IMPERIAL

High Intensity Mid-Infrared Laser Development for Laser Wakefield Acceleration

Annabel E. Gunn
October 2025

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

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

All work presented in this thesis is my own unless otherwise stated and acknowledged.

Annabel E. Gunn

2025

The copyright of this thesis rests with the author. Unless otherwise indicated, its contents are licensed under a Creative Commons Attribution-Non Commercial-No Derivatives 4.0 International Licence (CC BY-NC-ND). Under this licence, you may copy and redistribute the material in any medium or format on the condition that; you credit the author, do not use it for commercial purposes and do not distribute modified versions of the work. When reusing or sharing this work, ensure you make the licence terms clear to others by naming the licence and linking to the licence text. Please seek permission from the copyright holder for uses of this work that are not included in this licence or permitted under UK Copyright Law.

Abstract

This thesis presents a detailed investigation into the feasibility of a laser wakefield acceleration (LWFA) experiment driven by the 3.7 µm centred mid-infrared (MIR) arm of Chimera, a high-power, short-pulse, multi-beam optical parametric chirped pulse amplification system under development at Imperial College London. The work addresses the challenges of commissioning a developmental laser system for routine experimental operation, with emphasis on front-end stability, and the theoretical viability of MIR-driven LWFA through particle-in-cell (PIC) simulations. MIR drivers offer potential advantages for LWFA, including higher normalised vector potentials at lower intensities and increased accelerated charge, though experimental implementation remains limited by the scarcity of suitable high-power sources.

Guided by PIC simulations, MIR-driven LWFA is shown to be feasible within Chimera's current design constraints. Self-injection and acceleration of electrons to energies of $\sim$ 20 MeV are achieved at a plasma density of $n_e = 5.8 \times 10^{18}$ cm^{-3}, provided the laser delivers at least 15 mJ with a pulse duration below 70 fs. Pulse energy is identified as the dominant parameter, while pulse duration is comparatively forgiving, indicating that future development should prioritise energy scaling. Although the bandwidth at MIR generation – via difference-frequency generation in PPLN – could support compressed pulses as short as $\sim$ 30 fs, higher-order phase accumulation and bandwidth reduction in subsequent optical parametric amplifiers suggest an initial compressed pulse duration near $\sim$ 70 fs.

Simulations also reveal electron energies exceeding established scaling predictions, along with substantial accelerated charge. This is attributed to operation in a high-density regime enabled by short-pulse, long-wavelength laser-drivers. Early beam-loading effects suggest the onset of a hybrid beam-driven phase, potentially enhancing laser-plasma efficiency and beam current. The apparent insensitivity of this regime to acceleration length could simplify experimental implementation. If confirmed experimentally, this could position mid-infrared drivers as promising sources of high-charge, high-current electron beams for future medical and industrial applications.

Contents

Abstract **4**

Role of the author **13**

Acknowledgements **15**

1 Introduction **17**
1.1 Lasers . . . 19
1.2 Laser-Plasma Acceleration . . . 20
1.3 Motivation for the Mid-Infrared . . . 23
1.4 Organisation of Thesis . . . 23

2 Theory – Laser Technology **25**
2.1 Nonlinear Light-Matter Interactions . . . 25
2.2 Nonlinear Mixing . . . 29
2.2.1 Second Harmonic Generation (SHG) . . . 30
2.2.2 Difference Frequency Generation (DFG) . . . 35
2.2.3 Optical Parametric Amplification (OPA) . . . 37
2.3 Time-Bandwidth Product (TBP) . . . 39
2.4 Chirped Pulse Amplification (CPA) . . . 39
2.4.1 Stretching and Compressing a High Intensity Pulse . . . 44
2.5 Optical Parametric Chirped Pulse Amplification (OPCPA) . . . 45
2.6 Carrier Envelope Phase (CEP) . . . 47
2.7 Optical Damage . . . 49
2.7.1 Self-focusing . . . 49
2.7.2 B-Integral . . . 50
2.8 Kerr Lens Modelocking (KLM) . . . 51

3 Theory – Laser Wakefield Acceleration **54**
3.1 Single Particle Motion . . . 54
3.1.1 The Ponderomotive Force . . . 56
3.1.2 Relativistic Motion . . . 59
3.2 Nonlinear Optics in Underdense Plasmas . . . 61

3.2.1 EM Waves in Plasmas . . . 61
3.2.2 Relativistic Self-Focussing . . . 63
3.2.3 Pulse Compression and Photon Acceleration . . . 64
3.3 Laser Wakefield Acceleration (LWFA) . . . 66
3.3.1 Plasma Wave Formation . . . 66
3.3.2 Electron Injection . . . 68
3.3.3 The Bubble Regime . . . 69
3.3.4 Limitations on Electron Energy Gain . . . 70
3.3.5 Lu Scaling Model . . . 72
3.4 Direct Laser Acceleration . . . 74

4 Methods 75
4.1 Laser Diagnostics . . . 75
4.1.1 Third Order Autocorrelator Diagnostic (TOAD) . . . 77
4.1.2 Mid-Infrared Diagnostics . . . 79
4.2 Numerical Methods . . . 80
4.2.1 Particle-in-Cell (PIC) Simulations . . . 81
4.2.2 The PIC Algorithm . . . 82
4.2.3 Resolution Tests . . . 83

5 The Chimera Multi-Beam Mid-Infrared OPCPA Laser System at Imperial College London 89
5.1 System Overview . . . 92
5.2 The Femtolasers Rainbow Oscillator . . . 94
5.3 The Photonic Crystal Fiber . . . 98
5.4 Gain-Storage Laser Amplifiers . . . 102
5.4.1 The Linear Regenerative Amplifier . . . 102
5.4.2 The Ring Booster . . . 107
5.4.3 Flashlamp Pumped Chain . . . 109
5.5 Optical Parametric Amplification and Mid-Infrared Generation Stages . 112
5.5.1 800nm OPA sub-system . . . 112
5.5.2 Mid-Infrared Generation . . . 115
5.5.3 Mid-Infrared OPA Stages . . . 119
5.6 Stretching and Compressing the Mid-Infrared Pulse . . . 122
5.7 Experimental Output of the Chimera Laser System . . . 132
5.8 Multi-beam Capabilities of the Chimera laser system . . . 135
5.9 Summary . . . 136

6 Mid-Infrared Laser Wakefield Acceleration Simulations : An exploration of high electron energies 137
6.1 Plasma Density Dependence of Accelerated Electron Energy . . . 140
6.1.1 Normalised Vector Potential . . . 147
6.1.2 Carrier Envelope Phase (CEP) . . . 151

6.1.3 Direct Laser Acceleration (DLA) . . . 153
6.1.4 Bubble Size . . . 155
6.1.5 Beam Loading . . . 158
6.2 Discussion . . . 163
6.3 Summary . . . 165

7 Design Considerations for a Laser Wakefield Acceleration experiment driven by the Chimera laser system 167
7.1 Laser Parameters . . . 168
7.2 Plasma Considerations . . . 174
7.3 Discussion . . . 179
7.4 Summary . . . 181

8 Conclusions and Future Outlook 183

A Appendix A 186
A.1 Fundamental Constants . . . 186
A.2 The definition of the full-width-at-half-maximum . . . 186

B Appendix B 189
B.1 Relativistic Single Particle Motion . . . 189

Bibliography 192

List of Figures

1.1 An aerial view of the current and proposed particle colliders at CERN. . 18
1.2 A simplistic evolution of laser power increases as a result of technological advances. 20
1.3 A photograph of a plume of a gas at the moment a high intensity laser is fired into it to drive a plasma accelerator. 21

2.1 Energy level diagrams to describe SFG, DFG and OPA processes. 30
2.2 Second harmonic intensity signal as a function of coherence length. . . . 34
2.3 Parametric fluorescence as a result of spontaneous down conversion of the signal and idler beams in an OPA. 38
2.4 A cartoon showing the concept of chirped pulse amplification (CPA). . . 40
2.5 Electric field of a chirped transform-limited pulse. 42
2.6 The effects of second, third and forth order spectral phase on laser pulses. 43
2.7 Optical layout of a simple two-parallel-grating system used to introduce negative GDD. 44
2.8 Configuration of group velocities and wavevectors for non-collinear phase-matching in OPAs for group velocity, v_g and wavevector k matching. . . 47
2.9 The predicted electric field of the Chimera mid-infrared few-cycle laser pulse (30 fs at 3700 nm) to illustrate the effects of carrier envelope phase (CEP). 48
2.10 Electric field as a function of pulse duration and carrier envelope phase (CEP). 49
2.11 A schematic cavity representation of a typical Kerr lens modelocking oscillator. 52
2.12 A geometric representation of the four stability limits and corresponding cavity modes present within a Kerr lens modelocking cavity. 53

3.1 The trajectory of an electron while in the presence of an infinite, linearly polarised, monochromatic plane wave with field strength a_0 and wavenumber k in the lab and drift frames. 60
3.2 An electron density plot from an AM cylindrical simulation of LWFA using the PIC code Smilei. 70

3.3 The dephasing and depletion lengths plotted as functions of a_0 and plasma density n_e. 73

4.1 Optical layout of the Imperial College Third Order Autocorrelator Diagnostic (TOAD). 78
4.2 Axial plasma density and longitudinal electric field plotted for simulations with varying longitudinal resolution, azimuthal modes and particles per cell. 86
4.3 Final energy spectrum and evolution of the 90^{th} percentile of all electron longitudinal momenta through time for varying longitudinal resolution, azimuthal modes and particles per cell. 87

5.1 A simplified component diagram of the Chimera laser system. 92
5.2 The Femtosource Laser cavity. 94
5.3 Comparison of Chimera oscillator spectra. 96
5.4 Measured Chimera oscillator spectrum and pump power VS modelocked power efficiency plot. 97
5.5 The Chimera oscillator spectrum over a period of four consecutive days. 98
5.6 The Photonic Crystal Fiber (PCF). 99
5.7 Experimental output of the PCF. 100
5.8 A simplified top-down view of the optical path of the chirped Bragg grating. 101
5.9 The spectral profile and oscilloscope trace of the 54-pass diode pumped Nd:YLF linear regenerative amplifier. 103
5.10 Output energy of the Nd:YLF linear regenerative amplifier whilst varying diode temperature, recorded over 10,000 shots 105
5.11 Output energy of the Nd:YLF linear regenerative amplifier whilst varying diode driving current. 106
5.12 Beam profiles showing the mode before and after passing through the 54-pass linear regenerative amplifier. 107
5.13 Diode pumped Nd:YLF ring regenerative amplifier. 108
5.14 The Ring Booster Amplification Loop. 110
5.15 Beam profiles showing the output mode of the ring regenerative amplifier before and after the addition of significant beam shaping. 111
5.16 The Nd:YLF Flashlamp Pumped Chain. 113
5.17 A TOAD trace showing the ring amplifier output (at 1053 nm) after compression. 114
5.18 Beam profiles, taken using a WinCamD-LCM – 1" CMOS Beam Profiler, showing the input and output modes of the SHG stage. 114
5.19 The 800 nm seed spectrum before and after amplification in a BBO OPA stage. 115
5.20 The output MIR spectrum of the DFG stage as a function of PPLN temperature. 117
5.21 Experimental output of the MgO:PPLN DFG stage. 118

5.22 A beam profile of the amplified MIR beam (centred at 3700 nm) after the first MIR OPA stage (OPA-B1), measured with a Spiricon Pyrocam IIIHR beam profiler. 119
5.23 Experimental output of the OPA-B1 stage. 120
5.24 Saturated, bright MIR detection on a silicon ship CCD camera, generated by three-photon absorption in the detector. 121
5.25 A spectrum taken at the output of the second mid-infrared OPA on the Chimera laser system by Dr Nicholas Stuart. 122
5.26 The effect of third- and fourth-order spectral phase on the compressibility of an ultrashort pulse with $\lambda = 3.7$ µm and a transform-limited pulse duration of 30 fs. 123
5.27 The effect of third- and fourth-order spectral phase on the compressibility of an ultrashort pulse with $\lambda = 3.3$ µm and a transform-limited pulse duration of 66 fs. 124
5.28 Optical layout of the Chimera mid-infrared off-axis Öffner stretcher. . . 126
5.29 Image-on-screen and phase analysis output of the Optometrika raytracing code simulating an Öffner style stretcher system for the Chimera MIR laser. 127
5.30 Image-on-screen and phase analysis output of the Optometrika raytracing code simulating a Tracey style compressor system for the Chimera MIR laser. 130
5.31 The effect of high order spectral phase on the compressibility of an ultrashort pulse with $\lambda = 3.7$ µm and a pulse duration of 50 fs. 131
5.32 The compressor grating 3D printed mount designs. 132
5.33 MIR beam profile and high harmonic spectral data collected from a sample of crystalline ZnSe. 133
5.34 Measured MIR harmonics as a function of photon energies. 135

