.0088 n_{crit}$, compared to $0.01 n_{crit}$ without the correction. It is this effect which means the peak efficiency is at slightly lower than the expected density. This estimate is in excellent agreement with the simulation, which gives the peak efficiency in the linear stage for a density of $0.009 n_{crit}$.

The width of the resonance is broad for both the linear and non-linear cases. This is because the seed had the same bandwidth in both cases. It suggests that using a broad bandwidth seed is an effective strategy for increasing the tolerance to density variations in the plasma. The tolerance of approximately 20% agrees with the estimate (6.3) which gives 17%.

6.3 Effect of density gradients

Another important question is the effect of density gradients. Since gas targets often expand from an aperture, the density may vary spatially at the time the laser pulse arrives. If there is a gradient in plasma density, this will mean any Raman scatter will no longer be resonant once the density changes.

To investigate this separate effect, the simulations for the linear stage case were repeated with a plasma density given by $n_0(1+\alpha x/300\lambda)$ with $n_0 = n_{crit}(1-\omega_1/\omega_0)^2$. This means the density starts at the resonant value and varies linearly by a fraction α over the $300\lambda_0$ propagation. The results are plotted in Fig. 6.1b. Note that the point with $\alpha = 0$ corresponds to the uniform plasma point in Fig. 6.1a. Importantly, there is a significant tolerance to the density gradients on large spatial scales, even below the non-linear stage. The tolerance of 15% is similar to that for a uniform density shift.

The resonance condition is met separately at each point in the plasma. This means there is the same tolerance to density gradients as there is to a uniform shift

in density [102, 58]. As long as the density gradient scale-length is much longer than the extent of the amplified pulse, the plasma wave will be resonantly driven for the full duration of the seed pulse. This means the tolerance to density gradients is no worse than for a uniform density shift. The changing density will simply amplify different parts of the seed pulse spectrum.

The Rosenbluth gain threshold for the Raman instability in a density gradient was found in references [101, 64] to be

$$a_0^2 k_0 > \frac{1}{n}\frac{dn}{dx} \tag{6.9}$$

$$\alpha < 600\pi a_0^2 \tag{6.10}$$

If α exceeds this value, Raman scatter in the linear stage will be de-tuned. For the simulated parameters, this gives $\alpha > 0.19$. This agrees with the de-tuning and lower growth of the wider bandwidth seed pulse in Fig. 6.1b.

De-tuning the density in this way has previously been suggested as a method to reduce growth of Raman forward scatter, as the amplified pulse is more resistant due to its wider bandwidth. The amplified pulse grows strongly, even with a plasma density gradient. This was previously predicted and now observed in these simulations, up to a limit broadly agreeing with the result (6.3).

In addition to the advantages outlined in Chapter 4, shorter seed pulses have the advantage of greater resilience to plasma density gradients. This is because in a given density gradient, longer pulses will experience a larger density change across their extent, so the plasma wave will suffer more de-tuning. To maintain the efficiency, the plasma wave must be resonantly driven across the full extent of the seed. For example, a seed pulse in the linear stage requires plasma density uniformity to the 2% level derived in equation (6.5). The plasma must be uniform to this level over a length cT_1, or the plasma wave will become de-tuned within the extend of the amplified pulse. For shorter seed pulses, the uniformity is only required over a much smaller distance, so the allowable gradients should be larger. This is why the tolerance is so great in the simulations of a low amplitude seed in Fig. 6.1.

This consideration also helps with understanding the coherence requirements of the pump pulse [91]. For interactions at constant plasma density, $d\omega_0 = d\omega_1$. As long as the coherence time of the pump pulse is much longer than the seed, the

interaction will therefore be unaffected and the tolerance to density perturbations will be given by equation (6.3).

Transverse speckles in the pump pulse, due to induced spatial incoherence, may also play a part. If phase plates are used to increase the interaction volume, the pump amplitude may vary transversely, leading to spatially dependent Raman growth rates. These speckles will then be transferred to the seed beam [104], affecting its focal quality and seeding the filamentation instability.

6.4 Seeding of competing instabilities

As well as lowering the energy transfer efficiency, any spatial non-uniformity could seed filamentation of the pulse [115]. For this reason, transverse density modulations are also detrimental. This effect was investigated using a two-dimensional particle-in-cell simulation. The plasma density was initialised using a Fourier series

$$\frac{n(\mathbf{x})}{n_0} = 1 + 0.01 \sum_{j=1}^{10} \sin(2\pi j \mathbf{k_j}.\mathbf{x}/L), \tag{6.11}$$

where $L = 100\lambda_0$ is the simulation box width, $n_0 = 0.01 n_{crit}$ is the resonant density and the $\mathbf{k_j}$ are randomly directed unit vectors. This gives a wide range of spatial modes from the full box width down to the plasma wavelength. The peak amplitude of the modulation is 10%. Fig. 6.2a shows the plasma density profile that was initialised using this Fourier series. Although the pattern is periodic, the random alignments of each mode model a realistic experimental scenario. This will lead to longitudinal and transverse density gradients, some with scale length similar to the seed pulse duration.

The simulation used a moving window domain with length $125\lambda_0$ and width $100\lambda_0$. The boundary conditions were periodic, with cell size $\lambda_0/60 \times \lambda_0/5$ and 36 electron particles per cell. Ions were held static to reduce computational costs. The pump and seed had plane wave transverse profiles and $\omega_1 = 0.9\omega_0$. The initial seed had duration $25\lambda_0/c$ and peak amplitude $a_1 = 0.01$. The pump pulse had a constant amplitude of $a_0 = 0.02$.

The seed pulse was amplified for a propagation of $2400\lambda_0$ to simulate the filament growth over long durations. The final intensity and phase are shown across the transverse dimension in Figures 6.2b,c. These line-outs were taken at the position

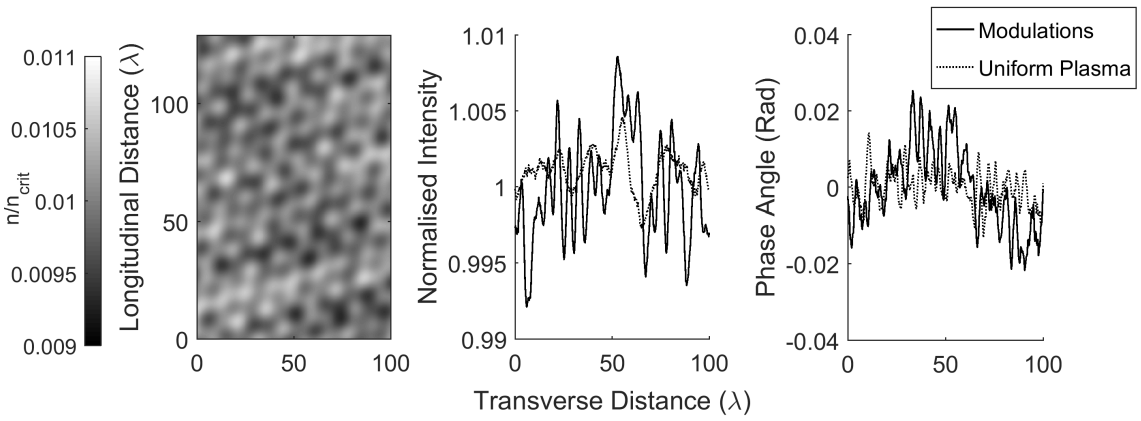


Figure 6.2: Simulated effects of spatial density modulations on the growth of ponderomotive filamentation. (a) The plasma density that was initialised in the 2D particle-in-cell simulation. The resonant density is $0.01n_{crit}$ and the simulation domain moved with the seed pulse. The pump had constant amplitude 0.02 and a plane wave transverse profile. The seed pulse had a plane wave transverse profile, frequency $0.9\omega_0$ and duration $25\lambda_0/c$. (b) A line-out of the intensity across the transverse dimension, at the position of peak intensity, after a $2700\lambda_0$ propagation. This is compared with the equivalent plot for the control simulation with uniform plasma density $0.01n_{crit}$. Panel (c) shows the phase of the wave across the transverse dimension, at the same time and longitudinal position.

of peak intensity. Also shown are the results of a similar control simulation with uniform density n_0.

Comparing to the case with uniform density, the presence of the density modulations has not introduced any additional transverse filamentation. The phase of the wave is largely unaffected. The amplitude is also still very uniform across the transverse dimension, although there are some filament modulations on a similar spatial scale to c/ω_p. The average intensity was less for the case with density variations. This was due to the loss in efficiency outlined in the previous section. The intensity has been normalised to compare the filaments.

Even with 10% density variations with many spatial modes, the amplification still proceeds, although the efficiency is lowered and there is some slightly enhanced transverse filamentation. This is promising for experimental designs, since this tolerance to spatial density variations should be achievable using gas cell targets. This is an important result for subsequent focussing of any amplified pulse, since transverse coherence is still retained and the focal spot should be similar to the case with no plasma interaction.

These results help to explain why previous experiments have not achieved efficiency close to η_{max}. In addition to the losses due to kinetic effects outlined in

Chapter 4, the amplification will become de-tuned by density gradients. The experimental interferometric results in reference [98] show that the plasma electron density varies from $1.2 - 1.6 \times 10^{19}\,\mathrm{cm}^{-3}$ across the 4 mm interaction. This is a variation of 14%, far more than the estimated tolerance of 2% to amplify a narrow bandwidth seed pulse in the linear stage, as the authors attempted.

In fact, this information may be used to estimate the experimentally measured efficiency. Equation 6.5 suggests the amplified pulse is amplified only at densities with 2% of its resonant density. Since the experimental density variation was 14%, the efficiency should only be 2/14 of its maximal value of 50%, found in the simulations in Chapter 4. This gives (2/14)*0.5 = 7%, very similar to the experimentally measured value of 6.4%. This suggests that de-tuning due to density gradients was at least partly to blame.

It seems that typical experimental constraints will not hinder the amplification in the non-linear stage. These new results, published in reference [108], are in agreement with the results of other authors [41], who showed how the pulse coherence is retained. In addition, I have now verified the expressions (6.3) and (6.5) for the tolerance to density variations.

The linear stage will be more adversely affected, unless the Rosenbluth gain condition is met. The results of this chapter give another argument for starting the seed pulse in the non-linear stage, as it will be more robust to density perturbations. This may be easily achieved using the new beam geometry given in Chapter 5.

Chapter 7

Seeding the Amplifier via Raman Forward Scattering

In Chapter 2, the simulations showed that a short, high intensity seed pulse has the highest performance. Longer pulses, especially those longer than γ_{RBS}^{-1}, tended to break into two amplified pulses of similar intensity. This reduces the peak intensity reached by the amplifier and produces an unwieldy, unpredictable output with large shot to shot variation.

It was found that seed pulses shorter than γ_{RBS}^{-1} improved both the efficiency and the peak intensity. Unfortunately for experiments, generating such a pulse is a difficult proposition, especially due to the shifted wavelength. Many previous experiments have generated a seed pulse through external methods, such as non-linear optics effects in gas cells [8]. However, the resulting pulses become lengthened by the gas cell and must be low power to avoid ionisation of the gas. These are not desirable characteristics [138].

Seed pulses could also be generated through solid state optical parametric amplifiers. This would allow more powerful, shorter seed pulses and more spectral control. This method has the disadvantage of requiring extra pumping lasers and a separate beam-line to develop the seed, increasing the complexity and timing issues. I now present a novel mechanism for seeding the amplifier, based on plasma Raman forward scattering of an ultra-short pulse. The scheme was tested at the Astra laser at the Central Laser Facility in the United Kingdom.

7.1 Seeding scheme

A schematic of the scheme is shown in Fig. 7.1. First, a high-power femtosecond laser pulse of frequency ω_0 is passed into a uniform plasma. It propagates for several

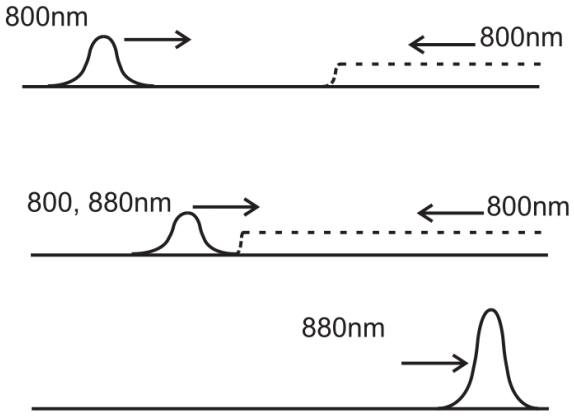


Figure 7.1: Schematic of the proposed seeding scheme. A high intensity femtosecond pulse excites Raman forward scatter in an initial section of uniform plasma, which then interacts with a counter-propagating pump of the same wavelength as the original seed. This allows generation of both laser pulses from a single beam-line, drastically reducing complexity and timing issues.

growth times of the Raman forward scatter [110, 15, 87, 134]. The growth rate of Raman backward scatter is greater, meaning much of the energy could be lost to spontaneous backwards scatter. However, keeping the seed pulse short will minimise this. With a plasma density of $n = 0.01 n_{crit}$ and a pulse amplitude of $a_1 = 0.1$, the linear backward scatter growth time is 30 fs. This means the pulse must be an ultra-short femtosecond pulse, such as those delivered by the Astra laser at the Central Laser Facility.

After this initial stage, the spectrum of the pulse will have a peak at $\omega_0 - \omega_p$ due to the forward scatter. The duration will be similar to the starting pulse. This light will act as a seed for the amplifier. At this point, a long pump pulse of frequency ω_0 overlaps this seed beam and the amplifier continues as in the standard scheme. Since the seed and pump beams have the same wavelength, they may be generated from the same laser system, minimising the cost and complexity. However, it is unknown whether the coherence and energy of the forward scattered seed light will be sufficient.

7.2 The Astra laser and set-up

To answer these questions, I tested the forward scatter stage of the scheme at the Astra laser facility in November 2016. The Astra laser is a titanium-sapphire amplifier that delivers a broad bandwidth pulse around 800 nm wavelength. It

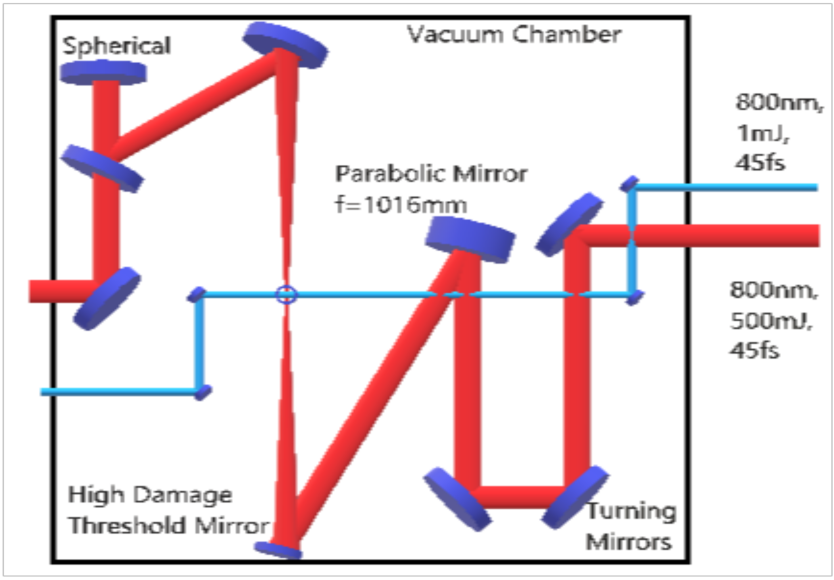


Figure 7.2: Schematic of the vacuum chamber. The 484 mJ, 11 TW main beam (red) was focussed by an f/17 off axis parabola to a 3 mm gas jet. The transverse interferometry probe (blue) was generated on a separate table and synchronised to arrive at the target with the main pulse. Both had wavelength 800 nm and duration 45 fs. The target was a 3 mm diameter gas jet at the focus.

uses a conventional chirped pulse amplification chain, delivering the stretched pulse at around 1 J to the target area. The compressor chamber and target chambers are adjacent and directly linked by a vacuum tube. The compressed pulse had a transform limited half width at full maximum duration of 45 fs and the pulse was focussed by a f/17 parabola with focal length 1016 mm. Since this focal length is longer than the target chamber, a high damage threshold mirror is used to reflect the pulse before focus. This is shown in the schematic Fig. 7.2.