6.1 Electron waterfall plot (energy-time distribution) and energy spectrum for a LWFA simulation at a starting plasma density of 2.8×10^{18} cm^{-3} . 142
6.2 A 3D waterfall plot showing the transverse laser evolution for a simulation with starting density 2.8×10^{18} cm^{-3}. 143
6.3 The 90$^{\text{th}}$ percentile of accelerated electron energies for each starting density $1.8 - 5.0 \times 10^{18}$ cm^{-3}. 144
6.4 The highest electron energy with plasma density for a selection of laser-driven wakefield experiments. 145
6.5 Figure 6.4 is replotted with individual results classified into experiments which are self-injecting/self-guided, externally guided, staged, or using downramp injection, optical injection and ionisation injection. 146
6.6 A laser kymograph for a simulation with starting density 2.8×10^{18} cm^{-3}. 148
6.7 Evolution of the normalised vector potential a_0 through a simulation with starting parameters: $a_0 = 1.9$, $n_e = 2.8 \times 10^{18}$ cm^{-3}. 150
6.8 Electron waterfall momentum spectra for LWFA simulations run while scanning the laser carrier envelope phase (CEP) term. 152

6.9 The laser electric field E_z for both cosine and sine carrier envelope phases ($\phi = 0$, $\phi = \pi/2$), 0.27 ps into the simulation. 153
6.10 A map over the electron momentum phase space for a simulation with starting density 2.8×10^{18} cm^{-3}. 154
6.11 The DLA $W_\perp$ and LWFA $W_\parallel$ contributions to the final energy of randomly selected tracked electrons versus their final energies. 155
6.12 Procedure for quantitatively extracting the plasma bubble radius from a PIC simulation. 156
6.13 The electron density map at 0.55 ps is overlaid with the calculated bubble centre, the central horizontal bubble axis and the ellipse inferred from tracking the bubble shape. 157
6.14 Electron density map, with starting plasma density $n_\mathrm{e} = 2.8 \times 10^{18}$ cm^{-3} at 1.69 ps. 158
6.15 Electron density map with starting plasma density $n_\mathrm{e} = 2.8 \times 10^{18}$ cm^{-3} with additional 1D density and longitudinal field E_x line-outs, showing beam-loading effects during the course of the simulation. 160
6.17 A colour map, used to explore possible trends in (a) number of optical cycles, (b) a_0, and (c) P/P_c across the full database is applied to the Figure 6.5 data. 163

7.1 Electron energy spectra for 1D scans of laser energy, E_L, pulse duration, τ_FWHM and focal spot, d_FWHM. 169
7.2 The 90$^\mathrm{th}$ percentile of accelerated electron energies against simulated normalised vector potential a_0 at a plasma density of $n_\mathrm{e} = 2.8 \times 10^{18}$ cm^{-3}. 170
7.3 The 90$^\mathrm{th}$ percentile of accelerated electron energies against simulated normalised vector potential a_0 at plasma densities of $n_\mathrm{e} = 1.8 \times 10^{18}$ cm^{-3} and $n_\mathrm{e} = 3.8 \times 10^{18}$ cm^{-3}. 171
7.4 The 90$^\mathrm{th}$ percentile electron energies extracted from a 3-dimensional parameter scan across laser energy, pulse durations and plasma density. 172
7.5 The 90$^\mathrm{th}$ percentile electron energies extracted from a 3-dimensional parameter scan at a plasma density of $n_\mathrm{e} = 5.8 \times 10^{18}$ cm^{-3}. 173
7.6 The 90$^\mathrm{th}$ percentile of accelerated electron energies against simulated normalised vector potential a_0 at a plasma density of $n_\mathrm{e} = 5.8 \times 10^{18}$ cm^{-3}. 174
7.7 Acceleration distance corresponding to the maximum (90$^\mathrm{th}$ percentile) energies reached by injected electrons. 175
7.8 Representative electron spectra across the explored parameter space. . . 177
7.9 Evolution of in density space of a simulation at $n_\mathrm{e} = 2.8 \times 10^{18}$ cm^{-3} between 1.7 - 6.0 ps. 178
7.10 The total charge at maximum average energy extracted over a 3-dimensional parameter scan across laser energy, pulse duration and plasma density. . 180

A.1 The temporal intensity and (normalised) field Gaussian representations of the 3.7 µm central wavelength Chimera laser pulse. 188

List of Tables

2.1 The square sum field expansion of the second order nonlinear polarisation, with terms identified by the resultant 'new' frequency. 31

4.1 The simulation parameters chosen for all following simulation results. . . 88

5.1 Summary of MIR OPCPA systems comparable to the Chimera laser system. 91

6.1 A comparison of parameters for a LWFA experiment driven by either the Gemini, Salle Noire or Chimera laser systems. 139

6.2 The full set of parameters used to simulate the Chimera laser system using the PIC code Smilei. 141

A.1 To 4 significant figures, the fundamental constants used within this thesis. 186

Role of the author

The Chimera laser system was designed and installed by Dr Nicholas Stuart. The author inherited the system upon the start of this PhD project and has since been the sole user, operator and developer, being responsible for all repairs, upgrades and optimisations. This includes the continued characterisation, development, operation and maintenance of the laser system, as well as the commissioning into experimental operation. A technical paper reviewing the Chimera laser system is currently in preparation.

The design of the high harmonic generation experiment performed using the Chimera laser system was conducted by the author and Jack Morse, with samples provided by Dr Mary Matthews. The author was responsible for beam delivery to target, as well as support with target alignment and optimisation.

All simulations presented in this thesis were performed and analysed by the author using the particle-in-cell code Smilei, with the CX1 and CX3 computing clusters at Imperial College London. Smilei is an open-source and collaborative code, co-developed by both physicists and experts in high-power computing.

The author has also been involved in two off-site experiments: one at the Gemini laser at Rutherford Appleton Laboratories, UK and one at the Extreme Light Infrastructure - Nuclear Physics, Romania. Experiments of this nature are highly synergistic, involving a very large experimental team. While the results obtained from these experiments are not being presented in this thesis, the author was heavily involved in the setup and running of both experiments. Publications resulting from this work are currently in preparation.

The author has presented results from this work at several High Power Lasers Christmas meetings, the ELI summer school, the 6th European Advanced Accelerator Concepts workshop and the Laser-Plasma Accelerators Workshop, the latter two of which were supported by the European Union's Horizon 2020 Research and Innovation programme under Grant Agreement No 101004730 "I.FAST", as won by the author. This work has been supported by an EPSRC DTP research studentship.

Publications involving the author:

- E. Gerstmayr, B. Kettle, M. J. V. Streeter, L. Tudor, O. J. Finlay, L. E. Bradley, R. Fitzgarrald, T. Foster, P. Gellersen, A. E. Gunn, O. Lawrence, P. P. Rajeev,

B. K. Russell, D. R. Symes, C. D. Murphy, A. G. R. Thomas, C. P. Ridgers, G. Sarri and S. P. D. Mangles, "High brightness multi-MeV photon source driven by a petawatt-scale laser wakefield accelerator", arXiv, 10.48550/arXiv.2506.23718, (2025) (Submitted 01/07/25, under review).

- A. E. Gunn, N. Stuart, J. Morse and R. A. Smith, "Design, performance and control of a short-pulse, high-power, mid-infrared OPCPA laser system", in prep.
- D. Duszynski, J. Grieves, J. Morse, A. E. Gunn and R. A. Smith, "LED-based autocorrelation at large crossing angles for precision pump-probe timing", in prep.

Acknowledgements

It is difficult to believe that before starting this PhD, I had never been inside a laser lab, let alone largely in charge of one. For this, I have Roland Smith to thank: despite having more water pipes burst on me than I could have possibly predicted, I also can not imagine a better way to spend the past four years. For the trust, the extreme patience, for the many, many coffees, and for always being the first person in the lab when something goes wrong (or there is water pouring through the ceiling...), I am extremely grateful. I would also like to thank my other supervisor, Stuart Mangles. It was rare that I came out of a meeting without something new to think about. Having two supervisors who care so deeply about their work has been a real privilege that I do not take for granted.

I am also completely indebted to Nick Stuart – laser whisperer supreme, and the brains behind Chimera. Despite there being no intersection in our H148 tenures, he has always been on hand to help with no judgment, even late at night when I touched something I shouldn't have.

To the people who made my experience on facility experiments an absolute pleasure: Brendan, Jon, Matt and Elias – thank you for putting up with my pestering, and for teaching me so much. To Ollie, Mickey, Rod, Jon, Nick and Spencer: for Friday wine nights, pub trips, crosswords and innumerable coffee breaks – never a dull moment! I want to particularly mention the Plasma Social 3.0 division... It's never 'just the one', is it? To Dariusz and Jack, thanks for coming to join 'Team Roly'. Jack – thank you for being so helpful and understanding in my last 6 months of lab chaos. To Cornelia and Emily – my wonderful physics friends – I can't believe how lucky I am to have met you both.

To my long-suffering friends outside of physics: thank you for sticking with me. To Helena. A better version of me in every way imaginable (other than height, which I must cling to) and probably the most remarkable person I know – in my grandma's words: "Helena's rather an amazing person I think. A jolly good friend".

My grandma passed away almost four years ago now, roughly two weeks into starting this PhD. A second mother and always my biggest supporter, I imagine her constant refrain throughout my time at Imperial would have been "Is your machine working yet?". Well, the answer is still no, but I dedicate this to her nonetheless.

To my parents, Caroline, Iain and John, and the rest of my family. Thank you for always being so supportive of the things I choose to do, and for pretending to listen when I tell you what I've been up to. I wouldn't be doing this without you. To Janet, for my weekly breakfast, and for stepping into the role of grandma for me without a second thought. To Christine and David. Having always been assured that the writing-up period was the most stress-inducing part of the whole PhD, I couldn't have asked for a more peaceful way to spend it.

Finally, to Bogg. I truly do not know what I would do without you.

1 Introduction

With its history spanning over centuries of academic research, the particle accelerator is one of the most fundamental and versatile tools accessible to the scientific community. Since the first demonstration of a potential difference produced by opposing charges by Johan Wilcke in 1762 [1], the advent of the electric battery by Alessandro Volta in 1800 [2] – for whom the 'volt' is so-named, to Rutherford scattering particles off a gold foil in 1911 [3], the manipulation and acceleration of charged particles has been a tool used to probe and understand the minutiae of our universe. Over a hundred years in the wake of Rutherford's keystone experiment, particle accelerators have dramatically evolved beyond what the savants of the 18^{th} century could have predicted. Today, following the completion of the Standard Model with the discovery of the Higgs boson in 2012 [4–6], high-energy particle colliders such as those at CERN and SLAC seek to continue to advance our collective knowledge of the cosmos.

Beyond the pursuit of *scientia gratia scientiae* – the concept of "knowledge for the sake of knowledge" – lie practical and transformative applications of energetic particles. Far from being confined to kilometres of underground tunnels, particle accelerators appear in many guises: in hospitals treating cancer through proton therapy [7]; in the production of isotopes for Positron-Emission-Tomography (PET) scans [8]; in techniques such as accelerator mass spectroscopy for radiocarbon dating [9]; in semiconductor fabrication [10]; in food sterilisation [11]; and even in art, where x-rays reveal images hidden beneath the surface of priceless paintings [12].

Relying on electric and magnetic fields to accelerate and guide particles, modern particle accelerators have evolved from electrostatic devices, such as the Nobel Prize winning Cockcroft-Walton generator in 1932 [13], to time-varying accelerators, initially proposed by Ising in 1924 [14] with his proposal considered the origin of the first 'true' particle accelerator [15]. These accelerators typically make use of resonant acceleration, using radio-frequency (RF) cavities – metal structures that sustain oscillating electromagnetic fields at a fixed frequency. In linear accelerators (linacs), a series of such cavities is spaced so that particles remain in an accelerating phase as they propagate. Circular accelerators, on the other hand, use magnetic fields to keep particles on a closed path. Although linacs require particles to pass through each accelerating structure only once, circular geometries allow them to circulate repeatedly through the same RF

cavities, gaining energy on every lap. This reuse of accelerating fields substantially raised the attainable energy limit, but both techniques remain fundamentally constrained by electrical breakdown in the accelerating structures, placing an upper limit on the electric field gradient that they can support. Circular accelerators, such as synchrotrons, face an additional limitation: when charged particles are bent along a curved trajectory, they unavoidably emit synchrotron radiation, often in the X-ray regime, losing energy on every turn. To compensate, either more RF power is required, or the bending radius must be increased, meaning that colliders must be made physically larger to reach higher energies efficiently. For these reasons, linacs and synchrotrons have, over the years, grown steadily larger. Superconducting RF cavities can, in principle, sustain gradients approaching 100 $\mathrm{MV\,m^{-1}}$ before breakdown; in practice, however, reliable operation is limited to around 20 - 30 $\mathrm{MV\,m^{-1}}$ and, the accelerating field being capped, the target particle energy directly determines the required accelerator length [16]. The Large Hadron Collider (LHC), currently the largest particle accelerator in the world, has a circumference of 27 km and cost over £3 billion to build. Its proposed successor, the Future Circular Collider (FCC), would surpass it both in scale and expense (see Figure 1.1), spanning a circumference of $\sim$ 100 km at an estimated cost of £20 billion.

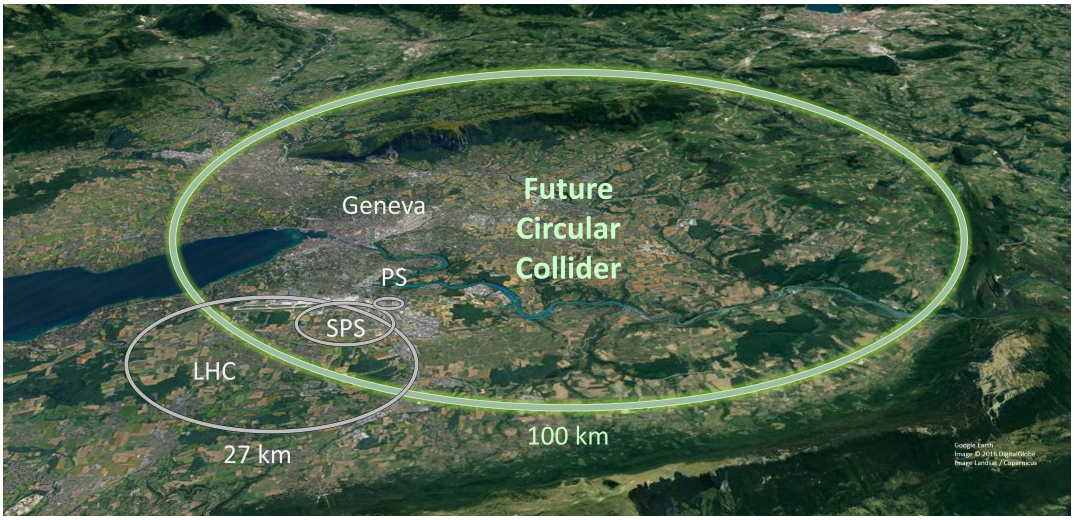


Figure 1.1: An aerial view of the current (LHC – 27 km circumference) and proposed (FCC – 100 km circumference) particle colliders at CERN. Credit: CERN.

The vast expense and diminishing returns associated with accelerating particles to higher energies in vacuum motivate the search for alternative approaches – ones not limited by the technological ceiling of electrical breakdown. Beyond advancing particle physics, many fields, particularly medicine, require less energetic beams; yet conventional accelerators remain too large, costly, and cumbersome for most hospitals to accommodate. The hope, therefore, is that compact, high-gradient accelerators might not only broaden applications in medicine, industry, and materials science, but also provide the higher accelerating gradients sought in particle physics. For this shift, we turn to a different kind of particle acceleration – one that harnesses the fourth state of matter: plasma.

As plasmas are already fully – or partially – ionised, they are not subject to the same constraints that limit RF cavities and can, in principle, sustain electric field gradients thousands of times larger than those achievable in conventional RF structures. Realising such extreme conditions, however, requires a driver capable of delivering ultra-intense, ultra-short bursts of energy. The advent of the pulsed laser in 1960, with Theodore Maiman demonstrating the first operational device and reportedly describing it as "a solution seeking a problem" [17, 18], and the subsequent development of high-power laser technology, provided just such a tool, opening the door to laser-plasma acceleration.

1.1 Lasers

Since Maiman's ruby laser, laser technology has found countless applications, spurring entire realms of scientific and medical research, as well as wide-ranging everyday use. At its core, a 'classical' laser relies on a population inversion between electron states within a gain medium – a material capable of amplifying radiation. This principle has given rise to many varieties of lasers, including gas lasers such as helium-neon and carbon dioxide, semiconductor lasers such as gallium arsenide (GaAs) laser diodes, and solid-state lasers such as titanium-sapphire (Ti:Sa), neodymium-doped yttrium lithium fluoride (Nd:YLF) and neodymium-doped yttrium aluminium garnet (Nd:YAG). A population inversion is established by a 'pump', which may, for example, be an electrical discharge or a light source, such as another laser or a flashlamp. The gain medium is enclosed in a resonant cavity formed by two mirrors, one highly reflective and the other partially transmitting, so that photons undergo multiple passes of amplification before emerging as an intense, quasi-monochromatic beam.

In the decades following its advent, there has been a steady drive to achieve ever higher peak laser powers and intensities. Shown in Figure 1.2, progress in this field has often advanced in leaps of several orders of magnitude, enabled by technological breakthroughs – firstly due to the invention of Q-switching in 1961 [19], only a year after the laser itself, followed shortly by the development of mode-locking in 1964 [20–24]. Progress slowed until the emergence of chirped pulse amplification (CPA) in 1985 by Donna Strickland and Gérard Mourou [25], revolutionising the field of high-intensity lasers and their applications, and later earning them the 2018 Nobel Prize in physics.

Beyond CPA, the field of nonlinear optics has provided additional tools for high-power laser development, enabling frequency conversion and parametric amplification. In particular, optical parametric chirped pulse amplification (OPCPA) has emerged as a powerful successor to CPA, offering broader gain bandwidths and higher peak powers while imposing lower thermal constraints than traditional population inversion-based gain media. As a result, femtosecond-scale lasers are now commonplace in fields as diverse as chemistry, biology and high-intensity science, with current efforts pushing toward the attosecond regime, in pursuit of ever higher intensities.

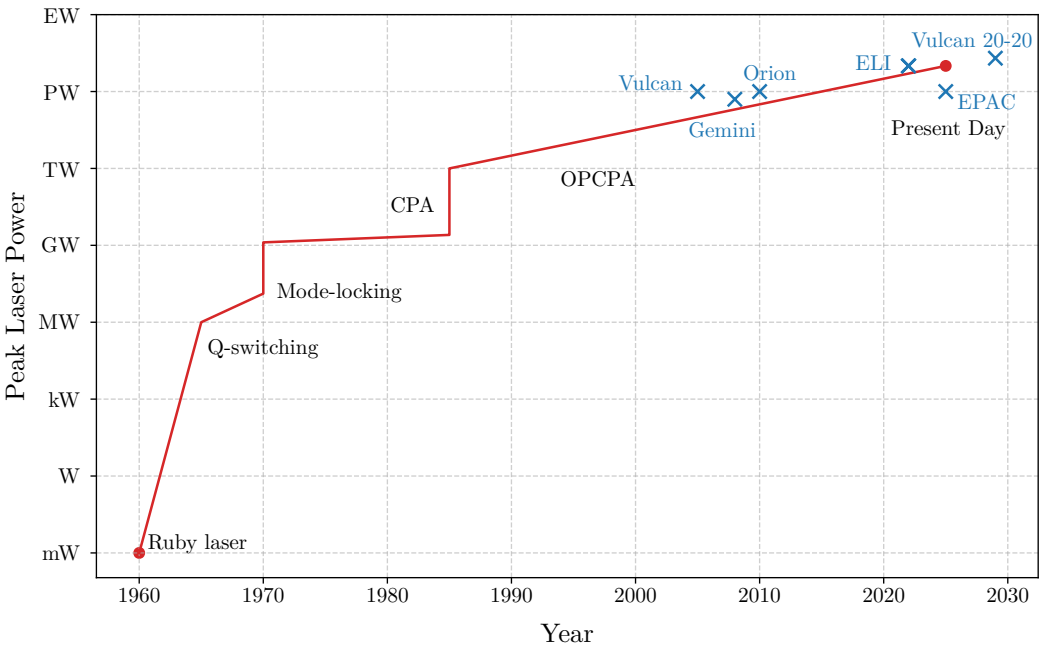


Figure 1.2: A simplistic evolution of laser power increases as a result of technological advances, from the advent of the ruby laser in 1960 to the present day. The current and future high-power laser systems available at UK national laboratories are shown in blue, with the European Extreme Light Infrastructure (ELI) included as a major international partner.

The unique properties of high-power femtosecond-scale lasers have allowed laser research to extend into a multitude of wide-ranging applications across the fields of science, technology, and medicine, exemplifying their variety and versatility. Examples of these include precision spectroscopy [26], high-harmonic generation [27], and laser-driven particle acceleration [28–31]; fiber-optic communications [32] and microelectronics [33]; and, precision surgery [34] and the fabrication of cardiological micro-stents [35]. The most powerful lasers of this kind currently operate in the visible and near-infrared (NIR) spectral range. Ti:Sa systems are capable of the shortest pulse durations, while Nd-doped materials typically produce the highest pulse energies of $\mathcal{O}(1\text{ kJ})$. The most powerful systems to date are Ti:Sa lasers at 800 nm, producing ∼10 PW peak power at ∼20 fs [36]. These extreme peak powers and ultrashort pulse durations provide precisely the conditions required to drive plasma waves, making high-power lasers an ideal tool for plasma-based particle acceleration.

1.2 Laser-Plasma Acceleration

The concept of laser-plasma acceleration was proposed by Toshiki Tajima and John Dawson in 1979, who, through computer simulation, showed that an intense electromagnetic pulse (such as a high-intensity laser) could drive plasma waves capable of accelerating electrons to gigaelectronvolt energies per centimetre of accelerating distance [28]. In their seminal paper, they explored several mechanisms for exciting such plasma waves, all based on the principle that a photon-derived ponderomotive force can drive a plasma wave

with a phase velocity matching that of the driver. Amongst the mechanisms discussed were beat-wave excitation, in which two laser beams with a small frequency difference $\Delta\omega = \omega_\mathrm{p}$ interfere to resonantly drive a plasma wave; Raman forward scattering (RFS); and, using a laser pulse length on the timescale of the plasma frequency, $\tau \approx \omega_\mathrm{p}^{-1}$, what is now known as standard laser wakefield acceleration (LWFA). In the LWFA regime, an intense, ultrashort laser pulse propagates into a low-density gas and, with its leading edge, ionises it into a plasma. As the laser propagates, electrons are expelled from its path by the ponderomotive force, leaving behind a positively charged cavity and setting up large-amplitude plasma oscillations in its wake. These plasma waves, oscillating at the characteristic plasma frequency, sustain strong longitudinal electric fields capable of accelerating electrons to relativistic energies. An example of the LWFA regime, where a high-intensity laser pulse drives a plasma wave, is shown in Figure 1.3.

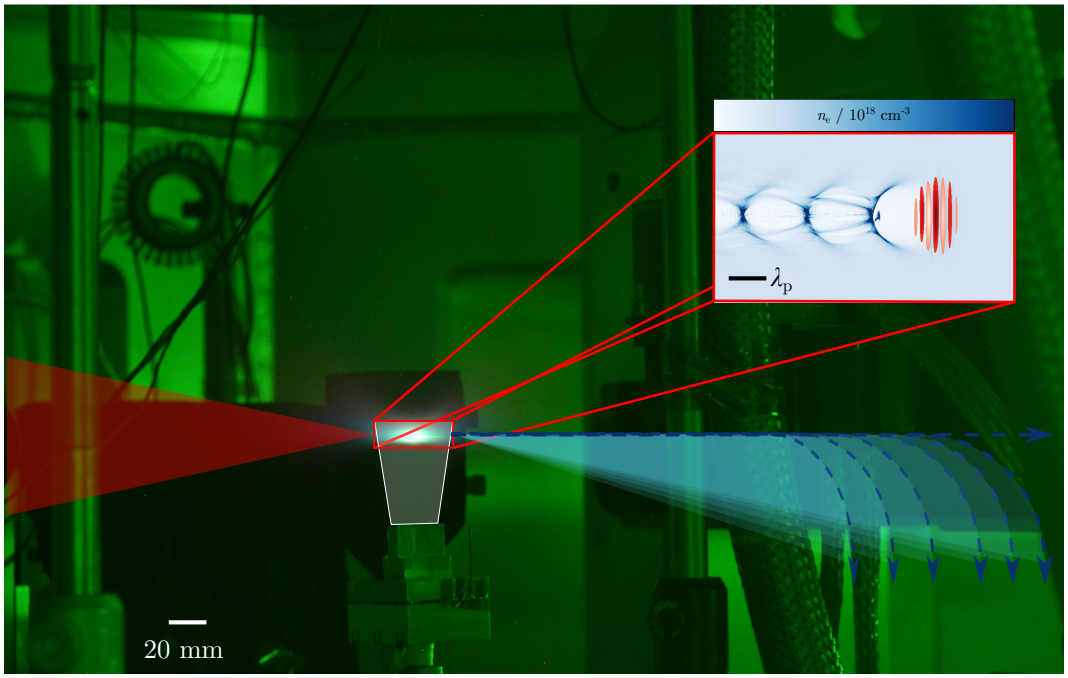


Figure 1.3: A photograph, taken at ELI-NP in Romania, of a plume of gas at the moment a high intensity laser is fired into it to drive a plasma accelerator. The laser pulse, illustrated in red, is focused into the gas (white), which is ionised by its leading edge to form a plasma – this can be seen in the photograph by its bright emission. The inset shows a simulated view of the interaction region, illustrating the initial stages of electron acceleration in a laser wakefield accelerator. A small population of electrons from the background plasma (blue) has been injected into the wave and is being accelerated behind the laser pulse (red). In the photograph, the accelerated electrons are represented by the blue dashed lines exiting the gas plume. Their curved trajectories represent an interaction with a magnet spectrometer so that their energy may be extracted. The green hue is due to a coloured filter in front of the camera, used to block out the laser light.

Many proof-of-principle experiments followed Tajima and Dawson's 1979 paper, but it was not for another 13 years that the first clear experimental demonstration of plasma electron acceleration (using beat-wave excitation) was achieved [37–39]. This delay was primarily due to the lack of sufficiently intense electromagnetic drivers, as lasers at the time were limited by the damage threshold of optical components; a barrier which

was overcome by the pioneering development of chirped pulse amplification (CPA) in 1985 [25], which enabled the amplification of ultra-intense, ultrashort laser pulses and prompted a rapid increase in experimentally available laser intensities. In the following decades, a series of landmark experiments confirmed the feasibility of plasma-based acceleration schemes for electrons [40–43], while intense laser-solid interactions also enabled proton and ion acceleration by alternative mechanisms such as target normal sheath acceleration (TNSA) [44–47].