The target was a circular 3 mm diameter methane gas jet nozzle. The nozzle axis was tilted 5° away from vertical. After the interaction, the main beam reflected from two fused silica mirrors with reflection coefficient approximately 10%. The light was then collimated with a spherical mirror of focal length 762 mm. The light then passed back through the second fused silica mirror, reflected from a third fused silica mirror and out through the target chamber window.

The outgoing beam was split and focussed to two charge-coupled devices (CCD), with one imaging a position within the Rayleigh range and one outside of the Rayleigh range. The light was also coupled to the slit of an Andor Shamrock-303I

optical spectrometer with an Andor Newton 920 electron multiplying CCD camera. The pixel spectral resolution was 0.13 nm and the grating spectral resolution was 0.21 nm from a 600 lines/mm grating. The spectral range was calibrated using an atomic emission source, giving a spectral range from $762.2 - 893.3$ nm. This allowed simultaneous measurement of the laser spectrum and any Raman scattered light.

The laser pulse energy was calibrated using a calorimeter. Since the calorimeter cannot be placed at the target position under vacuum, an indirect method was used. First, the uncompressed beam was measured at full power and low power. The average value was taken over several shots. Next, the compressed low power pulse was measured at the target position. Assuming the compression efficiency is independent of power, the energy of the high power pulse is then the ratio of the first two measurements, multiplied by the third measurement. This yields a value of 484 mJ$\pm$ 9.0%, where the error was primarily from the shot to shot variation between the three measurements. The peak laser power was 11 TW$\pm$ 9.0%.

The laser pulse intensity was maximised by adjusting the main parabolic mirror. The focal spot within the Rayleigh range was observed by reflecting the main beam to a high magnification objective and CCD within the target chamber. The magnification was calibrated by translating the camera a known distance, using a micrometer, then measuring the translation of the CCD image in pixels. This gave a calibration of 0.97 μm/pixel. The measured focal spot profile had a FWHM of 30 $\pm$ 1pixels or 29 μm$\pm$ 3%. This gives the focal beam waist as 24.6 μm $\pm$ 3%, the Rayleigh range as $2L_R = 4.8$ mm and the peak intensity was 1.2×10^{18} Wcm$^{-2}\pm$15%. The peak dimensionless laser amplitude was $a_0 = 0.75 \pm 7\%$.

The Rayleigh range was centred over the target by imaging the position of best focus on the exit mode camera and moving a thin wire until it was well imaged in this location. The target was then centred on this location by using an alignment camera vertically above the target.

To assess the plasma free electron density, I also used an interferometry diagnostic with a 45 fs, 800 nm wavelength probe pulse, passing transversely through the target in a collimated beam. The beam diameter was 1 cm and it was timed to arrive within 1 ns of the main pulse by using an optical photo-diode sensor and fast oscilloscope. Fine adjustment of the synchronisation was performed by observing the images of the plasma and adjusting the probe delay accordingly. The target was

imaged on to an Allied vision Pike F-032 CCD, with pixel size 7.4 μm and 640×480 pixels.

Before the image plane, the pulse was passed through a Mach-Zehnder interferometer and the images from the two separate arms were offset on the CCD, such that one image of the plasma overlapped with the region outside the plasma on the other image. This introduced an interference pattern between the two images. The relative delay on the two arms of the interferometer was adjusted to zero, such that the contrast of the fringes was maximal in the region of interest.

7.3 Interferometry results

The raw interferometry data for a typical shot is shown in Fig. 7.3a. The methane gas target had a backing pressure of 14 Bar$\pm$3%.

First, the phase shift as a function of position was extracted using a horizontal Fourier transform, selecting only the peak at positive k, corresponding to the fringes. This peak was shifted to $k = 0$ and inverse transformed, giving the complex phase as a function of longitudinal and transverse distance. The phase of the complex number is wrapped between $-\pi$ and π, exhibiting artificial discontinuities of 2π along branch cut lines in the image. To remove this, a 2D phase unwrapping algorithm was used to adjust for the discontinuities. This image is shown in Fig. 7.3b for a typical shot.

Retrieving the electron density from the phase shift requires making an assumption about the extent of the plasma into the other transverse dimension x, along the propagation of the probe pulse. Assuming the plasma is axially symmetric, the electron density as a function of radius can be extracted. However, deviations from axial symmetry can quickly introduce errors in the analysis, as well as being sensitive to noise. The electron density was therefore extracted by assuming the electron density profile had the same width in the x direction as the phase shift profile in the y direction. Line-outs of the phase shift were taken for each z position and the maximal phase shift ϕ_{max} and FWHM length L_x were determined. The electron density was then calculated as $n = n_{crit} \frac{\phi_{max} \lambda}{\pi L_x}$.

The peak electron density on axis was $1.5 \times 10^{19}\,\mathrm{cm}^{-3} \pm 20\%$. The error was determined by comparing the optical resolution to the fringe separation. The density drops towards the edge of the gas jet, however there is a significant distance of almost uniform plasma to excite Raman scattering. Across several millimetres, the

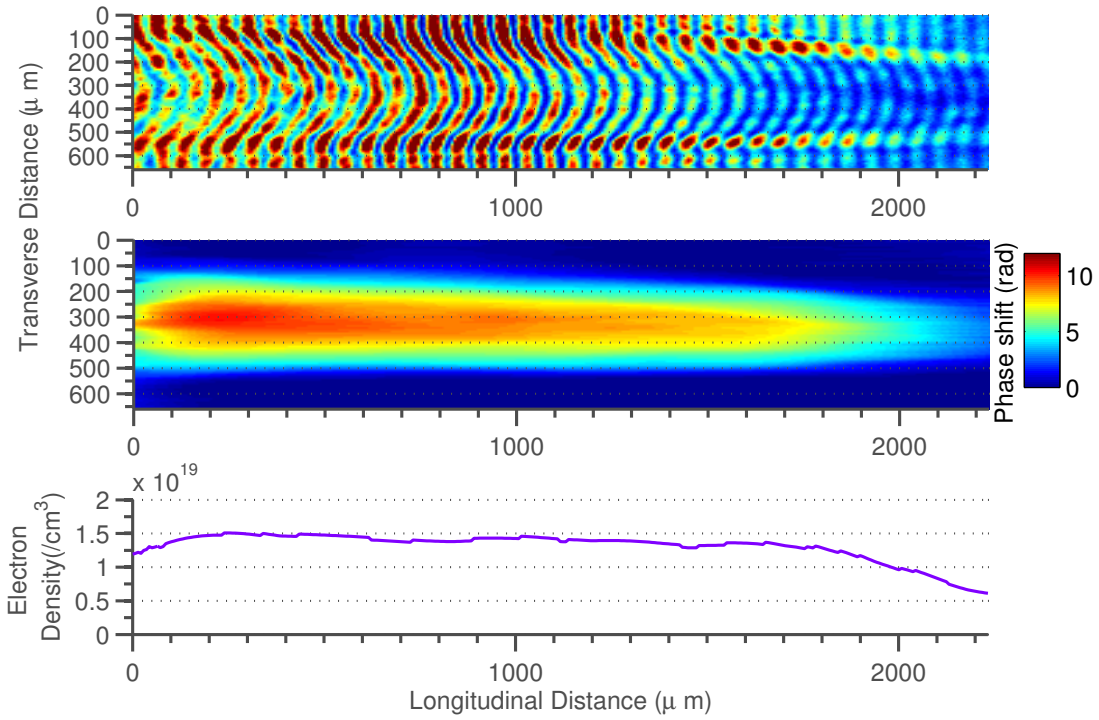


Figure 7.3: Results of the interferometry diagnostic within a methane gas target. (a) The raw interferometry data, taken from a side-on image of the plasma. The main laser pulse propagates towards the left. The pixel resolution is 7.25 μm and the optical resolution for the imaging system was 20 μm. The gas nozzle is centred on $z = 500\,\mu$m and has radius 1500 μm. The probe pulse, with wavelength 800 nm, passed transversely through the plasma. It was interfered with itself in a Mach-Zehnder interferometer. (b) The phase shift across the image, extracted via a horizontal Fourier transform. (c) The on axis electron density.

plasma density variation is smaller than the tolerance for the linear stage given by equation (6.5), which gives approximately 10%.

The plasma transverse width is approximately 400 μm which is much wider than the FWHM of the laser pulse. This is because the ionisation intensity is much less than the half-maximum intensity, so even the light far out in the wings of the focal spot is intense enough to ionise the methane.

It is interesting to note that methane gas ejected into a vacuum can form solid density nano-metre size clusters, bound by Van der Waals forces. This means the ionisation mechanism will more typical of a solid density target. Plasma clusters would ionise through collisional ionisation at much lower intensity than a uniform low density gas, which would field ionise. This may explain why the plasma in shots with a methane target was observed to be twice as wide as the plasma formed using a uniform nitrogen gas target.

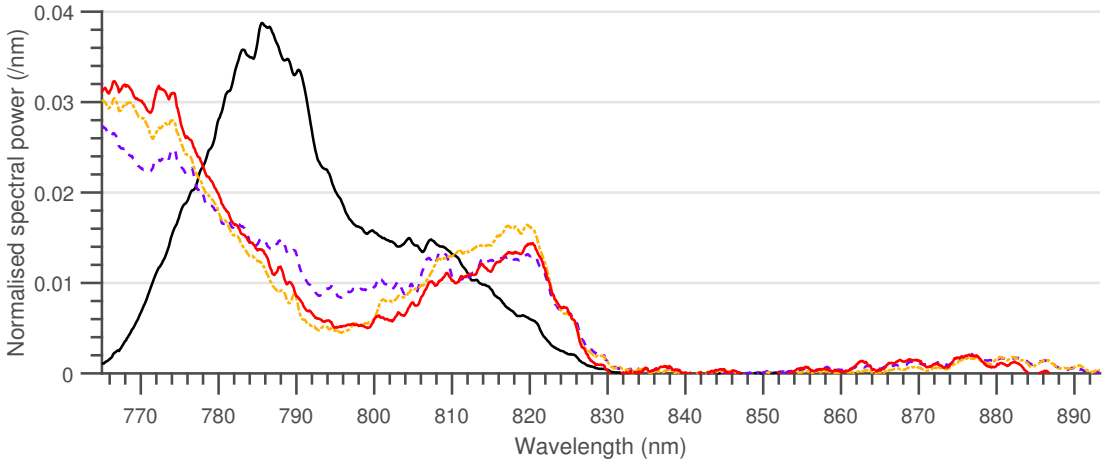


Figure 7.4: Spectra of the light collected from the target in the forwards direction. The black solid line shows a reference spectrum of the laser pulse with no plasma interaction. Spectra from three selected laser shots are shown in different colours. The spectra have been smoothed over a 1 nm width and normalised such that the integrated spectral power is 1. The blue shifted laser light around 770 nm is due to the ionisation front. The red shifted laser light at 820 nm is due to the wake-field density gradient. The small forward Raman scatter peak was observed at 880 nm.

7.4 Measurement of the Raman scatter

The linear forward Raman scatter growth distance is 25 μm, meaning there will be many exponentiations over the 3 mm gas jet and the Raman peak will grow from noise to a large amplitude. There could also be significant amounts of Raman scatter in the backwards direction. The back-scatter was not measured.

The forward scattered spectrum is shown in Fig. 7.4 for several separate laser shots. Also given is a reference laser spectrum with no plasma target. The spectrum of the laser pulse is altered by the plasma interaction. The laser bandwidth splits into two peaks, one caused by a density gradient at the ionisation front and the other by a gradient from the ponderomotive pressure. The ionisation front will also seed the forward Raman instability [45]. This process detrimentally affects laser wake-field accelerators, but is beneficial in this application.

There is also a small red-shifted Raman peak in the spectral region of 880 nm. This was evident on all shots, although the spectral power varied. The Raman peak had a wide bandwidth of 20 nm±30%. This compares with the linear stage analytical estimate of $\Delta\omega/\omega = \gamma_{RFS}/\omega_1 \simeq 0.005$ or $\Delta\lambda = 4.5$ nm. The values are inconsistent due to two reasons. Firstly, the linear stage derivation assumed a zero bandwidth pump pulse with infinite extent, whereas the experimental laser

pulse had a wide bandwidth of > 20 nm. Secondly, the laser pulse was of relativistic intensity and relativistic corrections were not considered in the derivation of Chapter 2.

The central wavelength of the Raman peak is 878 $\pm$3nm and that of the laser is 793 nm. The interferometry diagnostic gives the peak density as $0.0083 n_{crit} \pm 20\%$. Using that $\omega_1 = \omega_0 - \omega_p$ gives an expected Raman wavelength of 872 $\pm$ 8 nm. This estimate from interferometry is consistent with the measured peak in the spectrum.

Since the scattered light may not have the same angular distribution as the incoming light, the spectra of different shots cannot be absolutely calibrated. The total Raman scattered energy should be measured by integrating across the full aperture using a camera, rather than the spectrometer.

7.4.1 Energy and profile of the forward scatter

The energy of the Raman light was measured by a two stage process. We used a laser pulse calorimeter, oscilloscope and a long pass spectral filter to transmit only light above 850 nm wavelength. First, the main pulse energy was absolutely calibrated as discussed previously, yielding a value of 484.0mJ$\pm$9%. Second, the CCD after the interaction was used to measure the three different signals:

- The main beam with no target and no filter
- The main beam through the spectral filter with no target
- The light from the plasma target through the spectral filter

Each signal was integrated across the full aperture. The total signal of the Raman near-forward scatter is the third measurement minus the second one. This eliminates the main laser light which leaks through the filter. The signal is then divided by the relative quantum efficiency of the CCD, since the sensitivity is slightly less at the higher wavelengths of the Raman peak. I then found the fractional conversion efficiency by dividing by the first measurement. The optics had almost constant reflection and transmission across the spectral range, so no correction was needed.

Fig. 7.5 shows the transverse profile of the Raman light, imaged outside of the Rayleigh range. Comparing to the reference shot with no plasma interaction, the Raman light seems slightly offset to one side. The offset was consistent for over

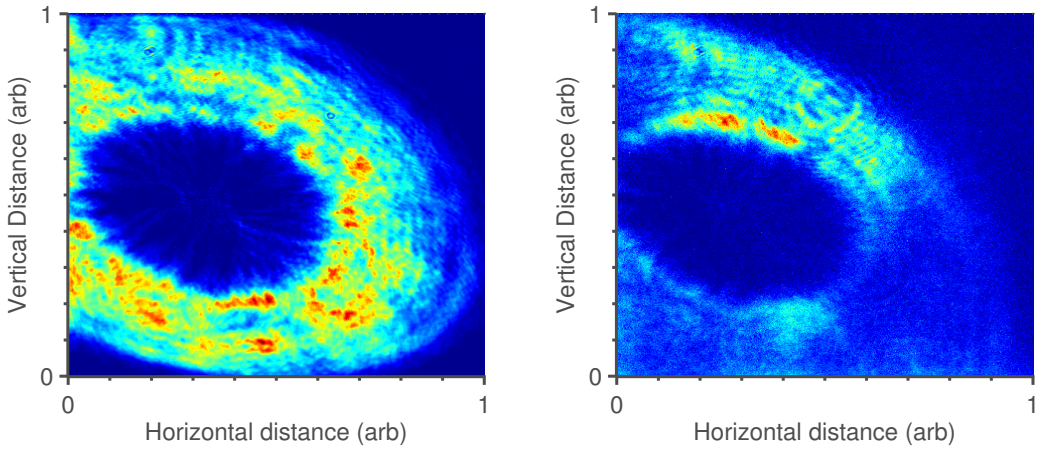


Figure 7.5: The transverse profile of the light collected in the forwards direction, imaged outside of the Rayleigh range. Note that due to the collection optics, there is a circular central region where no light was collected. (a) The main laser pulse with no target or spectral filter. (b) The light from the plasma interaction, collected through a long pass filter for 850 nm+. The colour scale has been adjusted for each plot, since the energy of the Raman light is much less than the original laser pulse.