In 2004, three independent experiments reported the first generation of high-quality, quasi-monoenergetic electron beams using LWFA. These beams exhibited both low divergence ($\sim$ 1 mrad) and high charge ($\sim$ 10 pC), with energies peaking at $\sim$ 100 MeV [29–31]. The results have since allowed the LWFA regime to dominate the field, with higher and higher electron energies being achieved through methods such as relativistic self focusing [48], capillary-discharge waveguides [49–55] and laser pre-ionised plasma channels [56–61]. Since then, continued advances in laser technology, target design, and diagnostics have enabled LWFA systems to routinely produce GeV-class electron beams. At the time of writing, energies approaching 10 GeV have been demonstrated using meter-scale, channel-guided laser plasma accelerators [62].

While these milestones suggest that the field has moved beyond its infancy, significant challenges remain. Chief amongst them are increasing the accelerated charge and energy, improving beam quality and stability, and scaling the technology to higher repetition rates, all of which are necessary before compact, high-gradient accelerators become viable for medicine and industry. Ongoing research is tackling these issues on several fronts: particle energies are being extended through pre-formed plasma channels [62] and multi-stage acceleration schemes [63]; wall-plug efficiency and repetition rate are being addressed with more efficient laser drivers [64, 65] and lower power systems [66]; and beam quality and charge are being improved through advances in controlled injection techniques [67–71] and tailored plasma profiles [72, 73].

To date, most experiments have relied on near-infrared Ti:Sa lasers, operating at a centre wavelength near 800 nm. Yet the wavelength of the driving laser remains a largely unexplored frontier. Fundamental scaling laws [74] indicate that long wavelength drivers could provide substantial benefits: increasing the effective normalised vector potential for a given intensity ($a_0 \propto \lambda$), (which also allows the same a_0 to be reached with lower-power, lower-intensity systems); lowering the critical density ($n_c \propto \lambda^{-2}$); and increasing the total accelerated charge ($N \propto \lambda$). These compelling potential advantages, coupled with rapid technological progress in mid-infrared (MIR) laser development – where we define MIR to be in the range 2.5 - 10 µm – make this regime a compelling direction for the field to now explore.

1.3 Motivation for the Mid-Infrared

While near-infrared (NIR) lasers dominate in fields such as telecommunications [32], multiphoton microscopy [75, 76], and experiments requiring extreme intensities such as inertial confinement fusion and laboratory astrophysics [77, 78], MIR systems open a different set of opportunities. Light in the 2.5 - 10 µm band, known as the 'fingerprint' region, can directly excite fundamental vibrational modes of molecules, making the MIR uniquely suited for vibrational spectroscopy and molecular identification [79, 80]. This means that the MIR can be used across applications such as gas [81] and human breath analysis [82], environmental monitoring and industrial process control [83], as well as security and defence through explosive detection [84]. Strong absorption of water around 3 µm has also led to precise surgical techniques such as corneal and dental ablation, and continues to motivate ongoing research into dermatological, gynaecological, cardiovascular and neurological applications [85]. The MIR is also of interest for atmospheric studies, there being two transmission windows (3 - 5 µm and 8 - 12 µm) in the MIR, prompting great interest in its use for free-space optical communications and LIDAR [86]. Finally, the attractive scaling of normalised vector potential with increasing wavelength provides ample opportunities within the field of strong-field and attosecond science, such as extending high-harmonic generation (HHG) to ever shorter wavelengths [87].

Despite this abundance of potential applications, short-pulse MIR sources remain far less common than their NIR counterparts. Many of these applications, especially those requiring high peak powers, rely on femtosecond-scale pulse durations, but producing such pulses in the MIR remains a veritable challenge. Conventional gain media capable of lasing with sufficient bandwidth in this region are scarce, limiting direct approaches. While solid-state lasers have made substantial progress in recent years, much of the field has instead turned to nonlinear optical techniques for MIR generation and amplification.

This thesis presents the first technical report of the Chimera laser system, based at Imperial College London. Chimera is a multi-beam, high-power, short-pulse system built on optical parametric chirped pulse amplification (OPCPA) technology, with one arm operating in the mid-infrared (centred at 3.7 µm) generated via the nonlinear mixing process, difference frequency generation. Here, we document and characterise the system, and evaluate its feasibility for driving an *in situ* laser wakefield acceleration experiment at Imperial College London, particularly through detailed numerical simulations used to identify the key laser performance requirements.

1.4 Organisation of Thesis

This thesis aims to investigate the practical realities of commissioning a short-pulse, high-power mid-infrared laser to the point where it can routinely deliver light to an experiment, with particular emphasis on its application to driving a laser wakefield acceleration experiment. The remainder of this thesis is organised as follows:

Chapter 2 and **Chapter 3**

Will review the theoretical background relevant to the work presented. Chapter 2 will focus on laser physics, while Chapter 3 will introduce the theory of laser wakefield acceleration. Topics will include nonlinear optical processes central to the Chimera laser system, such as difference frequency generation and optical parametric amplification, before extending the treatment of nonlinear optics into underdense plasmas. This will culminate in a description of the laser wakefield acceleration mechanism.

Chapter 4

Will describe the methods used in this work, including diagnostics implemented on the Chimera laser system and the numerical approaches employed for wakefield acceleration simulations. Convergence tests supporting the choice of simulation resolution will also be presented.

Chapter 5

Will document the development and characterisation of the short-pulse, high-power, multi-wavelength OPCPA Chimera laser system at Imperial College London, with emphasis on its mid-infrared arm, centred at 3.7 µm. The chapter will also report results from its first commissioning shots in a mid-infrared-driven solid-target high harmonic generation experiment. A technical paper detailing the design and performance of the Chimera laser system is currently being prepared for publication.

Chapter 6

Will present the results of a density-scan simulation campaign, which revealed unexpectedly high accelerated electron energies. Potential explanations for these observations are investigated, including the influence of carrier envelope phase effects, direct laser acceleration, and beam loading. The work presented in this chapter is ongoing; however is expected to lead to a publication.

Chapter 7

Will discuss the parameter space required for the Chimera mid-infrared pulse to successfully drive a laser wakefield acceleration experiment. This will include the minimum performance requirements of the Chimera laser as well as estimates of the plasma length required for effective acceleration.

Chapter 8

Will summarise the work presented in this thesis. Conclusions from the preceding chapters will be revisited, and prospects for future work will be outlined.

2 Theory – Laser Technology

This chapter develops the theoretical background required to understand the operation of the Chimera laser system. We begin by introducing Maxwell's equations and their role in describing the propagation of electromagnetic waves in linear and nonlinear media. From this, we turn to nonlinear optical interactions, outlining the principles of second harmonic generation, difference frequency generation, and optical parametric amplification. We consider the temporal aspects of ultrashort pulses, beginning with the time-bandwidth product and the limits it places on pulse duration, before discussing chirped pulse amplification as a means of achieving high intensities while avoiding material damage. This leads to a treatment of pulse stretching and compression schemes, and their implementation within the framework of optical parametric chirped pulse amplification. We also examine the role of carrier envelope phase in few-cycle systems, and conclude with a discussion of optical damage mechanisms and the Kerr-lens modelocking process, which underpins the operation of femtosecond oscillators.

2.1 Nonlinear Light-Matter Interactions

In their macroscopic differential forms, Maxwell's equations describing all electromagnetic phenomena are given as:

$$\nabla \times \mathbf{E} = -\frac{\partial \mathbf{B}}{\partial t} \tag{2.1}$$

$$\nabla \times \mathbf{H} = \mathbf{J} - \frac{\partial \mathbf{D}}{\partial t} \tag{2.2}$$

$$\nabla \cdot \mathbf{D} = \rho \tag{2.3}$$

$$\nabla \cdot \mathbf{B} = 0 \tag{2.4}$$

with

$$\mathbf{D} = \epsilon_0 \mathbf{E} + \mathbf{P} \tag{2.5}$$

$$\mathbf{B} = \mu_0 \mathbf{H} + \mathbf{M} \tag{2.6}$$

Here we denote electric and magnetic fields as **E** and **M**, magnetic vector **H**, free current density **J**, electric displacement **D**, charge density ρ and material polarisation **P**. In dielectric materials, we have $\mathbf{D} = 0$ and in vacuum, $\mathbf{P} = 0$. Following the derivations presented in [88], we now review the fundamental nonlinear phenomena relating to high-intensity laser systems, along with specific cases that are relevant to the Chimera laser system.

Before going any further, we note that in optics, the word 'polarisation' is given two meanings. We may refer separately to the polarisation of light and the polarisation of matter. The former describes the direction of the field vectors in light (as a transverse electromagnetic wave), and the latter refers to the polarisation of the medium (i.e. the result of charge displacement by an imposed time-varying electric field). This second characterisation is also defined as the induced dipole moment per unit volume, $\mathbf{P}_i = \epsilon_0 \chi^{(i)} \mathbf{E}^i$. It is this definition that we will continue to work with in this section. We may expand polarisation $\mathbf{P}_i$ in terms of electric field amplitudes $\mathbf{E}^i$ as a power series:

$$\mathbf{P} = \epsilon_0 [\chi^{(1)} \mathbf{E} + \chi^{(2)} \mathbf{E}^2 + \chi^{(3)} \mathbf{E}^3 + ...] \tag{2.7}$$

Where ϵ_0 is the vacuum permittivity and the tensor coefficients $\chi^{(i)}$ represent the electric susceptibilities of the medium, with units $\mathrm{m}^{(i-1)}\mathrm{V}^{(1-i)}$ for all $\mathrm{i} > 1$. Note that, other than the lowest order linear $\chi^{(1)}$, these coefficients are typically small, and the terms they accompany only become significant with the imposition of strong electric fields such that the higher terms are then approximately comparable in magnitude to the linear response. The minimum electric field strength amplitude required to exhibit such a response is $\mathcal{O}(10^4\ \mathrm{Vcm}^{-1})$, corresponding to intensities of $\mathcal{O}(10^4\ \mathrm{Wcm}^{-2})$.

Whilst a series expansion is essentially a mathematical tool, the remarkable and observable power of nonlinear optics is that individual coefficients have meaningful and distinct consequences linked to specific physical phenomena. The first order susceptibility, $\chi^{(1)}$, describes linear optical effects and gives rise to the "ordinary" material refractive index, used for example to determine the behaviour of lenses and other fixed refractive optical elements. $\chi^{(2)}$ is used to describe second-order (typically two-photon) effects such as second harmonic generation (SHG), sum frequency generation (of which SHG is a specific form) and difference frequency generation (DFG). $\chi^{(3)}$ describes third-order effects such as third harmonic generation, four-wave mixing, the Kerr effect, and so on.

One reason we can attribute different behaviours to distinct coefficients is that these higher-order terms are diminishingly small except under specific conditions. These conditions depend not only on symmetry considerations but also on the temporal and spatial overlap (or walk-off) between the driving field and the generated field, which must maintain phase coherence for the new field to build up appreciably. An example of this can be demonstrated by selecting the second-order coefficient $\chi^{(2)}$, which is non-zero only

when the material in use possesses no inversion symmetry. This is true for all even terms: no centrosymmetric crystal may exhibit even-ordered nonlinear effects. We will focus particularly on the second-order coefficient $\chi^{(2)}$, as many of the nonlinear phenomena most relevant to the Chimera laser system are of the second order.

From Maxwell's equations (assuming a non-magnetic medium with no free charges), we take the curl of Equation (2.1) and the derivative with respect to time of Equation (2.2) and extract a wave equation:

$$\nabla^2\mathbf{E} - \epsilon_0\mu_0\frac{\delta^2\mathbf{E}}{\delta t^2} = \mu_0\frac{\delta^2\mathbf{P}}{\delta t^2} \tag{2.8}$$

where $c = 1/\sqrt{\epsilon_0\mu_0} = 3\times10^8$ ms^{-1} is the vacuum speed of all electromagnetic waves.