50 laser shots, suggesting it was due to refraction from a plasma density gradient rather than a plasma instability. The Raman light is fairly uniform transversely, although there are a few speckles with intensity 2-3 times the average intensity.

The energy in the Raman forward scattered peak was 55 mJ±30%. This is 11% of the in-going laser energy. The Raman light has enough energy and sufficient transverse beam quality to be suitable as a seed pulse for a secondary stage amplifier. Shorter, more intense seeds have the best performance and so I used a frequency resolved optical gating (FROG) diagnostic to measure the pulse temporal profile.

A small sample of the collimated output beam was extracted and exited the chamber through a thin window. This was so the group velocity dispersion in the window would not adversely affect the duration of the pulse. The 850 nm+ spectral filter was used to isolate only the Raman light, which was subsequently analysed by the FROG or GRENOUILLE diagnostic as discussed in Chapter 3. An algorithm called FretGreen was used to invert the frog trace. The reverse process was performed to ensure a close match to the measured trace.

The temporal envelope of the Raman light for a typical shot is shown in Fig. 7.6. The FWHM duration is $T_1 = 150$ fs. Note this is an upper bound, since spectral dispersion in the window and spectral filter may have lengthened the pulse. The pulse is significantly shorter than seed pulses used in many previous Raman amplification experiments and the energy of 55 mJ is relatively high. The pulse

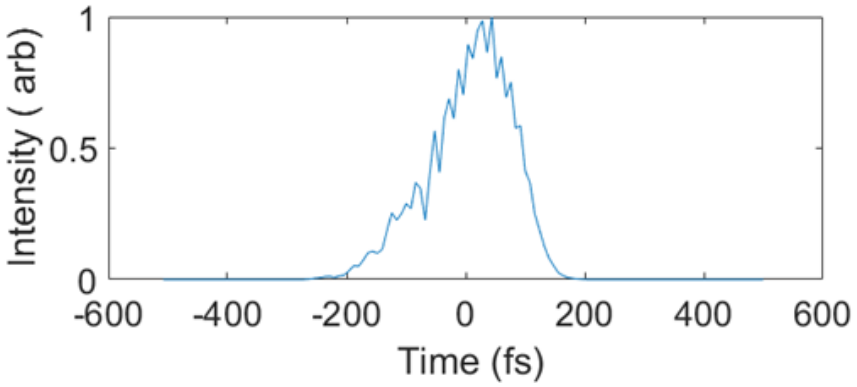


Figure 7.6: The inverted GRENOUILLE data for the forward scattered light, giving the intensity as a function of time. A spectral filter was used to isolate the Raman light only. The spectral range of the diagnostic was centred on 880 nm. The algorithm FretGreen was used to invert the data.

duration is shorter than the backwards Raman growth time for a typical pump pulse, $\gamma_{RBS}^{-1} \simeq 300\,\text{fs}$. As found in Chapter 4, these are all favourable qualities for effective seeding of a Raman amplifier.

This technique has several advantages over other schemes for generating Raman seed pulses. Previous experimental groups have used either Raman shifter crystals [8, 98] or optical parametric amplifiers [94]. The first method gives a fixed wavelength, at approximately 880 nm from a 800 nm pump. The second method has the advantage of being tune-able from 600 nm to 1300 nm and giving a broad bandwidth of up to 100 nm. However, it is more complex and difficult to align. These experiments used seeds with energy less than 1 mJ and duration approximately 500 fs [92]. The seed pulses developed in this chapter improve on the energy by a factor of more than 50, as well as automatically being resonant for Raman back-scatter with the pump pulse. Furthermore, the duration is a factor of 5 less than in previous experiments. This results in a much greater seed intensity, which Chapter 4 showed leads to a greater efficiency.

In summary, it may be possible to seed the amplifier using Raman forward scatter from a plasma, rather than a Raman shifter crystal as used in many previous experiments [98]. This experiment showed that plasma Raman forward scatter of a femtosecond relativistic pulse has many of the desirable characteristics found in Chapter 4. These include short duration, mJ energy and smooth spatial profile. The scheme could allow effective seeding of Raman amplifiers using the output of only one laser system.

This experiment was my original work, with experimental support from the research group. The idea for seeding via forward scattering was devised and appraised by Peter A. Norreys and Raoul M. G. M. Trines. The experimental analysis was conducted by myself, with a publication in preparation.

Chapter 8

Experimental Studies of Raman Backscatter

In this chapter I describe an ongoing experimental campaign at the Laboratory for Laser Energetics (LLE) in Rochester, NY, USA. The simulation results in previous chapters have been applied to optimise the new experiment. Although the full Raman amplifier scheme will not be in operation until 2019, this chapter presents promising preliminary results using only the pump pulse.

8.1 Experimental objectives

Previous experiments of Raman amplification have managed to access the non-linear stage and demonstrate an overall efficiency of 6.4% [98]. Key aspects of the non-linear stage were observed, including broadening of the amplified pulse bandwidth, shortening of the amplified pulse and measurable pump depletion. Despite these advances, the pulses had sub-terawatt power, far below the desired petawatt range. Future experiments must not only improve the efficiency but also scale up the amplifier to larger beam diameters. This presents two challenges. The first is maintaining a uniform wave-front across a wide aperture, which is typically difficult and requires optics to correct phase perturbations. Secondly, the plasma target must be uniform across a large volume, even after thermal expansion during the interaction.

Although the proof of principle has been demonstrated, experiments are still a long way from surpassing traditional chirped pulse amplification methods to reach petawatt powers. The goal of the present experiment is therefore to improve the efficiency of the amplifier beyond 6% and scale it up to terawatt or even petawatt powers.

This should be achievable by using the numerical results of previous chapters to optimise the efficiency of the amplifier. The study will also help design its various parameters such as the length, density and uniformity of the plasma and the amplitude, chirp and duration of the two laser pulses.

The numerical results in previous chapters have helped to identify the optimal part of the parameter space, information which is now built into the design of the new $1M scale experiment. The target parameters from Chapter 4 are $n = 0.02 n_{crit}$ and $a_0 = 0.02$, with a seed pulse shorter than $\omega_0 T_1 = 500$ and a power as high as possible. These requirements have been fully satisfied by the design parameters which are $\lambda_0 = 1054\,\mathrm{nm}$, $E_0 = 50\,\mathrm{J}$, $a_0 = 0.02$, $T_0 = 25\,\mathrm{ps}$, $\lambda_1 = 1100 - 1200\,\mathrm{nm}$, $n = 0.015 n_{crit} \simeq 1.5 \times 10^{19}\,\mathrm{cm}^{-3}$, $E_1 = 50\,\mathrm{mJ}$, $T_1 = 100\,\mathrm{fs}$, $a_1 = 0.01$ and $L = 4\,\mathrm{mm}$. Also in accordance with the findings of Chapter 4, the plasma will not be pre-heated by a third laser pulse. It will be ionised by the pump pulse.

Of further importance, the seed pulse is powerful enough to test the seed ionisation scheme, for the first time. This means the plasma will not be ionised until the seed pulse arrives, avoiding any heating and so maximising the efficiency.

The experiment has a timescale of several years and has a dedicated target chamber, allowing a full exploration of the non-linear stage of Raman scattering and as-yet untested schemes such as using the seed pulse to ionise the plasma, or using a diverging beam. The new set-up has several important advances over many previous experiments:

1. The pump pulse energy is much greater, up to 50 J
2. Adaptive phase correction optics will allow operation outside of the Rayleigh range, with much wider beams
3. The seed pulse power is expected to be much greater, around 1 TW
4. The seed pulse is powerful enough to ionise the plasma, allowing the seed ionisation scheme
5. The seed pulse is tune-able, allowing a wider scan of the density
6. The target is a hydrogen gas cell, allowing highly uniform and always fully ionised plasma [4, 126]

7. The diagnostic suite will measure the plasma temperature and use time resolved measurements

The flexibility, characterisation and long timescale of this experiment should help to isolate what caused the low efficiency in many previous experiments and help to improve it drastically.

8.2 The MTW laser facility

The multi-terawatt laser is a chirped pulse amplification laser system that was built to test laser technologies for the larger OMEGA EP system at LLE. It has since become an important laboratory for testing diagnostics, trialling new targets, training young scientists and is now used as a driver for more modern laser amplifiers and laser development. In the MTW laser, the initial high bandwidth pulses are selected from a mode-locked oscillator and a chirp is induced to stretch the pulse to several nanoseconds. From here, an optical parametric chirped pulse amplification system is used to increase the energy to approximately 200 mJ, before passing to Neodymium glass laser amplifiers. There is initially a set of rod shaped amplifiers. As the energy increases, the beam diameter must be increased due to damage thresholds, however a wide diameter rod is difficult to pump evenly across its interior, so the wide diameter beam must pass through slab shaped amplifiers which can be pumped from the ends as well as the side. Although the pulse energy increases, the laser bandwidth decreases to approximately 3 nm in this stage.

Depending on the number of amplifiers used, the beam now has energy between $0.2 - 50$ J and passes into a vacuum compressor chamber. The large gratings in the chamber introduce a frequency dependent phase delay. The compressor can achieve close to the transform limited duration of $\simeq 0.5$ ps by introducing a phase delay that exactly cancels that of the stretcher. However, for this experiment, the compressor was set to give a full width half maximum duration of 25 ps, meaning there was some residual chirp. This beam will be used as the pump pulse for the Raman amplifier.

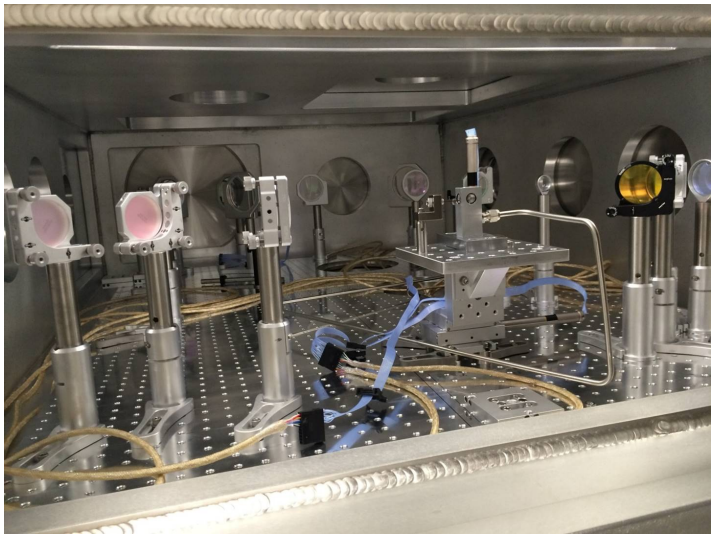


Figure 8.1: Interior of the target vacuum chamber, showing the gas cell target, pump beam optics and collection optics for the interferometer. Some of the backscattered light was also collected for spectral analysis.

8.3 OPAL beam-line and the Raman amplifier chamber

The MTW laser is being used a driver for OPAL, a chain of optical parametric chirped pulse amplifiers aiming to produce a 15 fs pulse with energy 5 J. The OPAL system will use 2 ns long beams at 527 nm to pump the optical parametric amplifiers. These stages overlap an intense pump beam with the main pulse in a crystal with a strong non-linear refractive index. This has the effect of amplifying the main beam, with the excess energy released as a third beam known as the idler. The last stage will use the MTW beam to supply approximately 100 J, expected to amplify the main pulse to several joules with an ultra broad bandwidth. Subsequently, the chirped pulse must be compressed using a conventional grating compressor.

Optical parametric amplifiers have the advantage of using a well understood non-linear optics mechanism, however the main disadvantage is that the amplified pulse still requires compression, with the associated large apparatus. As a by product from the interaction, the OPAL system will produce a broad bandwidth idler pulse across the near infra-red spectral region.

Since there is a well known relationship between the main pulse and the idler, it should have a well known spectral phase. This allows it to be compressed and used as a tune-able, high-power seed pulse for the Raman amplifier. Its energy will be approximately 50 mJ and will be compressed to 100 fs, giving a power of 0.5 TW.

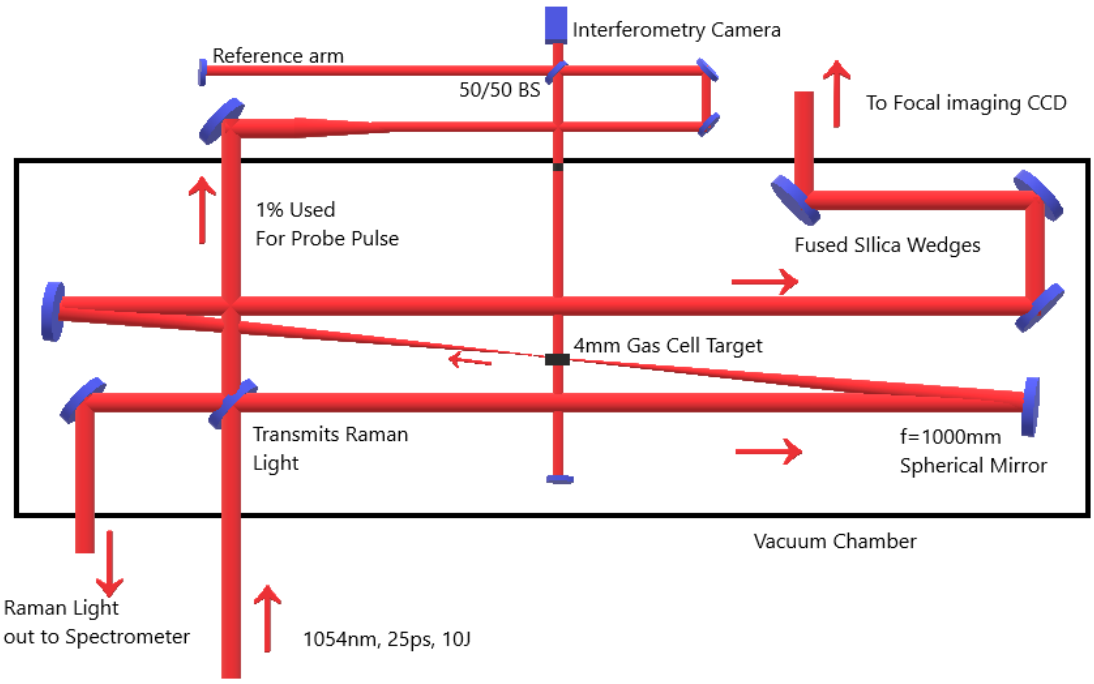


Figure 8.2: Schematic of the target chamber, showing the focussing optic, the gas cell target, interferometry probe optics and the collection optics for the main pulse. The back-scattered Raman light transmits through the four-way mirror and exits the chamber to a spectrometer.

This is comparable to the pump pulse power.

The target chamber is shown in Fig. 8.1 and the schematic in Fig. 8.2. The interferometry beam splitter (visible in Fig. 8.1) was later moved outside of the chamber to give additional alignment control. The gas cell target had variable length and the entry holes had variable width. The target position was controlled under vacuum through several actuators.

A small fraction of the main laser pulse exited the chamber to a variable delay line to be used in the Michelson interferometer. This diagnostic allows accurate determination of the plasma density, by passing a low power probe pulse transversely through two windows in the side of the gas cell. A thin pellicle beam splitter was used to minimise aberrations to the image of the target. Light propagating backwards from the target is collected and re-collimated. Any Raman light transmits through the main dichroic mirror and is coupled out of the vacuum chamber to the spectrometer.

During the commissioning of this new target chamber, it was possible to investigate the spontaneous Raman backscatter from the pump pulse only. Unfortunately,

the high-power seed pulse will not be available until 2019.

The pump pulse, with energy up to 10 J, was focussed using a spherical mirror of focal length 1000 mm and width 50 mm. This long focal length was chosen to maximise the focal beam waist. The focal spot quality was improved using a spatial filter aperture before the focus, outside of its Rayleigh range.