Ignoring the vectorial nature of the fields, we can simplify this expression for a plane wave propagating in the $+z$-direction to:

$$\frac{\delta^2 E}{\delta z^2} - \epsilon_0\mu_0\frac{\delta^2 E}{\delta t^2} = \mu_0\frac{\delta^2 P}{\delta t^2} \tag{2.9}$$

Assuming monochromaticity to simplify the mathematics, ensuring no time dependence of the respective complex amplitudes (in practice, we will also need to consider broadband, temporally confined pulses), we can solve this equation by taking the general case where we express both the electric field and polarisation waves as a complex sum of discrete frequency components:

$$\begin{aligned} E(z,t) &= \frac{1}{2}\sum_n \hat{E}_n \exp\{i\omega_n t\} + \text{c.c.} \qquad (n\in\mathbb{N}) \\ &= \frac{1}{2}\sum_n \tilde{E}_n \exp\{i(\omega_n t - k_n z)\} + \text{c.c.} \qquad (n\in\mathbb{N}) \end{aligned} \tag{2.10}$$

$$\begin{aligned} P(z,t) &= \frac{1}{2}\sum_n \hat{P}_n \exp\{i\omega_n t\} + \text{c.c.} \qquad (n\in\mathbb{N}) \\ &= \frac{1}{2}\sum_n \tilde{P}_n \exp\{i(\omega_n t - k_n z)\} + \text{c.c.} \qquad (n\in\mathbb{N}) \end{aligned} \tag{2.11}$$

Where c.c. refers to the respective complex conjugates. Inserting these expressions into Equation (2.9), and assuming a slowly varying envelope approximation to disregard the second spatial derivative of the field term, it follows that for the n^{th} component:

$$2ik_n\frac{\delta\tilde{E}_n}{\delta z} + \tilde{E}_n(k_n^2 - \epsilon_0\mu_0\omega_n^2) = \mu_0\omega_n^2\tilde{P}_n \tag{2.12}$$

When propagating in vacuum, $\mathbf{P}$, and therefore also $\tilde{P}$, is equal to zero and the equation becomes relatively simple to solve, simply by choosing a wavenumber which makes $\delta\tilde{E}/\delta z$ equal to zero. We therefore retrieve the familiar expression:

$$k_n = \sqrt{\epsilon_0\mu_0}\omega_n = \frac{\omega_n}{c} \tag{2.13}$$

When propagating through any material or medium, however, the polarisation term may not be set to 0. Instead we can choose to separate it into its respective linear and nonlinear contributions P^{L}, P^{NL}, such that in the case above, we have $\tilde{P} = \tilde{P}^{\mathrm{L}} + \tilde{P}^{\mathrm{NL}}$. At this point, we may simplify again by expressing the linear polarisation as the first order term of its Taylor expansion, giving $\tilde{P}_n^{\mathrm{L}} = \epsilon_0\chi^{(1)}\tilde{E}_n$, and Equation (2.12) becomes:

$$2ik_n\frac{\delta\tilde{E}_n}{\delta z} + \tilde{E}_n(k_n^2 - \epsilon_0\mu_0\omega_n^2(1+\chi^{(1)})) = \mu_0\omega_n^2\tilde{P}_n^{\mathrm{NL}} \tag{2.14}$$

In a linear approximation, $\tilde{P}^{\mathrm{NL}} = 0$ and we extract a new expression to describe the wavenumber of a wave passing through a material:

$$k_n = \sqrt{\epsilon_0\mu_0}\omega_n\sqrt{1+\chi^{(1)}} = \frac{\omega_n}{c}\sqrt{1+\chi^{(1)}} = \frac{n_n\omega_n}{c} \tag{2.15}$$

Where $\sqrt{1+\chi^{(1)}}$ is equal to the refractive index of light n_n travelling at frequency ω_n, and the susceptibility $\chi^{(1)}$ is defined as before. Using this new definition of the wavenumber k_n Equation (2.14) is reduced to:

$$\frac{\delta\tilde{E}_n}{\delta z} = -\frac{i\mu_0\omega_n}{2\sqrt{\epsilon_0\mu_0}n}\tilde{P}_n^{\mathrm{NL}} \tag{2.16}$$

$$= -\frac{i\mu_0 c\omega_n}{2n}\tilde{P}_n^{\mathrm{NL}} \tag{2.17}$$

Although simplified to be one-dimensional and to ignore any time dependence of the complex amplitudes (potentially a poor simplification when dealing with real – and short – laser pulses, as is the case for us), it is this differential equation that provides the basis for all further study of nonlinear optics.

While a detailed derivation will not be given in this thesis, the assumptions required to obtain the full wave equation for short pulses will be discussed. To begin, we choose to include only one frequency component ω, as we typically deal with signals very close to a single frequency. As such, we drop the subscript n from all physical variables. As

before, we separate the polarisation term into its linear and nonlinear contributions, $\tilde{P} = \tilde{P}^{\mathrm{L}} + \tilde{P}^{\mathrm{NL}}$. If we do not ignore the time dependence of the complex field and polarisation amplitudes, the wave equation given in Equation (2.9) becomes:

$$2ik\frac{\delta\tilde{E}}{\mathrm{d}z} + 2i\omega\epsilon_0\mu_0\frac{\delta\tilde{E}}{\mathrm{d}t} + \tilde{E}(k^2 - \epsilon_0\mu_0\omega^2) = \mu_0\left(\omega^2(\tilde{P}^{\mathrm{L}} + \tilde{P}^{\mathrm{NL}}) - 2i\omega\frac{\delta(\tilde{P}^{\mathrm{L}} + \tilde{P}^{\mathrm{NL}})}{\mathrm{d}t}\right) \tag{2.18}$$

Setting $\tilde{P}^{\mathrm{L}} = \epsilon_0\chi^{(1)}\tilde{E}$ and using $c = 1/\sqrt{\epsilon_0\mu_0}$ and $k = n_0\omega/c$, where $n_0 = \sqrt{1+\chi^{(1)}}$ is the low-intensity (linear) refractive index, we can simplify the above expression to:

$$\frac{\delta\tilde{E}}{\mathrm{d}z} + \frac{n_0}{c}\frac{\delta\tilde{E}}{\mathrm{d}t} = \frac{-i\mu_0\omega c}{2n_0}\left(1 - \frac{2i}{\omega}\frac{\delta}{\mathrm{d}t}\right)\tilde{P}^{\mathrm{NL}} \tag{2.19}$$

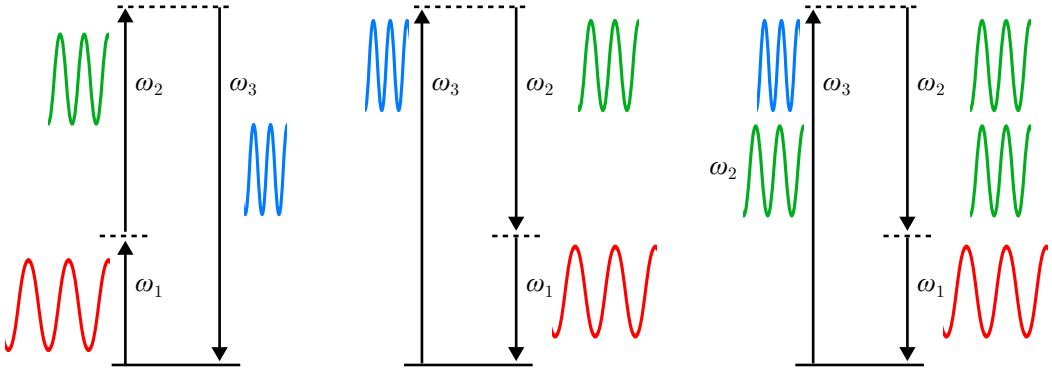
And identify a phase velocity v_{phase} as being equal to c/n_0. Happily, we can see that if we choose again to ignore any time dependencies, we can reduce this equation back to Equation (2.17).

Something which we have omitted from this derivation thus far is the fact that both the refractive index and the susceptibility are generally frequency dependent and complex, meaning that our previous relation between n_0 and $\chi^{(1)}$ requires a revision such that $v_{\mathrm{phase}} = c \,/\, n_0(\omega) = c \,/\sqrt{1+\chi^{(1)}(\omega)} = \omega/k$. The physical consequence of this frequency dependence is the well-known phenomenon of dispersion. We will not go into any further detail here, other than to define a group velocity $v_{\mathrm{group}} = \mathrm{d}\omega/\mathrm{d}k$ as the speed of the pulse envelope, note that v_{phase} describes the speed of the optical carrier wave and that to include dispersion in our equation, we often find it more convenient to switch to the frequency domain. The treatment of this will not be included here but can be found in [89].

2.2 Nonlinear Mixing

Here, we consider the basic frequency mixing processes of nonlinear optics. In optical parametric chirped pulse amplification (OPCPA) systems such as the Chimera laser, optical parametric amplification (OPA), a three-wave mixing process, is particularly relevant. Difference frequency generation (DFG), another three-wave mixing process, is also of the utmost importance, as it is this process that we exploit in order to first generate our mid-infrared seed pulse. The simplest example of a mixing process is second harmonic generation (SHG), whereby energy is transferred from the fundamental wave at ω to a wave at twice its frequency 2ω. This is a special case of the three-wave mixing process, sum frequency generation (SFG). Figure 2.1 shows a generalised photon picture of the energy exchange that takes place within each of these three-wave mixing processes.

In each case, the arrows pointing upwards depict a depleting (or pump) wave, and the arrows pointing downwards depict an enhanced one.



(a) Sum Frequency Generation (SFG)

(b) Difference Frequency Generation (DFG)

(c) Optical Parametric Amplification (OPA)

Figure 2.1: Energy level diagrams to describe SFG (a), DFG (b) and OPA (c) processes. Second harmonic generation (SHG) is a special case of SFG where $\omega_1 = \omega_2$. OPA is mathematically identical to DFG but is a seeded process.

To understand any mixing process with reference to the previously discussed wave equation, we need to find expressions for all relevant $\tilde{P}_n^{\mathrm{NL}}$. For the case of a three-wave process, we include $n = 1, 2, 3$, where the subscripts have no numerical significance other than to differentiate the frequencies. With the expansion for polarisation in Equation (2.7), we see that the second-order term comes in the form $\epsilon_0 \chi^{(2)} E^2$. Using Equation (2.10), we find the total field to be:

$$E(z,t) = \frac{1}{2}[\tilde{E}_1 \exp\{i(\omega_1 t - k_1 z)\} + \tilde{E}_1 \exp\{i(\omega_2 t - k_2 z)\} + \tilde{E}_3 \exp\{i(\omega_3 t - k_3 z)\} + \text{c.c.}] \tag{2.20}$$

Rather than beginning with the general three-frequency case, we will use second harmonic generation as an example. A special case of sum frequency generation, this is the most straightforward, and from this, we may extend the underpinning ideas to more complicated processes.

2.2.1 Second Harmonic Generation (SHG)

The simplest and best-known nonlinear optical phenomenon, second harmonic generation (SHG), also referred to as frequency doubling, is considered by many to be the 'second birth' of nonlinear optics, following the advent of the laser in 1960, which could finally provide the high field strength required to prompt a detectable nonlinear response. The first nonlinear experiment of the laser era was a demonstration of SHG by Peter Franken in 1961, who focused a ruby laser into crystalline quartz and produced a harmonic "as high

as a fraction of a percent of the fundamental" [90]. At this point, the conversion efficiency was not high enough for this result to be anything more than an exciting demonstration of the theoretical descriptions of multi-photon physics provided by Maria Göppert Mayer in 1931 [91]. Soon after Franken's result, however, the efficiency was greatly improved through the exploitation of material birefringence to 'phase match' the fundamental and harmonic waves and thus expand the field into practical applications [92]. Nowadays, with the phase matching problem solved and significant progress in reaching higher laser intensities, SHG is used widely, reaching efficiencies of over 50% and regularly exploited in applications ranging from optical microscopy [93] to shifting output wavelengths of lasers. Indeed, acting as a vehicle for a frequency shift is exactly how we make use of this process in the Chimera laser system.

To understand this process with reference to the wave equation, we need to find expressions for $\tilde{P}_1^{\mathrm{NL}}$ and $\tilde{P}_2^{\mathrm{NL}}$. Only these first two terms are necessary as we are considering only the terms $n = 1, 2$ for the respective fundamental and harmonic waves, ω_1 and ω_2. With the expansion for polarisation in Equation (2.7), we see that the second-order term comes in the form $\epsilon_0 \chi^{(2)} E^2$. Using Equation (2.10), we find the total field to be:

$$E(z,t) = \frac{1}{2}[\tilde{E}_1 \exp\{i(\omega_1 t - k_1 z)\} + \tilde{E}_1 \exp\{i(\omega_2 t - k_2 z)\} + \text{c.c.}] \tag{2.21}$$

Both $\tilde{P}^{\mathrm{NL}}$ terms can thus be found by evaluating the square of the sum of these fields. Renaming the fundamental frequency ω_1 as ω, and its second harmonic ω_2 as 2ω, each term in this expansion can be identified by its frequency order. This is shown in Table 2.1.

Field Term	Associated Frequency
$\tilde{E}_\omega^2 e^{i(2\omega t - 2k_\omega z)}$	2ω
$\tilde{E}_\omega^{*2} e^{i(-2\omega t + 2k_\omega z)}$	-2ω
$\tilde{E}_{2\omega}^2 e^{i(4\omega t - 2k_{2\omega} z)}$	4ω
$\tilde{E}_{2\omega}^{*2} e^{i(-4\omega t + 2k_{2\omega} z)}$	-4ω
$2\tilde{E}_\omega \tilde{E}_\omega^*$	DC
$2\tilde{E}_{2\omega} \tilde{E}_{2\omega}^*$	DC
$2\tilde{E}_\omega \tilde{E}_{2\omega} e^{i\left(3\omega t - (k_\omega + k_{2\omega})z\right)}$	3ω
$2\tilde{E}_\omega^* \tilde{E}_{2\omega}^* e^{i\left(-3\omega t + (k_\omega + k_{2\omega})z\right)}$	-3ω
$2\tilde{E}_\omega^* \tilde{E}_{2\omega} e^{i\left(\omega t - (k_{2\omega} - k_\omega)z\right)}$	ω
$2\tilde{E}_\omega \tilde{E}_{2\omega}^* e^{i\left(-\omega t + (k_{2\omega} - k_\omega)z\right)}$	-ω

Table 2.1: The square sum field expansion of the second order nonlinear polarisation, with terms identified by the resultant 'new' frequency.