8.4 Measurement of the Raman scatter

The pump pulse energy was slowly increased to 8 J, using only the rod amplifiers. The energy arriving at target was calibrated using a pulse calorimeter connected to an oscilloscope. The spatial filter improved the transverse beam profile at the focus, but meant that some of the laser energy was scattered. The focal profile had an elliptical shape, with a FWHM of $63 \times 30\,\mu$m. Use of the spatial filter reduced the energy reaching the target by 60 %

The beam was fully transmitted through the holes of the gas cell, without ionising its surfaces. Initially, the backing gas pressure was kept constant and the plasma density was varied by delaying the arrival of the laser pulse. As the gas cell fills, the gas density increases, so that shots with a greater delay had a higher gas density.

The light backscattered from the plasma was collected as it partially transmitted through a large mirror within the chamber. It was down-collimated through a telescope within the vacuum chamber, before passing out of the chamber window and coupled with an optical fibre to an Ocean Optics NIRQuest-512 modular spectrometer. The detector was cooled to reduce the noise amplitude. The spectral range was $1030 - 1208$ nm and the spectral resolution was 0.5 nm, small enough to resolve the linear bandwidth of the Raman backscatter.

A small fraction of the incoming laser light transmitted through an optic in the vacuum chamber. This light was collected, down-collimated to a width of 2 cm, frequency doubled in a non-linear crystal, then split with a beam splitter in a Michelson interferometer configuration. One arm contained the target. This yielded an interfered image of the plasma, similar to that shown in the previous chapter.

Due to the elliptical beam profile within the Rayleigh range, the assumption of an axially symmetric plasma was not valid and so an Abel inversion was not possible. To extract the plasma density, I made an estimate of the plasma channel width L from the height of the channel in the interferogram. This will have a significant

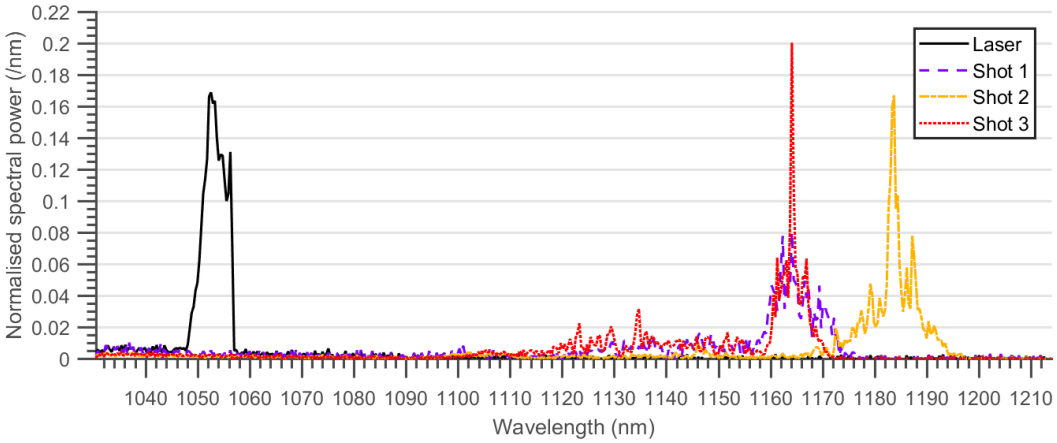


Figure 8.3: Spectrum of the scattered light from the target, in the backwards direction. In each case, the pump pulse had energy 1.4 J and intensity 8×10^{14} Wcm^{-2} or $a_0 = 0.025$. The plasma density was measured from the interferometry diagnostic as 1.2×10^{19} cm$^{-3} \pm 20\%$ for shots 1 and 3, and 1.5×10^{19}cm$^{-3} \pm 20\%$ for shot 2. The spectra have been relatively calibrated using the detector response function and detector noise has been subtracted. The spectra are not absolutely calibrated and they have been normalised. The black line shows the reference spectrum of the laser, which was removed using a spectral filter for all other shots.

error due to the ellipticity of the beam profile. The error was estimated from the relative ellipticity of the beam shape. The plasma density was then extracted using $n = n_{crit}\phi\lambda/L\pi$.

The plasma density extracted from the interferometry data was approximately 1.2×10^{19} cm$^{-3} \pm 20\%$ or $0.012 n_{crit}$. This shows the gas cell is capable of delivering a plasma density within the optimal range found in Chapter 4. The length of the gas cell was set to 7 mm.

The spectrum of the collected backscattered light for several different shots is shown in Fig. 8.3. In all shots, the laser pulse at the target had energy 1.2 J$\pm$5%, duration 25 ps and intensity 8.2×10^{14}Wcm$^{-2} \pm 20\%$.

The reflected laser light at 1054 nm was removed using a spectral filter. There is a clear peak due to Raman light around 1165 nm and 1184 nm for the higher density shot.

The spectrometer spectral calibration was completed by the manufacturer and supplied with the data collection software driver. Since the pinhole, grating, grating position, spectral range and detector were fixed, this calibration should remain accurate for the current purposes.

The spectrometer detector was calibrated in energy using a broadband black-

body emission source. We found the detector calibration function by comparing the measured spectrum to the black-body spectrum. The spectra in Fig. 8.3 have been corrected using this calibration curve. Due to the small slit size to attain high resolution, the spectrometer did not measure the full aperture of the Raman backscatter and so absolute energy calibration of the spectrometer is not accurate, since the angular distribution may vary from shot to shot.

The scattered spectrum appears to have some sub-structure. For all shots, there is a very sharp Raman scatter feature. This is the linear Raman scattering peak. The position of this peak increased with the plasma density, as expected.

There is some discrepancy with the expected positions of the peaks. For the shots shown in red and purple, the measured plasma density was $1.2 \times 10^{19}\text{cm}^{-3} \pm 20\%$, giving an estimated Raman wavelength of 1181 nm. The measured wavelength is beyond the lower end of the expected range.

The wavelength is similarly over-estimated for the higher density shot 2 in Fig. 8.3. The measured plasma density from interferometry was $1.5 \times 10^{19}\text{cm}^{-3} \pm 20\%$, giving an expected Raman peak at 1198 nm. Again, the measured peak is at the lower end of the expected range, suggesting there may be a systematic error in the interferometry analysis. This may be because the laser pulse profile within the Rayleigh range had an elliptical shape. This means the plasma could have been larger in the horizontal direction than the vertical direction. Since the analysis of the interferometry data assumed axial symmetry, this is likely to over estimate the true plasma density.

For each shot, the width of the Raman peak is consistent with the expected bandwidth of the Raman light, $\gamma_{RBS}/\omega_0 = 0.0015$. This leads to an expected Raman peak bandwidth of 1.5 nm, around three times the spectral resolution. The narrow Raman peak shows that the plasma density remains uniform across the full volume, since regions of higher or lower plasma density would scatter at different wavelengths.

In the measured spectra, the Raman features also had a lower amplitude, wider bandwidth pedestal. This could be due to the chirp of the pump pulse which will stimulate Raman scatter at different wavelengths at different times. The width of this pedestal is similar to the measured pump pulse bandwidth of 8 nm.

In the two lower density shots, there is also a broad spectral feature from $1120 - 1160$ nm. This is blue shifted compared to the main linear Raman peak. This feature

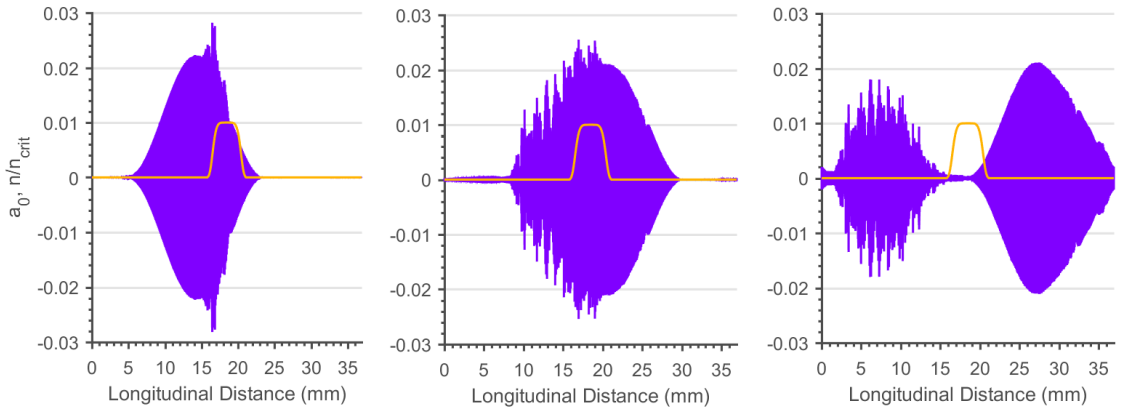


Figure 8.4: One-dimensional particle-in-cell simulation replicating the experiment. The dimensionless amplitude of the electric field is shown in purple, with the initial plasma density shown in yellow. The plasma density is normalised to the laser critical density. The panels show three separate time-steps, each separated by 20 ps. The laser propagates towards the right.

became relatively more prominent for a further shot at higher intensity, suggesting it increases for high pump intensity and/or plasma temperature. The broader feature could also be due to Raman scattering from lower density regions of the plasma.

8.5 Particle-in-cell simulation of the full plasma

To investigate the spectral features and the effect of the chirp in the pump pulse, the experiment was simulated using the particle-in-cell code Osiris. The one-dimensional fixed simulation domain covered the entire plasma, with a cell size of 17 nm containing 32 electrons and 32 hydrogen ions per cell.

The density profile had a super-Gaussian shape with a 1 mm size ramp either side of a 4 mm length plateau. The peak plasma density was set at $10^{19}\,\mathrm{cm}^{-3}$, with a constant temperature of 100 eV. In the experiment, the temperature increased as a function of time and space due to collisional damping, so the constant simulated value of 100 eV is an approximation.

A pump pulse of wavelength 1.054 μm, intensity $5 \times 10^{14}\,\mathrm{Wcm}^{-2}$ and dimensionless amplitude 0.02 was incident on the plasma. It had linear polarisation and the temporal profile was $a_0^2 \cos^2(\pi t/2T_1)$. The bandwidth of the pulse was 8 nm, produced with a linear chirp.

The evolution of the pulse across three time-steps is shown in Fig. 8.4. The pulse is predominantly transmitted through the plasma, although there is a large

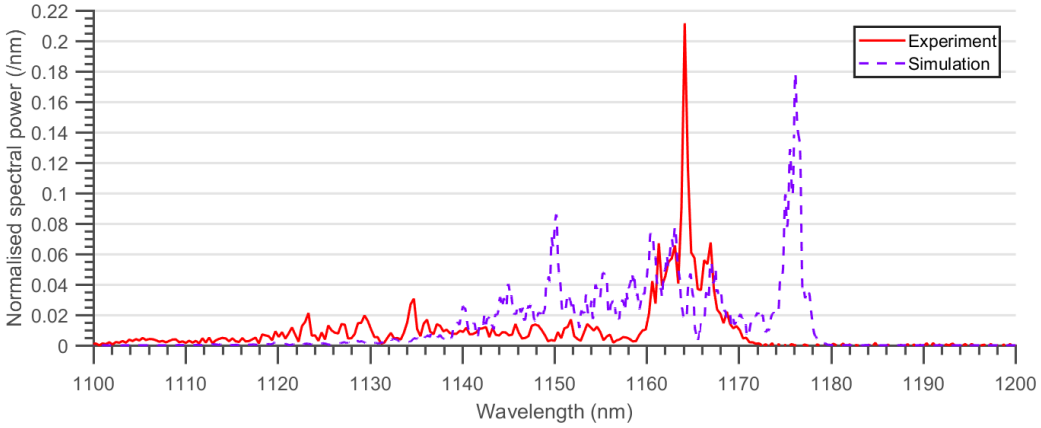


Figure 8.5: Simulated Raman backscatter spectrum, compared to the shot shown in red in Fig. 8.3. The electric field of the backscattered light from the simulation was analysed using the square modulus of its discrete Fourier transform.

amount of spontaneous backscatter. The spectrum of this backscatter is shown in Fig. 8.5 and compared to the experimental spectrum. It exhibits many of the same characteristics, such as the narrow linear Raman peak and the broad blue-shifted feature.

The narrow peak has a slightly longer wavelength than in the experimental spectrum, suggesting the simulated plasma density was slightly higher than in the experiment. However, the width is similar, in agreement with the experiment.

The broad, lower amplitude feature is also observed in the simulation [32]. To ascertain whether it comes from the bulk of the plasma or its low density edges, the simulation data was analysed using a short time Fourier transform, given by

$$F(\tau,\omega) = \left| \int e^{-i\omega t} e^{-(t-\tau)^2/2\sigma^2} E(t)\, dt \right|^2. \tag{8.1}$$

This gives a spectrogram similar to the discussion of the FROG signal in Chapter 3. The window width σ was taken as $2\lambda/c$. The spectrogram shows the time resolved spectrum of the signal. Fig. 8.6 shows the spectrogram of the backscattered electric field.

The backscattered light emanates mainly from the bulk of the plasma and not from the lower density edges. This is true both for the narrow, high amplitude peak and the broader feature.

This was further confirmed by repeating the simulation with a flat-top density profile with no lower density regions. The simulated spectrum was very similar, suggesting all spectral features come from the uniform density central section of the plasma.

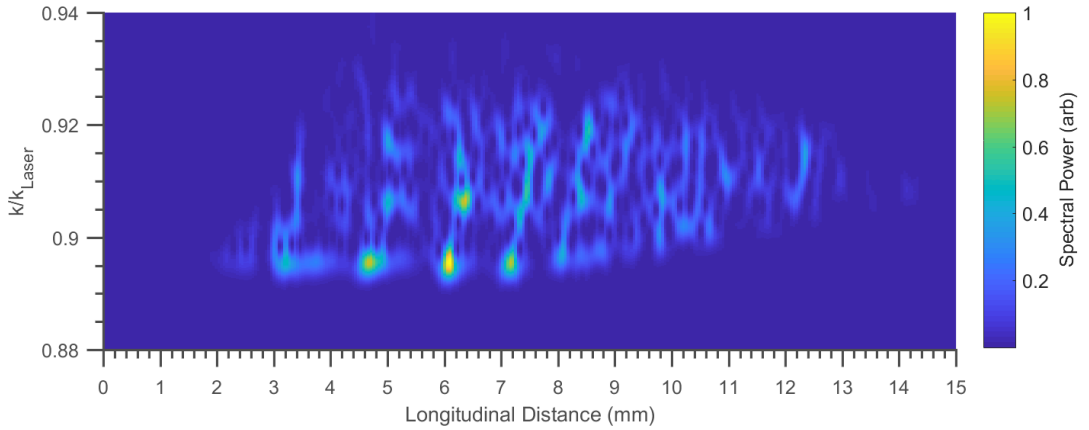


Figure 8.6: Windowed Fourier transform of the simulated backscatter electric field shown in Fig. 8.4c. This was acquired using the windowed Fourier transform (8.1). The laser pulse had fully propagated through the plasma. The backscatter is propagating towards the left. The spectral axis is normalised to the central frequency of the laser.

The spectrogram shows that many individual wave-packets grow through Raman scattering, then their growth saturates after several picoseconds. This could be because the chirp of the pump pulse de-tunes the resonance, or because the plasma wave amplitude grows to the point where it breaks and the growth of the wave-packet ceases. The plasma relaxes over several picoseconds, allowing a subsequent wave-packet to grow. This process repeats, meaning the spontaneous Raman backscatter has a picosecond scale temporal and spectral structure.

This behaviour is seen in the spectrogram, where the linear Raman peak at $k/k_{Laser} = 0.897$ is composed of several distinct pulses, each with a duration of approximately 1 ps and a separation of 4 ps. The linear Raman growth time $\gamma_{RBS}^{-1} \simeq 160$ fs, meaning 4 ps is sufficient time for each wave-packet to grow from noise.

The standard linear theory predicts a narrow bandwidth peak and fully coherent scattered light. However, the simulated backscatter is incoherent and composed of many individual wave-packets.

These findings are of significance to the inertial confinement fusion programme, since the peak intensity of the Raman scatter is much greater than the average intensity. Due to the picosecond scale temporal structure, a measurement of low reflectance from the plasma does not imply the plasma wave amplitude was low at all times. In fact, there could have been a cycle of a growing Raman wave-packet reaching the wave-breaking amplitude, followed by a swift decay [113]. This

duration of this cycle is typically several picoseconds, meaning there will be many such cycles within a typical inertial confinement fusion implosion.

Importantly, the plasma wave regularly reached high amplitude, despite a low overall reflectance of 14%. Since the decay of plasma waves through wave-breaking produces supra-thermal electrons, there may be greater numbers of these electrons than previously thought [22]. Furthermore, their angular distribution will be irregular and could pre-heat the target with a non-uniform pattern.

Since the plasma waves have been simulated to reach wave-breaking amplitude in just several picoseconds, methods to de-tune Raman scattering via phase modulation must do so within a picosecond time-frame.

This new experiment at the University of Rochester aims to study many aspects of Raman amplification and set a new record for the energy transfer efficiency. The design allows use of the seed ionisation scheme and a diverging seed pulse as described in Chapter 5. The initial results show that the plasma density is uniform and within the optimal range from the simulations. Furthermore, the spontaneous backscatter is composed of many incoherent wave-packets that grow and then become damped or de-tuned. This gives the Raman backscatter a picosecond scale temporal and spectral structure.

This experiment was devised, constructed and performed by the group of Dustin Froula at LLE. I attended and supported the experiment for 10 weeks, during the construction and data collection. The analysis and simulation were performed by myself and are unpublished.

Chapter 9

Application to X-ray Pulses

Sub-femtosecond X-ray pulses are opening a new branch of physics, aiming to understand ultra-fast events such as the dynamics of bound electrons and chemical transitions. The characteristic timescale of a bound electron in a 1 eV potential is $t = \hbar/E = 0.6\,\text{fs}$. Events shorter than this are required to temporally resolve the dynamics. However, sub-femtosecond events are shorter even than the time period of visible light, so short X-ray pulses are the only physical probe able to resolve events such as chemical reactions. The time period of a 100 eV photon is $40 \times 10^{-18}\,\text{s}$ or 40 as. X-ray or XUV pulses of only several cycles are therefore required.

9.1 Scaling to shorter wavelengths

Fortunately, the derivation given in Chapter 2 is independent of the pump and seed wavelengths [106]. So long as the ratio n/n_{crit} and the product $I\lambda^2$ are kept constant, the analytical results still hold.

This has important implications for compression of pulses in the ultra-violet or X-ray spectrum [71, 111, 86, 67, 26]. This is currently infeasible through standard methods because the damage threshold of optics decreases sharply beyond the visible part of the spectrum.

X-ray free electron lasers have led to a dramatic increase in the brightness of X-ray pulses in the last decade. For example, the Linac Coherent Light Source [28] and the European XFEL [42] achieve a brightness 6 orders of magnitude greater than previous synchrotron sources. In addition, the standard pulse duration is as short as 200 fs and the spectrum is monochromatic and tune-able.

A low emittance electron bunch is accelerated to several GeV by a linear accelerator. The bunch is then passed through a spatially oscillating transverse magnetic

field known as an undulator. The transverse acceleration of the bunch emits electromagnetic radiation with the wavelength of the undulator. The Lorentz transform to the laboratory frame gives a blue shift of many thousands, yielding an intense X-ray pulse of similar duration to the electron bunch. The electrons are then diverted and dumped.

Innovative schemes with short electron bunches have shortened the available pulse durations to just a few femtoseconds [17]. However, an order of magnitude shorter may be possible using the Raman amplifier mechanism.

Several other proposals have the potential for attosecond pulse generation, but most compromise on power and brightness. Several studies have highlighted the need for increased fluence in sub-femtosecond pulses [46].

A free electron laser offers the perfect conditions for stimulated X-ray Raman scattering, since the achievable dimensionless amplitude is of the order 0.01 after focussing. The collective emission of the electron bunches leads to transversely coherent pulses, allowing focussing using finely machined elliptical optics. Intensities of $6 \times 10^{17}\,\mathrm{Wcm}^{-2}$ have already been demonstrated [140]. Although the emission is transversely coherent, it is not temporally coherent because the electrons form independent micro-bunches.

Stimulated Raman scattering would also require a high-power X-ray seed. One option is a second free electron laser source, but a more practical option could be a high harmonic emission from a gaseous or solid target. Femtosecond XUV harmonics are routinely generated at facilities such as Artemis at the Central Laser Facility [125].

A schematic is shown in Fig. 9.1. The high harmonic source impinges on a thin plastic film, fully ionised and expanded by a pre-pulse laser to below solid density. The XFEL pump pulse is close to counter-propagating and focussed to peak intensity, as in reference [140]. The Rayleigh range will greatly exceed the length of the plasma. For a typical 200 fs XFEL pulse, the plasma thickness should be $40\,\mu$m. To allow full ionisation and reduce collisional damping, low Z elements and plastic targets will be preferable.

Alignment of the pulses should be possible due to the broad angular spread of the high harmonic radiation. The high harmonic radiation also has a broad spectrum, so the matching conditions should be automatically satisfied. In fact, the X-ray free electron laser pump pulse can also be finely tuned to maximise the

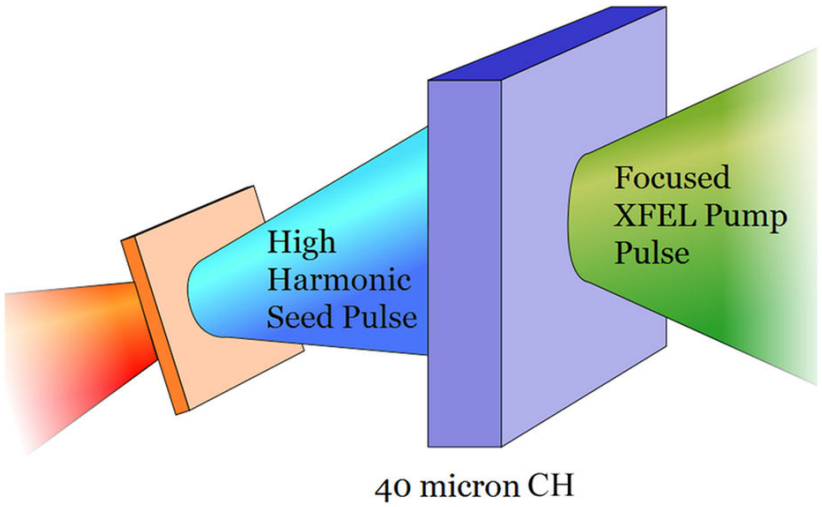


Figure 9.1: Schematic of a soft X-ray Raman amplifier. The required plasma density is slightly less than solid density. The free electron laser pulse can be focussed and used as a pump. It will amplify seed light from a high harmonic source. Since there are a range of harmonics and the free electron laser is tune-able, plasma wave resonance is easily achieved. The target may be expanded and heated by a pre-pulse to achieve the required density and temperature. The target should be close to fully ionised and consist of low Z ions.

resonance. In addition, the high harmonic radiation arrives as a train of pulses lasting for as long as the optical pulse that drives it. This alleviates any issues with timing consistency.

Although the self-similar solution in Chapter 2 is invariant with wavelength, the assumptions in its derivation are not. For example, collisional damping could lead to severe energy loss and greater temperatures. Furthermore, greater temperatures could increase Landau damping of the plasma wave and reduce efficiency. Fig. 9.2 shows the parameter space plot for pump wavelengths of 100 nm and 10 nm. Compared the the case with infra-red light, the high efficiency region narrows and is pushed towards higher plasma densities and temperatures. The dual effects of collisional damping at low temperature and Landau damping at high temperature act to close the viable parameter space for these higher densities.

As the permitted region narrows, there is a minimum wavelength, below which backward Raman scattering is infeasible [71]. This appears to be around 1 nm wavelength or 1 keV photon energy. Even for a 10 nm wavelength, the plasma temperature will be high, leading to a greater loss in efficiency [2].

To model this type of interaction in a highly collisional plasma, the particle-in-cell code must be adapted to include collisional damping. This is most easily done

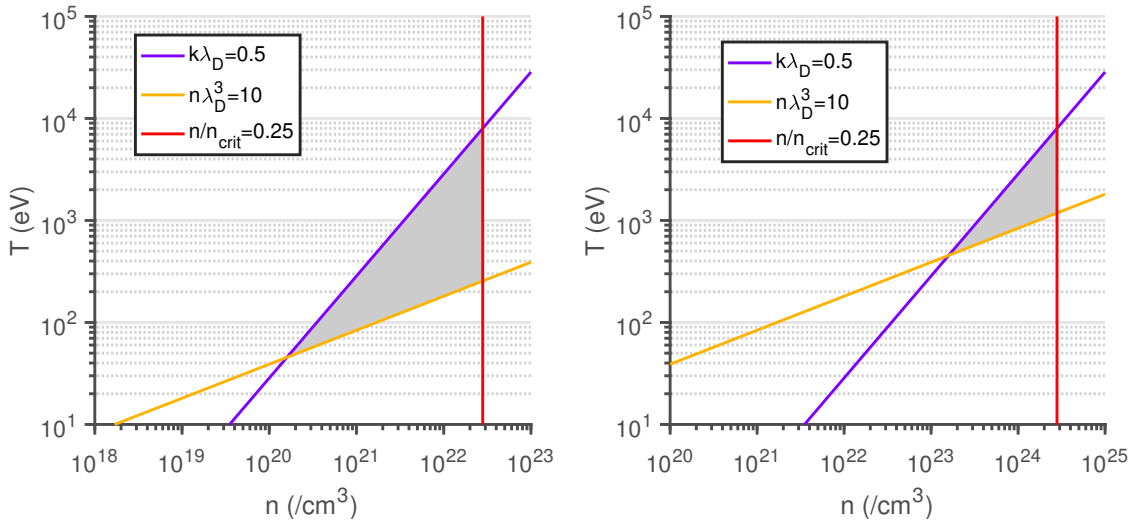


Figure 9.2: Parameter space for a pump wavelength of (a) 100 nm and (b) 10 nm. As the wavelength shortens, the plasma becomes denser and more collisional, closing the feasible parameter window, shaded in black. The lines show the limitations due to collisional damping (orange), Landau damping (purple) and the requirement $n/n_{crit} < 1/4$ for Raman scattering to occur (red).

through an extra routine that models the binary collisions between particles within the same or adjacent cells. Particles are paired randomly and their velocities are rotated using a calculation from the collision cross section. This method allows local collisional behaviour without a need for the grid to resolve the fine scale fields between particles.

Collisions will tend to increase the entropy of the plasma and the newly acquired random motion will reduce the transverse current in the laser pulse. This means it will become damped. To model this kind of interaction, the simulation must include ion and electron particles. Furthermore, a large number of particles are required to effectively model the velocity space within each laser wavelength.

The code Osiris models electron-electron, ion-ion and electron-ion collisions. The Debye length in equation (1.24) is approximated locally by measuring the temperature of the local particle distribution. The grid is split into a large number of collision cells, but due to computational restraints, this is a number less than the total number of grid cells.

Although Osiris does model field ionisation, it does not include detailed atomic physics and so absorption due to atomic transitions is neglected. The shaded region in Fig. 9.2 shows where $n\lambda_D^3 > 10$ and so collisions should be minimal on a timescale ω_p^{-1}. The plasma should be highly ionised and so losses to atomic absorption are

minimal. Attaining conditions shown in the grey area in Fig. 9.2 will require initial expansion and heating of the plasma by an external source. For a 10 nm wavelength, the required electron density is approximately solid density.

Since the collisional damping will occur over the full propagation of the pump pulse across the plasma, it is important to model the full volume of the plasma. Due to current computational restraints, this will require a one-dimensional simulation. The increase of the plasma temperature over time will increase the amount of Landau damping. In addition, the plasma wave will become de-tuned via the temperature dependent dispersion relation found in Chapter 1. The combination of increased Landau damping, increased collisionality and plasma wave de-tuning will act to lower the efficiency.

With the inclusion of this collision routine, the particle-in-cell code should effectively model conditions where $n\lambda_D^3 \simeq 10$. All kinetic effects will be automatically included as before, since the particle method models the full velocity space.

The simulated pump pulse wavelength was 10 nm (photon energy 124 eV). A wavelength this short is needed to produce sub-femtosecond pulses, but any shorter and efficiency losses become prohibitive. The length of the static simulation domain was 40 μm. The uniform plasma was initialised across a central 35 μm region, with constant density $2.2 \times 10^{23}\,\mathrm{cm}^{-3}$ or $n = 0.02 n_{crit}$, and a 400 eV initial temperature. The temperature was initialised high to allow $n\lambda_D^3 > 1$ from the start, a requirement for the collision model from Chapter 1 to be effective. Since the plasma is expanded by an optical pre-heating pulse, this is a realistic starting temperature. The cell size was 0.16 nm and there were 256 electrons and 128 Hydrogen ions per cell.

The seed was initialised with amplitude $a_1 = 0.02$ ($I_1 = 6 \times 10^{18}\mathrm{Wcm}^{-2}$), wavelength 12 nm and duration 5 fs. The pump pulse had a constant intensity of the same value, wavelength 10 nm and duration 240 fs. The polarisations were linear and aligned. Both pulses were monochromatic and fully coherent.

9.2 Simulation results

In the simulation, the intense pump pulse traverses the plasma and heats it through collisional damping. Fig. 9.3 shows the intensity over several time-steps. Panels (a,b) show the depletion of the pump pulse due to collisional damping as it traverses the plasma. For the initial temperature and density, the Debye number is $n\lambda_D^3 = 7$, giving the electron ion collision time as 2.5 fs and equation (1.96) gives the collisional

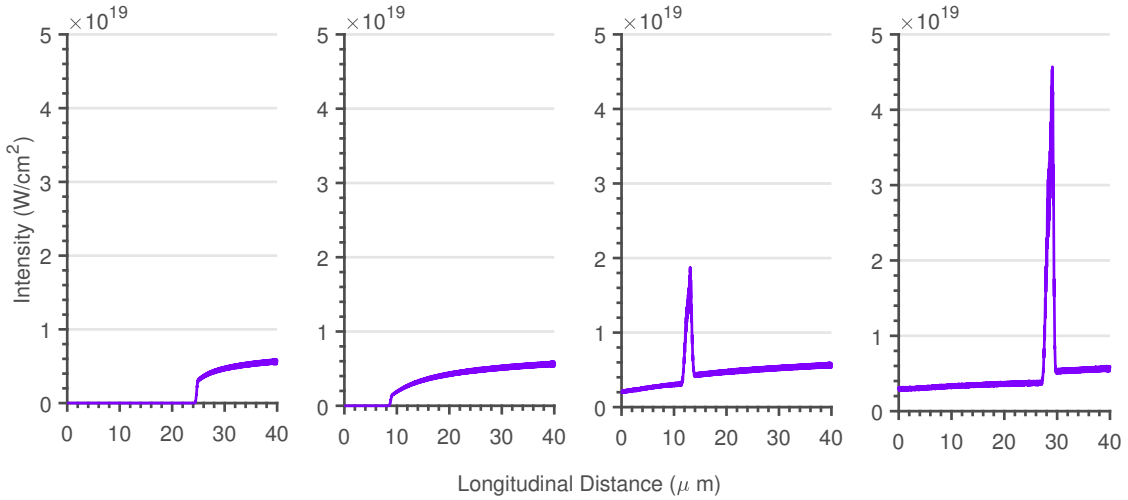


Figure 9.3: One-dimensional PIC simulation for an XUV pump of wavelength 10 nm. The panels show the pulse intensity across four time-steps, each separated by 53 fs. The collisional damping heats the thin target and depletes the pump, which travels to the left. The 5 fs long seed pulse, wavelength 12 nm, is initialised once the pump has crossed the plasma and it propagates to the right. Both pulses started with intensity $5.6 \times 10^{18}\,\mathrm{Wcm}^{-2}$.

damping distance $c/\gamma_{coll} = 37\,\mu$m, similar to the target thickness. This agrees with the depletion observed in the PIC simulation. As the target heats, it becomes less collisional and so only a small fraction of the pump energy is damped.