We can then match the relevant frequency terms in Equation (2.7) with the terms in Table 2.1 to find:

$$\begin{aligned}\frac{1}{2}\tilde{P}_{\omega}^{\mathrm{NL}}e^{i(\omega t-k_{\omega}z)} &= \frac{1}{4}\epsilon_0\chi^{(2)}\cdot 2\tilde{E}_{\omega}^{*}\tilde{E}_{2\omega}e^{i\left(\omega t-(k_{2\omega}-k_{\omega})z\right)}\\ \tilde{P}_{\omega}^{\mathrm{NL}} &= \epsilon_0\chi^{(2)}\tilde{E}_{\omega}^{*}\tilde{E}_{2\omega}e^{-i(k_{2\omega}-2k_{\omega})z}\end{aligned} \tag{2.22}$$

And:

$$\begin{aligned}\frac{1}{2}\tilde{P}_{2\omega}^{\mathrm{NL}}e^{i(2\omega t-k_{2\omega}z)} &= \frac{1}{4}\epsilon_0\chi^{(2)}\cdot \tilde{E}_{\omega}\tilde{E}_{\omega}e^{i\left(2\omega t-2k_{\omega}z\right)}\\ \tilde{P}_{2\omega}^{\mathrm{NL}} &= \frac{1}{2}\epsilon_0\chi^{(2)}\tilde{E}_{\omega}\tilde{E}_{\omega}e^{i(k_{2\omega}-2k_{\omega})z}\end{aligned} \tag{2.23}$$

For the fundamental and second harmonic frequencies. In both expressions, we may identify a phase-mismatch parameter $k_{2\omega}-2k_{\omega}$. This can be defined as:

$$\begin{aligned}\Delta k &= k_{2\omega}-2k_{\omega}\\ &= \frac{2\omega n_{2\omega}}{c}-2\frac{\omega n_{\omega}}{c}\\ &= \frac{2\omega}{c}(n_{2\omega}-n_{\omega})\end{aligned} \tag{2.24}$$

Inspecting the above, we find that Δk is equal to zero (i.e. the relative phase between the nonlinear polarisation and harmonic field is negligible, and the interaction is at its most efficient) only when the refractive indices are equal. Given the nature of dispersion, it is highly unlikely that our refractive indices will be the same in a real optical material, and thus, the phase-mismatch will prevent efficient energy flow from the fundamental to the harmonic waves. The severity of the phase-mismatch can be quantified by defining a distance known as the coherence length L_{coh}, for which the field and polarisation terms become π radians out of step. There is no well-defined definition of this parameter, so we will define it here as it is presented in [88] as:

$$L_{\mathrm{coh}} = \frac{\pi}{|\Delta k|} \tag{2.25}$$

Mathematically, to phase match requires manipulating a system such that $\Delta k = 0$ (and therefore $n_{2\omega} = n_{\omega}$) throughout a significant fraction of the length of the crystal in which the interaction occurs – this length typically being many 10s of L_{coh}. This can be achieved in a number of ways, including the use of birefringent phase matching (BPM), a technique which exploits material birefringence in non-centrosymmetric, optically anisotropic crystals such as BBO (β-barium borate) and KTA (potassium titanyle arsenate). The refractive indices of birefringent materials are polarisation dependent – a

phenomenon first documented as early as 1669 when Danish scientist Erasmus Bartholini observed a double refraction of light incident on the calcite crystals he was studying in Iceland [94]. In practical terms, this means that if the fundamental and harmonic waves are in different polarisation states, they can be quite efficiently (to $> 50\%$) phase matched. Birefringent phase matching can be further categorised into two types (Type I and Type II) depending on the polarisation states of the interacting waves. Type I phase matching occurs when the two input photons are of the same polarisation, either perpendicular (ordinary) or parallel (extraordinary) to the optical axis, and combine to generate a photon in the orthogonal polarisation state. Type II phase matching uses input photons of orthogonal polarisation states: one ordinary and one extraordinary. This allows finer control over the effective refractive indices within the nonlinear crystal in use, such that the phase matching condition can be satisfied and optimised through angle tuning.

Another method to achieve efficient phase matching is to reverse the orientation of the non-centrosymmetric crystal at every coherence length, causing a sign reversal of the $\chi^{(2)}$. This sounds complicated, but it was proposed in 1962, soon after the first demonstration of BPM [95], with the first experimental fabrication and results coming in 1965 [96, 97]. The technique is now known as periodic poling and is used to 'quasi-phase match' (QPM) fundamental and harmonic waves. The difference in signal intensity when phase matching using both methods versus having a phase mismatch is shown in Figure 2.2.

We return now with our newly defined parameters to the wave equation in Equation (2.19), to produce a set of coupled wave equations which govern second harmonic generation. These are:

$$\frac{\delta \tilde{E}_\omega}{\mathrm{d}z} = \frac{-i\omega\chi^{(2)}}{2cn_\omega}\tilde{E}^*_\omega \tilde{E}_{2\omega} e^{-i\Delta k z} \tag{2.26}$$

$$\frac{\delta \tilde{E}_{2\omega}}{\mathrm{d}z} = \frac{-i2\omega\chi^{(2)}}{4cn_{2\omega}}\tilde{E}_\omega \tilde{E}_\omega e^{i\Delta k z} \tag{2.27}$$

These coupled differential equations can be solved by using the Manley-Rowe relations, which, although originally derived for nonlinear capacitor and inductor oscillations within an electrical circuit [98], very precisely describe energy, momentum and photon number conservation within interacting fields. These are found by considering intensities rather than fields. Using:

$$I_n = \frac{1}{2} c\epsilon_0 n_n |\tilde{E}_n|^2 \tag{2.28}$$

We may rewrite the above equations as:

$$\frac{\mathrm{d}I_\omega}{\mathrm{d}z} = \frac{i\omega\epsilon_0\chi^{(2)}}{4}\left[\tilde{E}_\omega^2\tilde{E}_{2\omega}^* e^{i\Delta kz} + \text{c.c.}\right] \tag{2.29}$$

$$\frac{\mathrm{d}I_{2\omega}}{\mathrm{d}z} = -\frac{i2\omega\epsilon_0\chi^{(2)}}{8}\left[\tilde{E}_\omega^2\tilde{E}_{2\omega}^* e^{i\Delta kz} + \text{c.c.}\right] \tag{2.30}$$

These are the Manley-Rowe relations and, for the case of SHG, can be rewritten in the following form:

$$\frac{\mathrm{d}}{\mathrm{d}z}\left(\frac{4I_\omega}{\omega}\right) = -\frac{\mathrm{d}}{\mathrm{d}z}\left(\frac{8I_{2\omega}}{2\omega}\right) \tag{2.31}$$

We leave the frequency terms in to highlight the role they play when they are not harmonics of one another but rather two (or more) distinct frequencies ω_1 and ω_2. For the more general case:

$$\frac{\mathrm{d}}{\mathrm{d}z}\left(\frac{I_1}{2\omega_1}\right) = -\frac{\mathrm{d}}{\mathrm{d}z}\left(\frac{I_2}{\omega_2}\right) \tag{2.32}$$

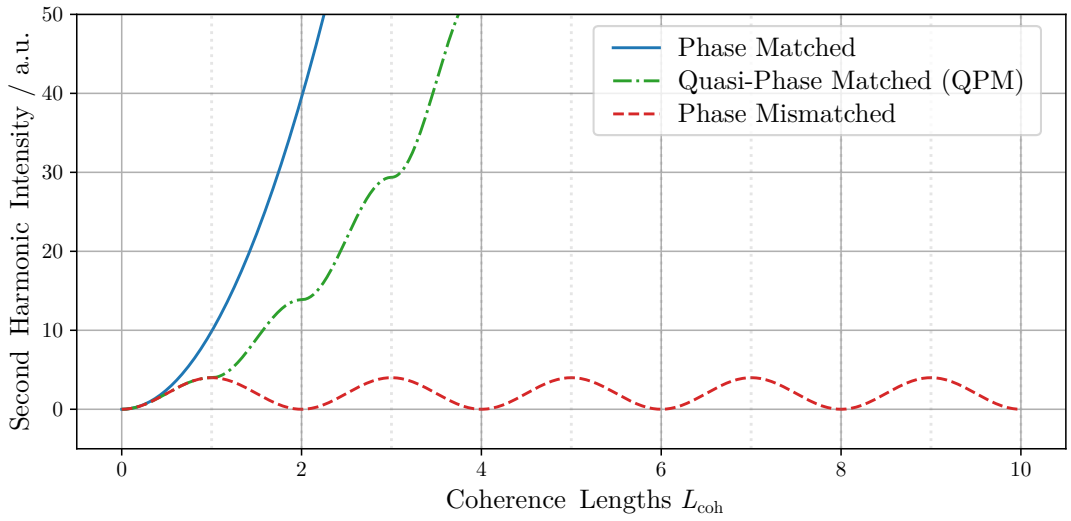


Figure 2.2: Second harmonic intensity signal as a function of coherence length. Birefringent phase matching (BPM) can be described as 'perfect' phase matching (although in reality it will be affected by crystal imperfections, temperature fluctuations and spatial walk off), and leads to higher intensities in second harmonic signal than for quasi-phase matching (QPM). The intensity growth of QPM initially follows the phase-mismatched curve. After reaching a single coherence length, the polarity is reversed, causing the field (and therefore intensity) to continue to grow instead of depleting through back-conversion. As the intensity goes as the square of the field, the intensity of the QPM signal after each n coherence lengths grows as the power of n.

We may also demonstrate this in a much simpler fashion by using the frequencies to think about the photon number picture, where we can say that for every photon gained at ω_2 (2ω), two photons are lost at ω_1 (ω). Using the conservation of energy, we can describe this with:

$$2\hbar\omega_1 \rightleftharpoons \hbar\omega_2 \tag{2.33}$$

Where $\hbar$ is Planck's reduced constant. Figure 2.1a shows this process schematically for the more general mixing process of sum frequency generation. The equation is written as reversible – this is to highlight how the process evolves as a result of the phase mismatch. After propagating over a single coherence length, the process starts to go in reverse, with the generated photons now being converted back to the fundamental pump field. This was a significant barrier to high efficiency processes, but is, these days, resolved by using either material birefringence or periodic poling to achieve quasi-phase matching. In the Chimera laser system, we use both of these methods. Our SHG and OPA stages are all phase-matched through material birefringence of angle-tuned BBO and KTA crystals. Our mid-infrared pulse is generated in a DFG stage using magnesium oxide-doped periodically poled lithium niobate (MgO:PPLN). In the photon picture, the phase-matching condition can also be described by the conservation of momentum:

$$2\hbar\vec{k}_1 = \hbar\vec{k}_2 \tag{2.34}$$

2.2.2 Difference Frequency Generation (DFG)

Difference frequency generation (DFG) is a second order, nonlinear, three-wave mixing process. When directed at a non-centrosymmetric crystal, an intense incoming beam, described as the 'pump', can be split into two lower-energy 'signal' and 'idler' output beams. As in the case of SHG, we aim to find a set of coupled wave equations to describe the process mathematically. Now, considering $n = 1, 2, 3$, we have a field of the form:

$$E(z,t) = \frac{1}{2}[\tilde{E}_1 \exp\{i(\omega_1 t - k_1 z)\} + \tilde{E}_1 \exp\{i(\omega_2 t - k_2 z)\} + \tilde{E}_3 \exp\{i(\omega_3 t - k_3 z)\} + \text{c.c.}] \tag{2.35}$$

We use this to solve for $\epsilon_0\chi^{(2)}E^2$, and find expressions for $\tilde{P}_1^{\text{NL}}$, $\tilde{P}_2^{\text{NL}}$ and $\tilde{P}_3^{\text{NL}}$. The maths is more involved for the case of three distinct frequencies, but the process remains the same. The nonlinear polarisation terms can be identified as $\tilde{P}_1^{\text{NL}} = \epsilon_0\chi^{(2)}\tilde{E}_3\tilde{E}_2^* e^{-i\Delta kz}$, $\tilde{P}_2^{\text{NL}} = \epsilon_0\chi^{(2)}\tilde{E}_3\tilde{E}_1^* e^{-i\Delta kz}$ and $\tilde{P}_3^{\text{NL}} = \epsilon_0\chi^{(2)}\tilde{E}_1\tilde{E}_2 e^{i\Delta kz}$, yielding three coupled differential equations:

$$\frac{\delta\tilde{E}_1}{\mathrm{d}z} = \frac{-i\omega_1\chi^{(2)}}{2cn_1}\tilde{E}_3\tilde{E}_2^* e^{-i\Delta kz} \tag{2.36}$$

$$\frac{\delta\tilde{E}_2}{\mathrm{d}z} = \frac{-i\omega_2\chi^{(2)}}{2cn_2}\tilde{E}_3\tilde{E}_1^* e^{-i\Delta kz} \tag{2.37}$$

$$\frac{\delta\tilde{E}_3}{\mathrm{d}z} = \frac{-i\omega_3\chi^{(2)}}{2cn_3}\tilde{E}_1\tilde{E}_2 e^{i\Delta kz} \tag{2.38}$$

Where this time, the phase mismatch is $\Delta k = k_3 - k_2 - k_1$ (and $k_3 > k_2 > k_1$). The Manley-Rowe relations describing the rate of creation and destruction of photons can be found and turned once more into a statement of energy conservation:

$$-\frac{\mathrm{d}}{\mathrm{d}z}\left(\frac{I_3}{\omega_3}\right) = \frac{\mathrm{d}}{\mathrm{d}z}\left(\frac{I_2}{\omega_2}\right) = \frac{\mathrm{d}}{\mathrm{d}z}\left(\frac{I_1}{\omega_1}\right) \tag{2.39}$$

This can again be described more simply using the photon picture, where the photon frequencies ω are described by:

$$\hbar\omega_\mathrm{p} \rightleftharpoons \hbar\omega_\mathrm{s} + \hbar\omega_\mathrm{i} \tag{2.40}$$

Here we adjust the subscripts to refer to pump $3 \to \mathrm{p}$, signal $2 \to \mathrm{s}$ and idler $1 \to \mathrm{i}$, and $\hbar$ is Planck's reduced constant. Note that $\omega_\mathrm{p} > \omega_\mathrm{s} > \omega_\mathrm{i}$. This relationship tells us that for every photon created with energy $\hbar\omega_\mathrm{s}$, a photon is also created at $\hbar\omega_\mathrm{i}$, while energy at $\hbar\omega_\mathrm{p}$ is depleted. This is shown schematically in Figure 2.1b.