The electron temperature rises from 400 eV to 1300 eV after 100 fs. Panels (c,d) show the propagating seed pulse, which started at the same intensity as the pump. The amplification is effective and the pump pulse is visibly depleted. However, the pump depletion of 30 % is considerably less than that simulated in the optical regime. The amplified pulse reaches a final intensity of $5 \times 10^{19}\,\mathrm{Wcm}^{-2}$, approximately 10 times the pump intensity. The final duration of the pulse is 5 fs and the energy transfer efficiency was 17%, considerably less than the high efficiency found for equivalent scaled parameters in the optical regime. This is despite the seed pulse entering the non-linear stage. The low efficiency is due to collisional damping of the pump pulse and the associated temperature increase, leading to excessive damping of the wave due to particle trapping. Magnetising the plasma may alleviate these problems [111]. The pulse is not shortened by the non-linear stage and remains at approximately 5 fs.

Although the optimal plasma density from Chapter 4 was $n = 0.02 n_{crit}$, Fig. 9.2 shows that the optimal region is pushed towards higher density in the x-ray spectrum. The simulation was therefore repeated with a free electron density

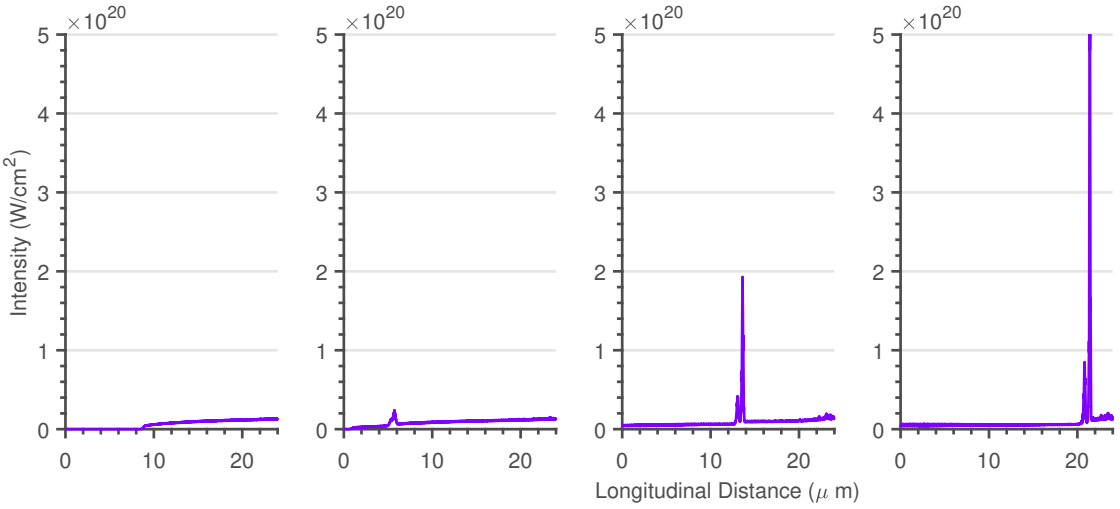


Figure 9.4: One-dimensional PIC simulation for an XUV pump of wavelength 10 nm. The panels show the pulse intensity across four time-steps, each separated by 26 fs. The collisional damping heats the thin target and depletes the pump, which travels to the left. The 2.5 fs long seed pulse, wavelength 13 nm, is initialised once the pump has crossed the plasma and it propagates to the right. Both pulses started with intensity $1.3 \times 10^{19}\,\mathrm{Wcm}^{-2}$.

close to solid density, at $0.04 n_{crit}$ or $n = 4.4 \times 10^{23}\,\mathrm{cm}^{-3}$. The plasma thickness was correspondingly reduced to $20\,\mu$m. The pump pulse intensity was $I_0 = 1.3 \times 10^{19}\,\mathrm{W/cm}^{-2}$. The seed pulse had the same intensity as the pump and duration 2.5 fs.

Fig. 9.4 shows the development of the amplified pulse intensity across four time-steps. With this higher density, conditions are closer to the centre of the optimal region in Fig. 9.2. The plasma was initialised with temperature 800 eV, giving only minimal collisional damping. The seed pulse enters the non-linear stage and a short spike grows at the front.

Due to the higher plasma density, the main amplified spike is shorter and reaches a duration of 300 as after the $20\,\mu$m propagation. The final intensity greatly exceeded that of the pump pulse. The pump depletion within the interaction was 55%.

This is a promising result, since a well focussed free electron laser pulse can be effectively compressed. Unfortunately, the energy efficiency is lower than in the optical regime. However, this means the pump pulse remains only partially depleted after the interaction. If the seed radiation consists of a train of pulses, each will be effectively amplified to a similar level. A train of mJ energy, sub-femtosecond pulses could lead to several snapshots of processes occurring on the femtosecond

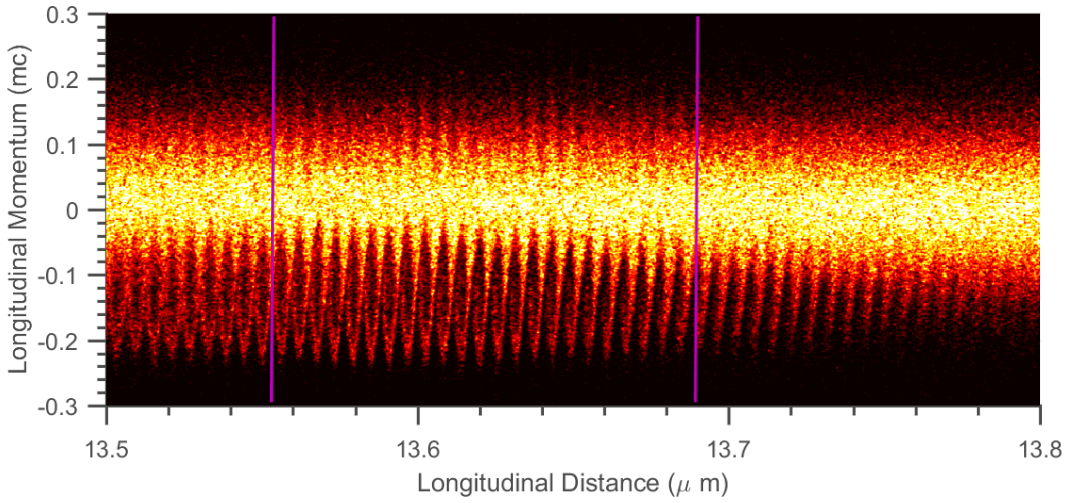


Figure 9.5: Electron phase-space data from the simulation, at the third time-step in Fig. 9.4. The FWHM of the amplified pulse is indicated by the solid purple lines, and it propagates towards the right. The plasma density was $4.4 \times 10^{23}\,\mathrm{cm}^{-3}$ and the pump intensity was $1.3 \times 10^{19}\,\mathrm{Wcm}^{-2}$.

timescale.

There is also evidence that the wave-breaking process helps to shorten the final duration of the X-ray pulse, since the wave breaks within the extent of the seed pulse and stops the energy transfer. Fig. 9.5 shows the longitudinal electron phase space. The plasma has reached a temperature of 1.6 keV via collisional absorption. Its Maxwellian distribution is seen on the right of the image, before the interaction. The plasma wave has a wavelength of 5.6 nm. This high temperature means the plasma wave breaks near the peak of the amplified pulse, although it does retain some coherence for many cycles afterwards. Since the energy transfer at the rear will be reduced by this loss of coherence, the amplified pulse is shortened to less than 1 fs.

Despite the promising results, it may be difficult to reach the X-ray intensity that was simulated, especially for the high harmonic seed. Extending the scheme to shorter wavelengths and shorter duration output will also be challenging due to losses in the dense plasma. The presence of atomic absorption, collisional damping and Landau damping act to render the process infeasible for wavelengths shorter than 10 nm, even in the simulations. This problem is further compounded by the lower dimensionless amplitude of x-ray pulses at shorter wavelengths, since $a_0 \propto \sqrt{I\lambda^2}$. In addition, the energy in high harmonic pulses has a scaling law that decreases sharply with higher harmonics, meaning the technique is only feasible for

the XUV region of the spectrum and not for X-rays.

The work in this Chapter is my original research, although the idea for amplification at X-ray wavelengths was previously put forward in reference [71].

In summary, the one-dimensional particle-in-cell simulations have been extended to shorter wavelengths by modelling collisional damping. This raises the required temperature and density, increasing the losses and decreasing the efficiency. Despite this problem, a free electron laser pump pulse can be compressed to sub-femtosecond duration within a high temperature solid density plasma. Initial experimental studies may be possible using a XUV seed pulse from a high harmonic source. The foil target will also require heating and expanding with use of an optical pre-pulse. Despite these complexities, it may be possible to efficiently compress x-ray free electron laser pulses to just a few laser cycles, and possibly sub-femtosecond duration. Furthermore, the simulated output has unprecedented intensity for any X-ray or XUV beam.

Chapter 10

Conclusions and Outlook

Reaching peak intensities in excess of $10^{24}\,\mathrm{Wcm}^{-2}$ could be possible via amplification of laser pulses in a plasma. For the first time, advances in performance of simulation codes has allowed a large scale parameter scan using a fully kinetic approach. Importantly, plasma amplifiers are also robust to experimental perturbations and work with various beam geometries. The impediment of generating a suitable seed pulse has been addressed by using forward Raman scatter of a short relativistic pulse, also verified experimentally.

10.1 Overall conclusions

- The analytical solution for late times is validated by particle-in-cell simulations. The condition is accurately given by (2.44)
- Below the wave-breaking threshold, the efficiency is close to η_{max}
- Above the wave-breaking threshold, the amplified pulse duration is shorter than the analytical solution and the efficiency decreases below η_{max}. A fit to the simulations gives $\eta \propto a_0^{-1.1}$
- In the near visible spectrum, the optimal efficiency for realistic plasma temperature is approximately 0.6, at an electron density of $0.02 n_{crit}$.
- For finite temperature, the wave-breaking threshold worsens and is well approximated by equation (4.22)
- The non-linear solution is only valid for pulses shorter than γ_{RBS}^{-1}
- For a fixed energy and bandwidth, the shortest seed pulses have the best performance, despite starting further from the non-linear solution

- Filamentation and forward Raman scattering limit the maximum achievable amplitude to far below $a_1 = 1$. Higher densities and lower pump intensity increase the efficiency, but also worsen the limitations of the competing instabilities.
- Growth of competing plasma instabilities limit the maximum compression ratio and the maximum power amplification to 300-1000.
- For a given pump pulse duration, there is a unique set of parameters that maximises the efficiency and compression ratio
- An amplifier was simulated in a focussing geometry and coherence is retained. However, relativistic self focussing may move the position of best focus.
- The analytical solution holds as long as the focal position is far from the extent of the interaction.
- If the amplification starts at the focus, there is a different temporal envelope of the amplified pulse
- By amplifying a diverging pulse and increasing the pump power linearly, the amplified pulse intensity can be kept constant while increasing its power and energy
- Starting the amplifier at the focus minimises the required seed pulse energy
- Variations in the plasma density profile will not have a significant effect on the efficiency or coherence, so long as the non-linear stage has been reached
- Spontaneous Raman backscatter may have picosecond scale temporal and spectral structure, even from uniform plasma.
- Seeding the amplifier could be possible using forward scatter from a laser at the same wavelength as the pump.
- It was experimentally verified that a seed generated in this way has sufficient energy and ultra-short duration.
- Simulations show the technique effectively compresses X-ray free electron laser pulses to attosecond duration

10.2 Outlook

Current experiments have verified the non-linear stage of the Raman scattering by measuring the shortening of the amplified pulse and the broadening of its spectrum.

The next step is to reach high overall efficiency, which means the majority of the amplifier length must have significant pump depletion. This work has showed this is possible by starting with a low energy seed at its focus and using the seed ionisation scheme as the pulse diverges. This will have the added benefit of reducing unwanted pump pulse instabilities.

This work has highlighted the importance of the electron temperature to the delicate interaction region. Future experiments must measure the plasma conditions and consider the detrimental kinetic effects to the plasma wave.

All of the elements described in this work are currently available at experimental facilities. Large scale uniform gaseous targets have already been developed for several applications. Fully coherent, tune-able mJ energy seed pulses can be developed through CPA techniques. Once high overall efficiency is demonstrated, it could be possible to scale up the amplifier using the diverging pulse scheme. This will require temporal shaping of a kJ level pump pulse, something that is routinely achieved at facilities such as the National Ignition Facility. Starting the amplification at focus means that a much lower energy seed is required.

The maximum efficiency can approach 50%, similar to conventional pulse compressors. The maximum compression ratio is limited by competing instabilities, however it can typically reach a value of 1000. The conditions maximising both the efficiency and compression ratio have been identified.

One potentially major issue is the transverse beam quality of the pulses outside of the Rayleigh range. Any modulations could quickly seed plasma instabilities. However, the diverging beam scheme permits very low energy seed pulses, so this may not be an issue if corrective optics are used.

The output power of such an amplifier is not limited by the damage thresholds of solid state optics. This has the potential to revolutionise the size and cost of high power laser systems and push the boundaries of high intensity experiments. To achieve these goals, the design must avoid competing instabilities, collisional damping, plasma wave-breaking and de-tuning of the interaction. This is now known to be possible through sufficient characterisation and control of the Raman instability.

Bibliography

[1] TM Antonsen Jr and P Mora. "Self-focusing and Raman scattering of laser pulses in tenuous plasmas". In: *Physical review letters* 69.15 (1992), p. 2204.

[2] AA Balakin, NJ Fisch, GM Fraiman, VM Malkin, and Z Toroker. "Numerical modeling of quasitransient backward Raman amplification of laser pulses in moderately undercritical plasmas with multicharged ions". In: *Physics of Plasmas* 18.10 (2011), p. 102311.

[3] AA Balakin, GM Fraiman, NJ Fisch, and VM Malkin. "Noise suppression and enhanced focusability in plasma Raman amplifier with multi-frequency pump". In: *Physics of Plasmas* 10.12 (2003), pp. 4856–4864.

[4] AA Balakin, GM Fraiman, NJ Fisch, and S Suckewer. "Backward Raman amplification in a partially ionized gas". In: *Physical Review E* 72.3 (2005), p. 036401.

[5] TG Blackburn, CP Ridgers, John G Kirk, and AR Bell. "Quantum radiation reaction in laser–electron-beam collisions". In: *Physical review letters* 112.1 (2014), p. 015001.

[6] Stepan S Bulanov, Timur Zh Esirkepov, Alexander GR Thomas, James K Koga, and Sergei V Bulanov. "Schwinger limit attainability with extreme power lasers". In: *Physical review letters* 105.22 (2010), p. 220407.

[7] SV Bulanov, T Zh Esirkepov, D Habs, F Pegoraro, and T Tajima. "Relativistic laser-matter interaction and relativistic laboratory astrophysics". In: *The European Physical Journal D* 55.2 (2009), p. 483.

[8] Weifeng Cheng, Yoav Avitzour, Yuan Ping, Szymon Suckewer, Nathaniel J Fisch, Min Sup Hur, and Jonathan S Wurtele. "Reaching the nonlinear regime of Raman amplification of ultrashort laser pulses". In: *Physical review letters* 94.4 (2005), p. 045003.

[9] Daniel S Clark and Nathaniel J Fisch. "Operating regime for a backward Raman laser amplifier in preformed plasma". In: *Physics of Plasmas* 10.8 (2003), pp. 3363–3370.

[10] Daniel S Clark and Nathaniel J Fisch. "Particle-in-cell simulations of Raman laser amplification in preformed plasmas". In: *Physics of Plasmas* 10.12 (2003), pp. 4848–4855.