This picture is an interesting one, as we might immediately ask how the photon with energy $\hbar\omega_\mathrm{p}$ knows which frequencies it ought to split into. The answer, of course, lies in where the most efficient phase matching may take place – the phase matching (momentum conservation) condition must also be satisfied:

$$\hbar\vec{k}_\mathrm{p} = \hbar\vec{k}_\mathrm{s} + \hbar\vec{k}_\mathrm{i} \tag{2.41}$$

Where $\vec{k}$ is used to denote the wavevector of the beams. Despite our choice of frequency (and therefore wavelength) being limited by where we may find efficient phase matching, this method opens the door not just to the possibility of more exotic 'off-harmonic' wavelength generation, but to the possibility of using this process whilst maintaining vast amounts of bandwidth, particularly compared to classic gain storage laser materials. In the Chimera DFG stage, we can create and amplify over 1 µm of bandwidth in the mid-infrared, allowing for later compression potentially into the few-cycle regime.

2.2.3 Optical Parametric Amplification (OPA)

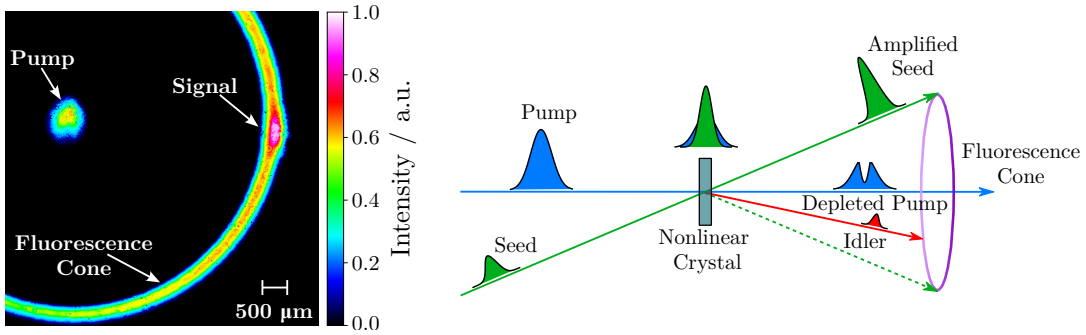
Optical parametric amplification (OPA) is based on the second order nonlinear optical process DFG. Mathematically, OPA and DFG are the same: an intense incoming pump beam with frequency ω_p separates itself into two lower energy output beams. The difference between these processes is in the initial conditions. OPA is necessarily a seeded process, whereas DFG requires only one incident beam [88, 99]. Figure 2.1 compares the energy level diagrams of the two processes.

Parametric amplifiers have several advantages over classic gain storage laser amplifiers. They can provide a broad bandwidth, determined not by the laser transition but by a combination of phase matching and transparency of the nonlinear crystal, and are not limited to narrow wavelength ranges. They are thus seen not just as an attractive way to amplify broad bandwidths of 'unusual' wavelengths of light (such as within the mid-infrared) but also as a way to avoid gain narrowing within an amplifying medium [100]. β-barium borate (BBO) and lithium triborate (LBO) are nonlinear crystals that are frequently used in OPA processes as they not only exhibit useful optical qualities such as high $\chi^{(2)}$ and low group velocity dispersion but also have very high damage thresholds [99]. Another common crystal is potassium titanyl arsenate (KTA). KTA is a more specialised case: it is transparent between 0.4 - 5 µm with optimal performance at 3.5 µm, hence providing an ultra-broad bandwidth, and it also exhibits a very high damage threshold, making it well suited for high power MIR laser amplification [101].

Despite producing a much higher gain per unit length than classic laser amplifiers, an OPA does not typically generate much heat. This is because the process has an extremely low quantum defect: the energy of each pump photon is exactly split into a signal and an idler photon, with no excess energy left behind in the medium. As a result, there is virtually no net energy transfer to the crystal, and heating arises only from minor parasitic absorption rather than from the amplification mechanism itself. This means that they can be more easily used within high-repetition-rate, high-power systems without the necessity of external cooling mechanisms. The accumulation of high-order spectral phase terms (often resulting from material dispersion) is also reduced when using OPAs, as they are generally thin (millimetre to few-centimetre scale) compared to standard gain storage amplifiers, and typically are used only with a single pass. Minimising these higher-order terms is something which is of great importance when it comes to the compressibility of the pulse. OPAs can produce high pulse energies while maintaining the ability to generate a signal with a shorter duration than the pump, and have the capacity to produce very high contrast, high-quality pulses. This is because the amplifier has no energy storage: it is only active while being pumped in a strictly limited (usually picosecond - nanosecond) temporal window. This also means that pre-pulses, accumulated in earlier parts of the system, are less likely to be amplified. This is a particularly attractive quality in a system wishing to drive high-intensity experiments, as it reduces the probability of a pre-pulse reaching a target or damaging a fragile structure. The disadvantages compared to more standard gain amplification methods are that a

complex high-quality pump system (essentially an entire additional laser) as well as very precise synchronisation of the pump and seed pulses are required (to within the temporal window of the pump) and that these crystals have to be grown with sufficient quality, meaning that their size and therefore aperture is more limited [102].

While OPAs have no long timescale spontaneous emission (usually the cause of a high noise floor of standard gain amplifiers), they do parametrically fluoresce. Analogous to the amplified spontaneous emission (ASE) in storage gain amplifiers, this fluorescence is a result of spontaneous down conversion of the signal and idler beams when the OPA is pumped at high intensity and can be a significant source of incoherent noise [103, 104]. This can be considered as an interaction of the pump field with a background of vacuum fluctuation virtual photons, and is particularly problematic as it only occurs in the direction of phase matching – it is emitted as a solid cone of angles which represent the phase-matching angles and frequencies surrounding the pump. The result of this is that it only occurs in the direction of the real signal pulse, meaning that it can produce a significant (incoherent) pedestal around the amplified signal that can be detrimental to its compressibility. Figure 2.3 shows a real example of parametric fluorescence seen on the 800 nm OPA stage of the Chimera laser system. As this is an effect that worsens with increasing gain, the maximum gain that an OPA stage can therefore achieve is limited both by parametric fluorescence and by the bandwidth that it may support.



(a) (b)

Figure 2.3: Parametric fluorescence as a result of spontaneous down conversion of the signal and idler beams in an OPA.
(a) A spatial profile, taken using a WinCamD-LCM – 1" CMOS Beam Profiler, showing a real example of parametric fluorescence seen at the output of the 800 nm OPA stage on the Chimera laser system. The pump has been filtered by a factor of 10^6 relative to the signal to prevent saturation of the camera.
(b) A cartoon to depict the process of parametric fluorescence. The process occurs because of the finite probability that the pump will decay into signal and idler pairs across the solid cone of phase-matching angles within the amplifier.

There are five (one double-passed) OPA stages currently in use on the Chimera laser system: one of these uses a BBO crystal and the other four use KTA crystals – it is the KTA which amplifies the MIR light. The choice of crystal here is not arbitrary. Different crystals have properties suitable for different wavelengths, pulse durations and bandwidths. The length and size of the crystal are also important, as is the angle at

which the incoming laser pulse is incident. These factors may be modelled, along with other important considerations such as pump depletion effects, by solving the coupled wave equations derived here.

2.3 Time-Bandwidth Product (TBP)

When designing a system to generate an ultrashort pulse, one is always constrained by a fundamental limit known as the time-bandwidth product (TBP). This limit has its origins in the uncertainty principle, but rather than describing fundamental information that can be known about a system, it instead describes a relationship between time (temporal width Δt) and a spread of frequencies (bandwidth $\Delta\nu$) that may be present within an electromagnetic wave. Here we define our temporal width Δt and bandwidth $\Delta\nu$ specifically to mean the full-width at half-maximum (FWHM) value of the temporal and spectral intensity profiles. See Appendix A.2 for more details of this definition. The relationship states that if you require a short pulse of light, you must also have sufficient bandwidth to support it. Mathematically, this is described as:

$$\Delta\nu\Delta t = \text{TBP} \geq \frac{h}{4\pi} \tag{2.42}$$

Where h is Planck's constant. The numerical limit to the time-bandwidth product is dependent on the detailed shape of the pulse it describes. For example, a Gaussian pulse has a transform limit TBP $\simeq 0.44$, while a sech^2 pulse has a transform limit TBP $\simeq 0.315$. The transform limit is a lower bound; it tells you how short you can potentially make the pulse; however, in practice, it is not possible to be perfectly transform-limited. It is possible to make the pulse longer than the transform limit by stretching it out in time: this is what chirping does – the amount of frequency chirp or phase structure within a pulse will also contribute to the time-bandwidth product.

2.4 Chirped Pulse Amplification (CPA)

Inspired by a technique used to increase operational power in radar, chirped pulse amplification, also referred to as CPA, was suggested and demonstrated in the optical regime by Donna Strickland and Gérard Mourou in 1985, earning them the 2018 Nobel Prize in physics [25]. Before this method became standard practice, the powers accessible by direct laser amplification were limited by intense laser-induced damage to optical components, particularly via the Kerr effect and self-focusing. The only way to combat this within systems wishing to deliver a high-power pulse was to introduce increasingly large optics to spread out the beam spatially; however, this was not always practical (or cheap!), especially as it became desirable to push laser systems to higher and higher intensities and energies. To use CPA in a system, one must first send in a short 'seed' pulse, typically from a modelocked oscillator. To amplify these pulses safely, they must be stretched (chirped) in time, with factors typical of order 10^3-10^5, subsequently amplified to energies required for experimental operation and then recompressed (negatively

chirped) in time to as close as possible to a pulse duration determined by its transform limit (see Section 2.3). The CPA method is illustrated in Figure 2.4. Now supporting the amplification and compression of lasers well into the petawatt regime, CPA has revolutionised the field of lasers and optics, it being the technology that underpins virtually all high-power laser systems [36].

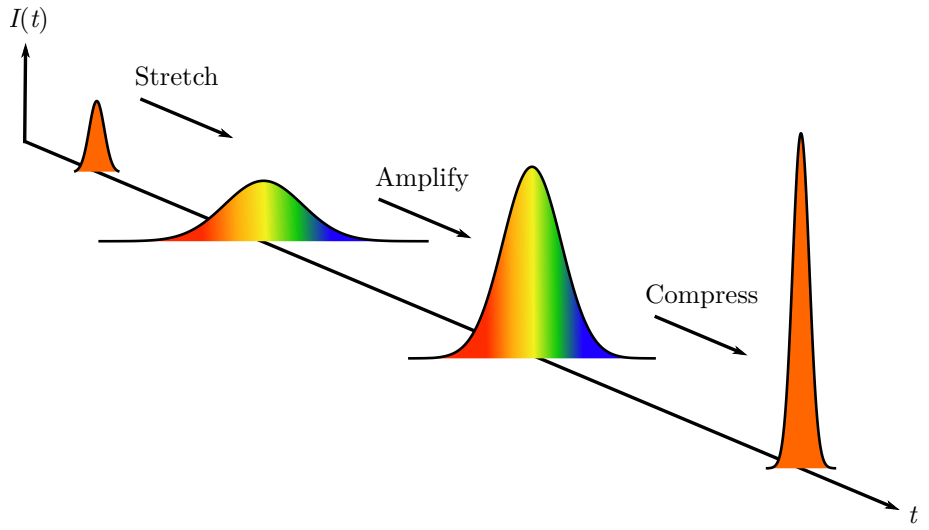


Figure 2.4: A cartoon showing the concept of chirped pulse amplification (CPA). As the pulse moves in time, it is stretched by temporally delaying different frequency components, amplified and recompressed so that a high-intensity pulse can interact with a target without causing any optical damage to a system. The rainbow fill is used to illustrate the process of chirping the pulse – its wavelength components have been dispersed in time.

In Section 2.1, we considered the relationship between linear susceptibility and refractive index using the equation $n(\omega) = \sqrt{1 + \chi^{(1)}(\omega)}$. At the time, we omitted the fact that both the linear susceptibility and refractive index are frequency dependent – a dependence that gives rise to the phenomenon of dispersion. When we account for this frequency dependence, we find that any change in the refractive index with frequency leads to a corresponding change in the phase velocity, given by $v_{\text{phase}} = c \;/\; n(\omega) = c \;/\sqrt{1 + \chi^{(1)}(\omega)} = \omega/k(\omega)$. This effect is known as chromatic (or phase velocity) dispersion, which may be classified as 'normal' when $\mathrm{d}n(\omega)/\mathrm{d}\omega > 0$ or 'anomalous' when $\mathrm{d}n(\omega)/\mathrm{d}\omega < 0$. A second velocity, the group velocity, can also be derived from this relationship: $v_{\text{group}} = \mathrm{d}\omega/\mathrm{d}k$. These two velocities are important in describing a propagating pulse $E(z,t)$: the high-frequency optical oscillations move at the phase velocity, while the envelope of the pulse travels at the group velocity. If the pulse envelope broadens and exhibits a frequency sweep (i.e. there is a change in the instantaneous frequency across the duration of the pulse), the pulse is said to be chirped. In a non-dispersive medium, the group and phase velocities are identical.