[11] Daniel S Clark and Nathaniel J Fisch. "Regime for a self-ionizing Raman laser amplifier". In: *Physics of Plasmas* 9.6 (2002), pp. 2772–2780.

[12] Daniel S Clark and Nathaniel J Fisch. “Simulations of Raman laser amplification in ionizing plasmas”. In: *Physics of Plasmas* 10.12 (2003), pp. 4837–4847.

[13] DS Clark and NJ Fisch. “Raman laser amplification in preformed and ionizing plasmas”. In: *Laser and Particle Beams* 23.1 (2005), pp. 101–106.

[14] TP Coffey. “Breaking of large amplitude plasma oscillations”. In: *The Physics of Fluids* 14.7 (1971), pp. 1402–1406.

[15] CD Decker, WB Mori, and T Katsouleas. “Particle-in-cell simulations of Raman forward scattering from short-pulse high-intensity lasers”. In: *Physical Review E* 50.5 (1994), R3338.

[16] S Depierreux, V Yahia, C Goyon, G Loisel, P-E Masson-Laborde, N Borisenko, A Orekhov, O Rosmej, T Rienecker, and C Labaune. “Laser light triggers increased Raman amplification in the regime of nonlinear Landau damping”. In: *Nature communications* 5 (2014).

[17] Y Ding, A Brachmann, F-J Decker, D Dowell, P Emma, J Frisch, S Gilevich, G Hays, Ph Hering, Z Huang, et al. “Measurements and simulations of ultralow emittance and ultrashort electron beams in the linac coherent light source”. In: *Physical review letters* 102.25 (2009), p. 254801.

[18] IY Dodin, GM Fraiman, VM Malkin, and NJ Fisch. “Amplification of short laser pulses by Raman backscattering in capillary plasmas”. In: *Journal of Experimental and Theoretical Physics* 95.4 (2002), pp. 625–638.

[19] T Döppner, DA Callahan, OA Hurricane, DE Hinkel, T Ma, H-S Park, LF Berzak Hopkins, DT Casey, P Celliers, EL Dewald, et al. “Demonstration of high performance in layered deuterium-tritium capsule implosions in uranium hohlraums at the National Ignition Facility”. In: *Physical review letters* 115.5 (2015), p. 055001.

[20] G Doumy, F Quéré, O Gobert, M Perdrix, Ph Martin, P Audebert, JC Gauthier, J-P Geindre, and T Wittmann. “Complete characterization of a plasma mirror for the production of high-contrast ultraintense laser pulses”. In: *Physical Review E* 69.2 (2004), p. 026402.

[21] James F Drake, Predhiman K Kaw, Yee-Chun Lee, G Schmid, Chuan S Liu, and Marshall N Rosenbluth. “Parametric instabilities of electromagnetic waves in plasmas”. In: *The Physics of Fluids* 17.4 (1974), pp. 778–785.

[22] RP Drake, RE Turner, BF Lasinski, KG Estabrook, EM Campbell, CL Wang, DW Phillion, EA Williams, and WL Kruer. “Efficient Raman sidescatter and hot-electron production in laser-plasma interaction experiments”. In: *Physical review letters* 53.18 (1984), p. 1739.

[23] Matthias Dreher, Eiichi Takahashi, Jürgen Meyer-ter-Vehn, and K-J Witte. “Observation of superradiant amplification of ultrashort laser pulses in a plasma”. In: *Physical review letters* 93.9 (2004), p. 095001.

[24] B Dromey, S Kar, M Zepf, and P Foster. “The plasma mirrorfffdfffdfffda subpicosecond optical switch for ultrahigh power lasers”. In: *Review of Scientific Instruments* 75.3 (2004), pp. 645–649.

[25] Matthew R Edwards, Nathaniel J Fisch, and Julia M Mikhailova. "Strongly enhanced stimulated Brillouin backscattering in an electron-positron plasma". In: *Physical review letters* 116.1 (2016), p. 015004.

[26] Matthew R Edwards, Julia M Mikhailova, and Nathaniel J Fisch. "X-Ray Amplification by Stimulated Brillouin Scattering". In: *arXiv preprint arXiv:1705.08599* (2017).

[27] Matthew R Edwards, Zeev Toroker, Julia M Mikhailova, and Nathaniel J Fisch. "The efficiency of Raman amplification in the wavebreaking regime". In: *Physics of Plasmas* 22.7 (2015), p. 074501.

[28] Paul Emma, R Akre, J Arthur, R Bionta, C Bostedt, J Bozek, A Brachmann, P Bucksbaum, Ryan Coffee, F-J Decker, et al. "First lasing and operation of an ångstrom-wavelength free-electron laser". In: *nature photonics* 4.9 (2010), pp. 641–647.

[29] B Ersfeld, J Farmer, G Raj, and DA Jaroszynski. "The role of absorption in Raman amplification in warm plasma". In: *Physics of Plasmas* 17.8 (2010), p. 083301.

[30] B Ersfeld and DA Jaroszynski. "Superradiant linear Raman amplification in plasma using a chirped pump pulse". In: *Physical review letters* 95.16 (2005), p. 165002.

[31] Kent Estabrook and WL Kruer. "Theory and simulation of one-dimensional Raman backward and forward scattering". In: *The Physics of fluids* 26.7 (1983), pp. 1892–1903.

[32] Kent Estabrook, WL Kruer, and MG Haines. "Nonlinear features of stimulated Brillouin and Raman scattering". In: *Physics of Fluids B: Plasma Physics* 1.6 (1989), pp. 1282–1286.

[33] John P Farmer and Alexander Pukhov. "Raman amplification in the coherent wave-breaking regime". In: *Physical Review E* 92.6 (2015), p. 063109.

[34] JP Farmer, B Ersfeld, and DA Jaroszynski. "Raman amplification in plasma: Wavebreaking and heating effects". In: *Physics of Plasmas* 17.11 (2010), p. 113301.

[35] JP Farmer and A Pukhov. "Fast multidimensional model for the simulation of Raman amplification in plasma". In: *Physical Review E* 88.6 (2013), p. 063104.

[36] Jérôme Faure, Yannick Glinec, A Pukhov, S Kiselev, S Gordienko, E Lefebvre, J-P Rousseau, F Burgy, and Victor Malka. "A laser–plasma accelerator producing monoenergetic electron beams". In: *Nature* 431.7008 (2004), pp. 541–544.

[37] NJ Fisch and VM Malkin. "Generation of ultrahigh intensity laser pulses". In: *Physics of Plasmas* 10.5 (2003), pp. 2056–2063.

[38] Ricardo A Fonseca, Luís O Silva, Frank S Tsung, Viktor K Decyk, Wei Lu, Chuang Ren, Warren B Mori, S Deng, S Lee, T Katsouleas, et al. "OSIRIS: a three-dimensional, fully relativistic particle in cell code for modeling plasma based accelerators". In: *International Conference on Computational Science*. Springer. 2002, pp. 342–351.

[39] DW Forslund, JM Kindel, and EL Lindman. "Nonlinear behavior of stimulated Brillouin and Raman scattering in laser-irradiated plasmas". In: *Physical Review Letters* 30.16 (1973), p. 739.

[40] DW Forslund, JM Kindel, and EL Lindman. "Theory of stimulated scattering processes in laser-irradiated plasmas". In: *The Physics of Fluids* 18.8 (1975), pp. 1002–1016.

[41] Gennady M Fraiman, Nikolai A Yampolsky, Vladimir M Malkin, and Nathaniel J Fisch. "Robustness of laser phase fronts in backward Raman amplifiers". In: *Physics of Plasmas* 9.8 (2002), pp. 3617–3624.

[42] Gianluca Geloni, E Saldin, L Samoylova, E Schneidmiller, H Sinn, Th Tschentscher, and M Yurkov. "Coherence properties of the European XFEL". In: *New Journal of Physics* 12.3 (2010), p. 035021.

[43] SH Glenzer, BJ MacGowan, P Michel, NB Meezan, LJ Suter, SN Dixit, JL Kline, GA Kyrala, DK Bradley, DA Callahan, et al. "Symmetric inertial confinement fusion implosions at ultra-high laser energies". In: *Science* 327.5970 (2010), pp. 1228–1231.

[44] S Gordienko, A Pukhov, O Shorokhov, and T Baeva. "Coherent focusing of high harmonics: A new way towards the extreme intensities". In: *Physical review letters* 94.10 (2005), p. 103903.

[45] DF Gordon, B Hafizi, P Sprangle, RF Hubbard, JR Penano, and WB Mori. "Seeding of the forward Raman instability by ionization fronts and Raman backscatter". In: *Physical Review E* 64.4 (2001), p. 046404.

[46] M Hentschel, R Kienberger, Ch Spielmann, Georg A Reider, N Milosevic, Thomas Brabec, Paul Corkum, Ulrich Heinzmann, Markus Drescher, and Ferenc Krausz. "Attosecond metrology". In: *Nature* 414.6863 (2001), pp. 509–513.

[47] C Hernandez-Gomez, SP Blake, O Chekhlov, RJ Clarke, AM Dunne, M Galimberti, S Hancock, R Heathcote, P Holligan, A Lyachev, et al. "The vulcan 10 pw project". In: *Journal of Physics: Conference Series*. Vol. 244. 3. IOP Publishing. 2010, p. 032006.

[48] DJ Hoarty, E Hill, P Beiersdorfer, P Allan, CRD Brown, MP Hill, LMR Hobbs, SF James, J Morton, N Sircombe, et al. "Measurements of plasma spectra from hot dense elements and mixtures at conditions relevant to the solar radiative zone". In: *AIP Conference Proceedings*. Vol. 1811. 1. AIP Publishing. 2017, p. 050001.

[49] Min Sup Hur, RR Lindberg, AE Charman, JS Wurtele, and H Suk. "Electron kinetic effects on Raman backscatter in plasmas". In: *Physical review letters* 95.11 (2005), p. 115003.

[50] C Joshi, T Tajima, JM Dawson, HA Baldis, and NA Ebrahim. "Forward Raman instability and electron acceleration". In: *Physical Review Letters* 47.18 (1981), p. 1285.

[51] S Yu Kalmykov and Gennady Shvets. "Stimulated Raman backscattering of laser radiation in deep plasma channels". In: *Physics of plasmas* 11.10 (2004), pp. 4686–4694.

[52] Allan N Kaufman and Bruce I Cohen. "Nonlinear interaction of electromagnetic waves in a plasma density gradient". In: *Physical Review Letters* 30.26 (1973), p. 1306.

[53] Jincheol Kim, Hae June Lee, Hyyong Suk, and In Soo Ko. "Characteristics of pulse compression in laser pulse amplification by stimulated Raman backscattering". In: *Physics Letters A* 314.5 (2003), pp. 464–471.

[54] Ben King, Antonino Di Piazza, and Christoph H Keitel. "A matterless double slit". In: *Nature Photonics* 4.2 (2010), pp. 92–94.

[55] Ben King, Antonino Di Piazza, and Christoph H Keitel. "Double-slit vacuum polarization effects in ultraintense laser fields". In: *Physical Review A* 82.3 (2010), p. 032114.

[56] RK Kirkwood, E Dewald, C Niemann, N Meezan, SC Wilks, DW Price, OL Landen, J Wurtele, AE Charman, R Lindberg, et al. "Amplification of an ultrashort pulse laser by stimulated Raman scattering of a 1 ns pulse in a low density plasma". In: *Physics of Plasmas* 14.11 (2007), p. 113109.

[57] William L Kruer. "The physics of laser plasma interactions". In: (1988).

[58] WL Kruer, Kent Estabrook, BF Lasinski, and AB Langdon. "Raman backscatter in high temperature, inhomogeneous plasmas". In: *The Physics of Fluids* 23.7 (1980), pp. 1326–1329.

[59] J Kruger and M Bovyn. "Relativistic motion of a charged particle in a plane electromagnetic wave with arbitrary amplitude". In: *Journal of Physics A: Mathematical and General* 9.11 (1976), p. 1841.

[60] Lev Davidovich Landau. "On the vibrations of the electronic plasma". In: *Zh. Eksp. Teor. Fiz.* 10 (1946), p. 25.

[61] G Lehmann, KH Spatschek, and G Sewell. "Pulse shaping during Raman-seed amplification for short laser pulses". In: *Physical Review E* 87.6 (2013), p. 063107.

[62] RR Lindberg, AE Charman, and JS Wurtele. "Self-consistent Langmuir waves in resonantly driven thermal plasmas". In: *Physics of Plasmas* 14.12 (2007), p. 122103.

[63] John Lindl, Otto Landen, John Edwards, Ed Moses, and NIC Team. "Review of the national ignition campaign 2009-2012". In: *Physics of Plasmas* 21.2 (2014), p. 020501.

[64] CS Liu, Marshall N Rosenbluth, and Roscoe B White. "Raman and Brillouin scattering of electromagnetic waves in inhomogeneous plasmas". In: *The Physics of Fluids* 17.6 (1974), pp. 1211–1219.

[65] E Lundström, Gert Brodin, Joakim Lundin, Mattias Marklund, R Bingham, J Collier, JT Mendonça, and P Norreys. "Using high-power lasers for detection of elastic photon-photon scattering". In: *Physical review letters* 96.8 (2006), p. 083602.

[66] P Maine, D Strickland, P Bado, M Pessot, and G Mourou. "Generation of ultrahigh peak power pulses by chirped pulse amplification". In: *IEEE Journal of Quantum electronics* 24.2 (1988), pp. 398–403.

[67] VM Malkin and NJ Fisch. In: *Physical Review E* 80.4 (2009), p. 046409.

[68] VM Malkin and NJ Fisch. "Backward Raman amplification of ionizing laser pulses". In: *Physics of Plasmas* 8.10 (2001), pp. 4698–4699.

[69] VM Malkin and NJ Fisch. "Key plasma parameters for resonant backward Raman amplification in plasma". In: *Eur. Phys. J. Special Topics* 223 (2014), p. 1157.

[70] VM Malkin and NJ Fisch. "Manipulating ultraintense laser pulses in plasmas". In: *Physics of plasmas* 12.4 (2005), p. 044507.

[71] VM Malkin, NJ Fisch, and JS Wurtele. "Compression of powerful x-ray pulses to attosecond durations by stimulated Raman backscattering in plasmas". In: *Physical Review E* 75.2 (2007), p. 026404.

[72] VM Malkin, G Shvets, and NJ Fisch. "Detuned Raman amplification of short laser pulses in plasma". In: *Physical Review Letters* 84.6 (2000), p. 1208.

[73] VM Malkin, G Shvets, and NJ Fisch. "Fast compression of laser beams to highly overcritical powers". In: *Physical review letters* 82.22 (1999), p. 4448.

[74] VM Malkin, G Shvets, and NJ Fisch. "Ultra-powerful compact amplifiers for short laser pulses". In: *Physics of Plasmas* 7.5 (2000), pp. 2232–2240.

[75] VM Malkin, Z Toroker, and NJ Fisch. "Laser duration and intensity limits in plasma backward Raman amplifiers". In: *Physics of Plasmas* 19.2 (2012), p. 023109.

[76] VM Malkin, Z Toroker, and NJ Fisch. "Saturation of the leading spike growth in backward Raman amplifiers". In: *Physics of Plasmas* 21.9 (2014), p. 093112.

[77] VM Malkin, Yu A Tsidulko, and NJ Fisch. "Stimulated Raman scattering of rapidly amplified short laser pulses". In: *Physical review letters* 85.19 (2000), p. 4068.

[78] Stuart PD Mangles, CD Murphy, Zulfikar Najmudin, Alexander George Roy Thomas, JL Collier, Aboobaker E Dangor, EJ Divall, PS Foster, JG Gallacher, CJ Hooker, et al. "Monoenergetic beams of relativistic electrons from intense laser–plasma interactions". In: *Nature* 431.7008 (2004), pp. 535–538.

[79] JM Manley and HE Rowe. "Some general properties of nonlinear elements-Part I. General energy relations". In: *Proceedings of the IRE* 44.7 (1956), pp. 904–913.

[80] George Marx. "Interstellar vehicle propelled by terrestrial laser beam". In: *Nature* 211.5044 (1966), pp. 22–23.

[81] Claire Ellen Max, Jonathan Arons, and A Bruce Langdon. "Self-modulation and self-focusing of electromagnetic waves in plasmas". In: *Physical Review Letters* 33.4 (1974), p. 209.

[82] C McGuffey, AGR Thomas, W Schumaker, T Matsuoka, V Chvykov, FJ Dollar, G Kalintchenko, V Yanovsky, A Maksimchuk, K Krushelnick, et al. "Ionization induced trapping in a laser wakefield accelerator". In: *Physical Review Letters* 104.2 (2010), p. 025004.

[83] Jena Meinecke, HW Doyle, Francesco Miniati, AR Bell, R Bingham, R Crowston, RP Drake, Milad Fatenejad, Michel Koenig, Yasuhiro Kuramitsu, et al. "Turbulent amplification of magnetic fields in laboratory laser-produced shock waves". In: *Nature Physics* 10.7 (2014), pp. 520–524.

[84] José T Mendonça, Bo Thidé, and Holger Then. "Stimulated Raman and Brillouin backscattering of collimated beams carrying orbital angular momentum". In: *Physical review letters* 102.18 (2009), p. 185005.

[85] RD Milroy, CE Capjack, and CR James. "Plasma laser pulse amplifier using induced Raman or Brillouin processes". In: *The Physics of Fluids* 22.10 (1979), pp. 1922–1931.

[86] SK Mishra and A Andreev. "Amplification of ultra-short laser pulses via resonant backward Raman amplification in plasma". In: *Physics of Plasmas* 23.8 (2016), p. 083108.

[87] WB Mori, CD Decker, DE Hinkel, and T Katsouleas. "Raman forward scattering of short-pulse high-intensity lasers". In: *Physical review letters* 72.10 (1994), p. 1482.

[88] GA Mourou, NJ Fisch, VM Malkin, Z Toroker, EA Khazanov, AM Sergeev, T Tajima, and B Le Garrec. "Exawatt-Zettawatt pulse generation and applications". In: *Optics Communications* 285.5 (2012), pp. 720–724.

[89] Tatsufumi Nakamura, James K Koga, Timur Zh Esirkepov, Masaki Kando, Georg Korn, and Sergei V Bulanov. "High-power γ-ray flash generation in ultraintense laser-plasma interactions". In: *Physical review letters* 108.19 (2012), p. 195001.

[90] R Nuter and V Tikhonchuk. "Prepulse suppression and optimization of backward Raman amplification with a chirped pump laser beam". In: *Physical Review E* 87.4 (2013), p. 043109.

[91] SP Obenschain, CJ Pawley, AN Mostovych, JA Stamper, JH Gardner, AJ Schmitt, and SE Bodner. "Reduction of Raman scattering in a plasma to convective levels using induced spatial incoherence". In: *Physical review letters* 62.7 (1989), p. 768.

[92] C-H Pai, M-W Lin, L-C Ha, S-T Huang, Y-C Tsou, H-H Chu, J-Y Lin, J Wang, and S-Y Chen. "Backward Raman amplification in a plasma waveguide". In: *Physical review letters* 101.6 (2008), p. 065005.

[93] Y Ping, I Geltner, NJ Fisch, G Shvets, and S Suckewer. "Demonstration of ultrashort laser pulse amplification in plasmas by a counterpropagating pumping beam". In: *Physical Review E* 62.4 (2000), R4532.

[94] Y Ping, I Geltner, and S Suckewer. "Raman backscattering and amplification in a gas jet plasma". In: *Physical Review E* 67.1 (2003), p. 016401.

[95] Yuan Ping, Weifeng Cheng, Szymon Suckewer, Daniel S Clark, and Nathaniel J Fisch. "Amplification of ultrashort laser pulses by a resonant Raman scheme in a gas-jet plasma". In: *Physical review letters* 92.17 (2004), p. 175007.

[96] Y Ping, RK Kirkwood, T-L Wang, DS Clark, SC Wilks, N Meezan, RL Berger, J Wurtele, NJ Fisch, VM Malkin, et al. "Development of a nanosecond-laser-pumped Raman amplifier for short laser pulses in plasma". In: *Physics of Plasmas* 16.12 (2009), p. 123113.

[97] Jun Ren, Weifeng Cheng, Shuanglei Li, and Szymon Suckewer. "A new method for generating ultraintense and ultrashort laser pulses". In: *Nature Physics* 3.10 (2007), pp. 732–736.

[98] Jun Ren, S Li, A Morozov, S Suckewer, NA Yampolsky, VM Malkin, and NJ Fisch. "A compact double-pass Raman backscattering amplifier/compressor". In: *Physics of Plasmas* 15.5 (2008), p. 056702.

[99] C Riconda, S Weber, L Lancia, J-R Marquès, GA Mourou, and J Fuchs. "Spectral characteristics of ultra-short laser pulses in plasma amplifiers". In: *Physics of Plasmas* 20.8 (2013), p. 083115.

[100] CP Ridgers, Christopher S Brady, R Duclous, JG Kirk, K Bennett, TD Arber, APL Robinson, and AR Bell. "Dense electron-positron plasmas and ultraintense γ rays from laser-irradiated solids". In: *Physical review letters* 108.16 (2012), p. 165006.

[101] Marshall N Rosenbluth. "Parametric instabilities in inhomogeneous media". In: *Physical Review Letters* 29.9 (1972), p. 565.

[102] MN Rosenbluth and CS Liu. "Excitation of plasma waves by two laser beams". In: *Physical Review Letters* 29.11 (1972), p. 701.

[103] M Roth, D Jung, K Falk, N Guler, O Deppert, M Devlin, A Favalli, J Fernandez, D Gautier, M Geissel, et al. "Bright laser-driven neutron source based on the relativistic transparency of solids". In: *Physical review letters* 110.4 (2013), p. 044802.

[104] C Rousseaux, K Glize, SD Baton, L Lancia, D Bénisti, and L Gremillet. "Experimental evidence of backward raman scattering driven cooperatively by two picosecond laser pulses propagating side by side". In: *Physical review letters* 117.1 (2016), p. 015002.

[105] C Rousseaux, G Malka, JL Miquel, F Amiranoff, SD Baton, and Ph Mounaix. "Experimental validation of the linear theory of stimulated Raman scattering driven by a 500-fs laser pulse in a preformed underdense plasma". In: *Physical review letters* 74.23 (1995), p. 4655.

[106] James D Sadler, Ricky Nathvani, Piotr Oleśkiewicz, Luke A Ceurvorst, Naren Ratan, Muhammad F Kasim, Raoul MGM Trines, Robert Bingham, and Peter A Norreys. "Compression of X-ray free electron laser pulses to attosecond duration". In: *Scientific reports* 5 (2015).

[107] James D Sadler, Raoul MGM Trines, Max Tabak, Dan Haberberger, Dustin H Froula, Andrew S Davies, Sara Bucht, Luís O Silva, E Paulo Alves, Frederico Fiúza, et al. "Optimization of plasma amplifiers". In: *Physical Review E* 95.5 (2017), p. 053211.

[108] James D Sadler, Marcin Sliwa, Thomas Miller, Muhammad F Kasim, Naren Ratan, Luke Ceurvorst, Alex Savin, Ramy Aboushelbaya, Peter A Norreys, Dan Haberberger, et al. "Robustness of raman plasma amplifiers and their potential for attosecond pulse generation". In: *High Energy Density Physics* 23 (2017), pp. 212–216.

[109] JD Sadler, LO Silva, RA Fonseca, K Glize, MF Kasim, A Savin, R Aboushelbaya, MW Mayr, B Spiers, RHW Wang, R Bingham, RMGM Trines, and PA Norreys. "Advantages to a Diverging Raman Amplifier". In: *Communications Physics* (Submitted 2018).

[110] AS Sakharov and VI Kirsanov. "Theory of Raman scattering for a short ultrastrong laser pulse in a rarefied plasma". In: *Physical Review E* 49.4 (1994), p. 3274.

[111] Yuan Shi, Hong Qin, and Nathaniel J Fisch. "Laser-pulse compression using magnetized plasmas". In: *Physical Review E* 95.2 (2017), p. 023211.

[112] G Shvets, NJ Fisch, A Pukhov, and J Meyer-ter-Vehn. "Superradiant amplification of an ultrashort laser pulse in a plasma by a counterpropagating pump". In: *Physical review letters* 81.22 (1998), p. 4879.

[113] G Shvets and JS Wurtele. "Instabilities of short-pulse laser propagation through plasma channels". In: *Physical review letters* 73.26 (1994), p. 3540.

[114] A Simon, W Seka, LM Goldman, and RW Short. "Raman scattering in inhomogeneous laser-produced plasma". In: *The Physics of fluids* 29.5 (1986), pp. 1704–1718.

[115] AA Solodov, VM Malkin, and NJ Fisch. "Random density inhomogeneities and focusability of the output pulses for plasma-based powerful backward Raman amplifiers". In: *Physics of Plasmas* 10.6 (2003), pp. 2540–2544.

[116] Max Tabak, James Hammer, Michael E Glinsky, William L Kruer, Scott C Wilks, John Woodworth, E Michael Campbell, Michael D Perry, and Rodney J Mason. "Ignition and high gain with ultrapowerful lasers". In: *Physics of Plasmas* 1.5 (1994), pp. 1626–1634.

[117] K Tanaka, LM Goldman, W Seka, MC Richardson, JM Soures, and EA Williams. "Stimulated Raman scattering from UV-laser-produced plasmas". In: *Physical Review Letters* 48.17 (1982), p. 1179.

[118] Z Toroker, VM Malkin, and NJ Fisch. "Backward Raman amplification in the Langmuir wavebreaking regime". In: *Physics of Plasmas* 21.11 (2014), p. 113110.

[119] Z Toroker, VM Malkin, and NJ Fisch. "Seed laser chirping for enhanced backward Raman amplification in plasmas". In: *Physical review letters* 109.8 (2012), p. 085003.

[120] Zeev Toroker, VM Malkin, AA Balakin, GM Fraiman, and NJ Fisch. "Geometrical constraints on plasma couplers for Raman compression". In: *Physics of Plasmas* 19.8 (2012), p. 083110.

[121] RMGM Trines, F Fiuza, R Bingham, RA Fonseca, LO Silva, RA Cairns, and PA Norreys. "Simulations of efficient Raman amplification into the multi-petawatt regime". In: *Nature Physics* 7.1 (2011), p. 87.

[122] RMGM Trines, F Fiúza, R Bingham, RA Fonseca, LO Silva, RA Cairns, and PA Norreys. "Production of picosecond, kilojoule, and petawatt laser pulses via raman amplification of nanosecond pulses". In: *Physical review letters* 107.10 (2011), p. 105002.

[123] RMGM Trines, EP Alves, E Webb, J Vieira, F Fiuza, RA Fonseca, LO Silva, J Sadler, N Ratan, L Ceurvorst, et al. "Essential criteria for efficient pulse amplification via Raman and Brillouin scattering". In: *arXiv preprint arXiv:1611.04485* (2016).

[124] Yu A Tsidulko, Vladimir M Malkin, and Nathaniel J Fisch. "Suppression of superluminous precursors in high-power backward Raman amplifiers". In: *Physical review letters* 88.23 (2002), p. 235004.

[125] IC Turcu, Emma Springate, Chris A Froud, Cephise M Cacho, John L Collier, William A Bryan, GRA Nemeth, Jon P Marangos, John WG Tisch, Ricardo Torres, et al. "Ultrafast science and development at the Artemis facility". In: *ROMOPTO 2009: Ninth Conference on Optics: Micro-to Nanophotonics II*. SPIE. 2010.

[126] D Turnbull, S Li, A Morozov, and S Suckewer. "Possible origins of a time-resolved frequency shift in Raman plasma amplifiers". In: *Physics of Plasmas* 19.7 (2012), p. 073103.

[127] D Turnbull, S Li, A Morozov, and S Suckewer. "Simultaneous stimulated Raman, Brillouin, and electron-acoustic scattering reveals a potential saturation mechanism in Raman plasma amplifiers". In: *Physics of Plasmas* 19.8 (2012), p. 083109.

[128] J Vieira, RMGM Trines, EP Alves, RA Fonseca, JT Mendonça, R Bingham, P Norreys, and LO Silva. "Amplification and generation of ultra-intense twisted laser pulses via stimulated Raman scattering". In: *Nature communications* 7 (2016).

[129] Gregory Vieux, Andrey Lyachev, Xue Yang, Bernhard Ersfeld, John Patrick Farmer, Enrico Brunetti, RC Issac, Gaurav Raj, Gregor H Welsh, SM Wiggins, et al. "Chirped pulse Raman amplification in plasma". In: *New Journal of Physics* 13.6 (2011), p. 063042.

[130] G Vieux, S Cipiccia, DW Grant, N Lemos, P Grant, C Ciocarlan, B Ersfeld, MS Hur, P Lepipas, GG Manahan, et al. "An ultra-high gain and efficient amplifier based on Raman amplification in plasma". In: *Scientific Reports* 7 (2017).

[131] Hoanh X Vu, DF DuBois, and B Bezzerides. "Transient enhancement and detuning of laser-driven parametric instabilities by particle trapping". In: *Physical review letters* 86.19 (2001), p. 4306.

[132] T-L Wang, DS Clark, DJ Strozzi, SC Wilks, SF Martins, and RK Kirkwood. "Particle-in-cell simulations of kinetic effects in plasma-based backward Raman amplification in underdense plasmas". In: *Physics of Plasmas* 17.2 (2010), p. 023109.

[133] T-L Wang, D Michta, RR Lindberg, AE Charman, SF Martins, and JS Wurtele. "Feasibility study for using an extended three-wave model to simulate plasma-based backward Raman amplification in one spatial dimension". In: *Physics of Plasmas* 16.12 (2009), p. 123110.

[134] SC Wilks, WL Kruer, EA Williams, P Amendt, and DC Eder. "Stimulated Raman backscatter in ultraintense, short pulse laser–plasma interactions". In: *Physics of Plasmas* 2.1 (1995), pp. 274–279.

[135] NA Yampolsky and NJ Fisch. "Effect of nonlinear Landau damping in plasma-based backward Raman amplifier". In: *Physics of Plasmas* 16.7 (2009), p. 072105.

[136] NA Yampolsky, NJ Fisch, VM Malkin, EJ Valeo, R Lindberg, J Wurtele, J Ren, S Li, A Morozov, and S Suckewer. "Demonstration of detuning and wavebreaking effects on Raman amplification efficiency in plasma". In: *Physics of plasmas* 15.11 (2008), p. 113104.

[137] Nikolai A Yampolsky and Nathaniel J Fisch. "Limiting effects on laser compression by resonant backward Raman scattering in modern experiments". In: *Physics of Plasmas* 18.5 (2011), p. 056711.

[138] Nikolai Andreevich Yampolsky, VM Malkin, and NJ Fisch. "Finite-duration seeding effects in powerful backward Raman amplifiers". In: *Physical Review E* 69.3 (2004), p. 036401.

[139] Xu Yang, Gregory Vieux, E Brunetti, Bernhard Ersfeld, John P Farmer, Min Sup Hur, RC Issac, Gaurav Raj, SM Wiggins, GH Welsh, et al. "Chirped pulse Raman amplification in warm plasma: towards controlling saturation". In: *Scientific reports* 5 (2015).

[140] Hirokatsu Yumoto, Hidekazu Mimura, Takahisa Koyama, Satoshi Matsuyama, Kensuke Tono, Tadashi Togashi, Yuichi Inubushi, Takahiro Sato, Takashi Tanaka, Takashi Kimura, et al. "Focusing of X-ray free-electron laser pulses with reflective optics". In: *Nature Photonics* 7.1 (2013), pp. 43–47.

[141] Matthew Zepf, EL Clark, FN Beg, RJ Clarke, AE Dangor, A Gopal, K Krushelnick, PA Norreys, M Tatarakis, U Wagner, et al. "Proton acceleration from high-intensity laser interactions with thin foil targets". In: *Physical review letters* 90.6 (2003), p. 064801.