To chirp an optical pulse requires us to instantaneously force shorter wavelength components to travel over a different distance to the longer wavelength components, such that different pulse frequencies then occur at different times. Mathematically, this can be described by introducing a frequency-dependent phase $\phi(\omega)$, given by $\phi(\omega) = k(\omega)z =$

$n(\omega)\omega z/c$, where z is the length of the material through which the pulse propagates. This phase variation with frequency leads to temporal and spectral reshaping of the pulse, resulting in a chirp. If, in the spectral domain, we expand the frequency-dependent phase $\phi(\omega)$ using a Taylor series, we find:

$$\phi(\omega) \equiv -k(\omega)z = \phi_0 + \frac{\mathrm{d}\phi}{\mathrm{d}\omega}\Big|_0(\omega - \omega_0) + \frac{1}{2}\frac{\mathrm{d}^2\phi}{\mathrm{d}\omega^2}\Big|_0(\omega - \omega_0)^2 + \dots \tag{2.43}$$

We can identify each order in the phase expansion with a specific physical effect. The zeroth order term, $\phi_0 = -k(\omega_0)z$, represents the absolute phase while the first order term $\mathrm{d}\phi/\mathrm{d}\omega\Big|_0$ corresponds to the group delay, $-\tau_{\text{group}} = \phi_1$. Together, these determine the carrier envelope phase (CEP) and the arrival time of the pulse envelope. The second order term $\mathrm{d}^2\phi/\mathrm{d}\omega^2\Big|_0$ represents the group delay dispersion (GDD), defined as $-\mathrm{d}\tau_{\text{group}}/\mathrm{d}\omega = \phi_2$. GDD quantifies the spread of group delays across the pulse's bandwidth and introduces a quadratic phase across the spectrum. This leads to a linear frequency sweep in the time domain – i.e., a linearly chirped pulse, characteristic of the stretched pulses most commonly used in CPA systems. Higher-order terms (e.g. third- and fourth-order dispersion) become increasingly significant for ultrashort or high-intensity pulses and can lead to asymmetric broadening or temporal distortion if not properly managed. Each term has a distinctive effect upon the shape of the pulse in time.

Other than a spectral dispersion in time, chirping does not detrimentally affect the shape of the pulse envelope after compression – its spectral content is not lost. It is for this reason that the second-order phase manipulation is used in the CPA process. We can illustrate this by using an example of a real chirp we impose on our pulse in the Chimera laser system (using idealised rather than specifically measured parameters). Practically, we can chirp a pulse by introducing GDD. This can be achieved using material dispersion through bulk media (e.g. fused silica or glass), or dispersion-engineered optics such as chirped mirrors, diffraction gratings, or prism pairs. A positively (negatively) or up (down) chirped pulse is one whose frequencies increase (decrease) with time. When our 800 nm pulse exits our oscillator, it has a transform-limited FWHM pulse duration of 7 fs before interacting with a pair of chirped mirrors. These are a Thorlabs DCMP175 - dispersion-compensating mirror set, with a dispersion per reflection of $\phi'' = -175$ fs^2. The pulse completes 6 bounces on this chirped mirror pair, accumulating a total GDD of -1050 fs^2. The chirped pulse duration can be calculated from its original duration as well as the ϕ'', using [105]:

$$t_2 = t_1\sqrt{1 + \left(\frac{4\ln(2)\phi''}{t_1^2}\right)^2} \tag{2.44}$$

To 3 significant figures, this gives a stretched pulse duration of 416 fs. Figure 2.5 shows the effect of one bounce on the oscillator pulse (so a stretch to 69.7 fs). The wavelength has been scaled by a factor of 5 for visual clarity of the chirp.

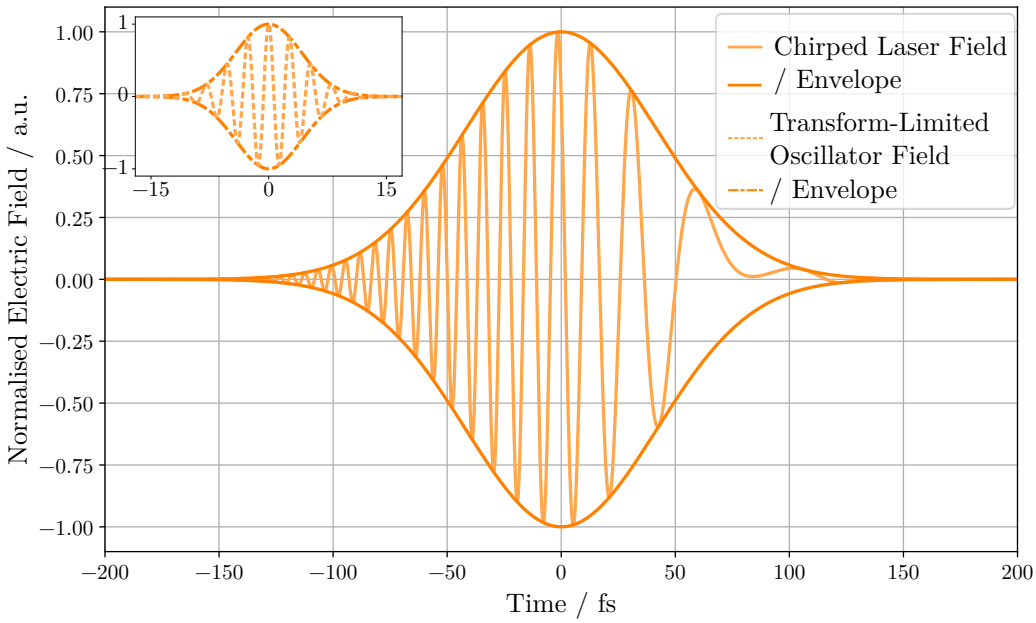


Figure 2.5: Electric field of a chirped transform-limited pulse. The parameters are those of the transform limit of the Chimera oscillator: 7 fs at a centre wavelength of 800 nm, with $\sim$ 135 nm of bandwidth. The chirp is modelled for the oscillator pulse performing a single bounce off a chirped mirror. This provides the addition of $\phi'' = -175$ fs^2 second order phase, resulting in a stretch to 69.7 fs. The inset shows the original 7 fs transform-limited pulse (note the change in scale on the x-axis). To improve the visual clarity, the wavelength of the chirped laser field has been scaled by a factor of 5. Notice that despite the pulse being stretched in time, the pulse envelope has not been adversely affected.

Higher-order phase terms, like third and fourth order dispersion (TOD and FOD), typically do have a detrimental effect upon the pulse envelope, even when only present in very small amounts. The effects of higher-order dispersion on the temporal pulse intensity are shown in Figure 2.6. When designing large OPCPA systems such as Chimera, it is very important to statically minimise these terms as much as is feasible. By chirping an optical pulse before amplifying and recompressing it, CPA allows the propagation of high-intensity, ultrashort, pulsed laser light through a system, whilst avoiding optical damage or pulse distortions caused by nonlinear processes.

As the refractive index of optical materials varies with wavelength, all such chirping methods inevitably also introduce higher-order dispersion – particularly TOD and FOD. These higher-order phase terms become increasingly problematic in high-bandwidth systems aiming to compress pulses to just a few optical cycles, as they limit the achievable temporal compression. To quantify this effect for a given material, we use the well-known Sellmeier equation, which models the wavelength dependence of the refractive index of a simple optical material (one without absorption or resonance features) as:

$$n(\lambda) = \left(1 + \frac{B_1\lambda^2}{\lambda^2 - C_1} + \frac{B_2\lambda^2}{\lambda^2 - C_2} + \frac{B_3\lambda^2}{\lambda^2 - C_3}\right)^{1/2} \tag{2.45}$$

Where $B_{1,2,3}$, $C_{1,2,3}$ are empirically determined coefficients (typically supplied by the material manufacturers). Once the refractive index as a function of wavelength is known, dispersion parameters such as GDD and higher-order dispersion can be calculated by taking derivatives of $n(\lambda)$ with respect to wavelength or frequency. This allows us to determine both the material lengths required to achieve a desired chirp and the extent of higher-order phase that will accumulate. Despite the simplicity and robustness of transmissive bulk media from an alignment standpoint, they tend to introduce significant higher-order dispersion – particularly third-order phase-per-millimetre. This limits their usefulness for stretching pulses beyond a few picoseconds. In high-power systems requiring greater stretch factors, stretcher and compressor grating pairs are more typically used, usually combined with a broadband reflective imaging system to create a stretcher.

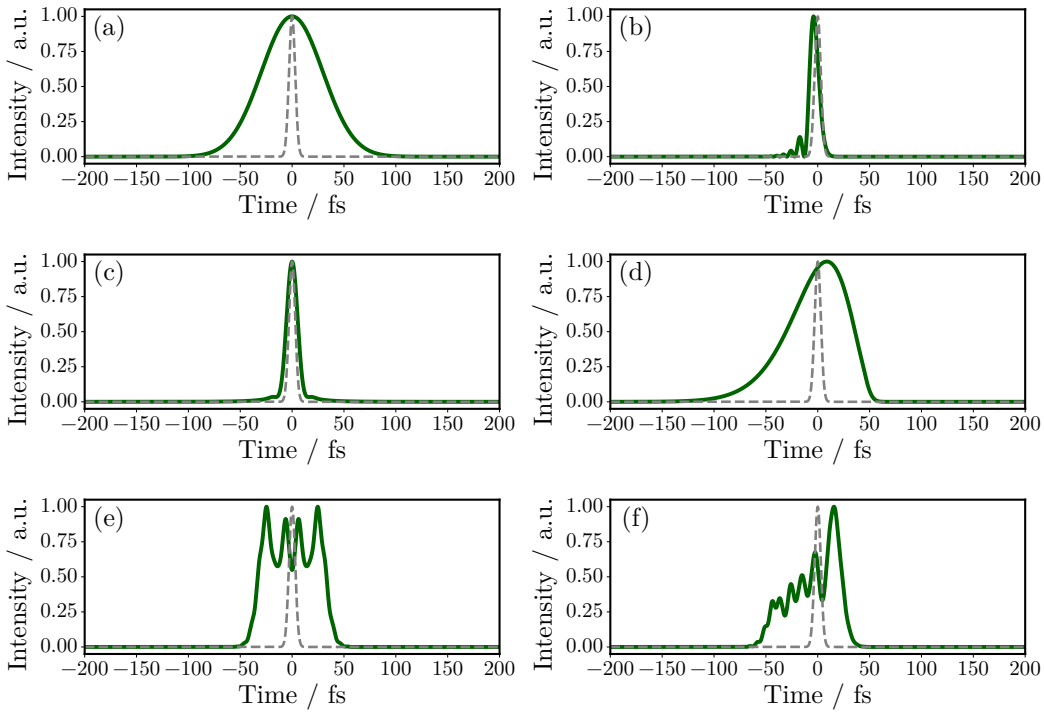


Figure 2.6: The effects of second ($\phi'' = -175$ fs^2), third ($\phi''' = -300$ fs^3) and forth ($\phi'''' = -4000$ fs^4) order spectral phase on a 7 fs, 800 nm, transform-limited laser pulse. The grey dashed curves show the original transform-limited pulse intensity profile. The green curves show the intensity profiles after the phase terms have been applied. The respective phase terms are abbreviated as group delay dispersion (GDD), third-order phase (TOD) and fourth-order dispersion (FOD).
(a) GDD only. The effect on the pulse envelope is a uniform broadening. (b) TOD only. TOD introduces a side lobe to the pulse – this can materialise in pre- or post-pulses. (c) FOD only. The addition of FOD produces a pedestal on the pulse, reducing the contrast. (d) GDD + TOD. (e) GDD + FOD. (f) GDD + TOD + FOD.

2.4.1 Stretching and Compressing a High Intensity Pulse

For high-intensity laser systems aiming to compress pulses to durations of just a few femtoseconds, diffraction gratings are typically used to stretch and subsequently compress the pulse. Commonly used imaging configurations in the stretcher include Treacy [106], Öffner [107], and Martinez [108]. For a reflective grating, the condition for constructive interference is given by the grating equation:

$$m\lambda = d(\sin\theta \; + \; \sin\gamma) \tag{2.46}$$

Where m is the diffraction order, λ is the wavelength, d is the groove spacing (inverse of the groove density), θ is the angle of incidence, and γ is the diffraction angle, both from the grating normal. The zeroth-order term corresponds to specular (non-dispersive) reflection and thus represents a loss mechanism in practical CPA systems.

Shown in Figure 2.7, the simplest stretcher or compressor geometry introduces GDD using a pair of parallel diffraction gratings. In this setup, the first grating spectrally disperses the pulse, introducing a frequency-dependent time delay, while the second grating recollimates the beam. This configuration introduces a negative chirp to the pulse – longer wavelengths travel a greater optical path length than shorter wavelengths, resulting in a temporally stretched pulse.

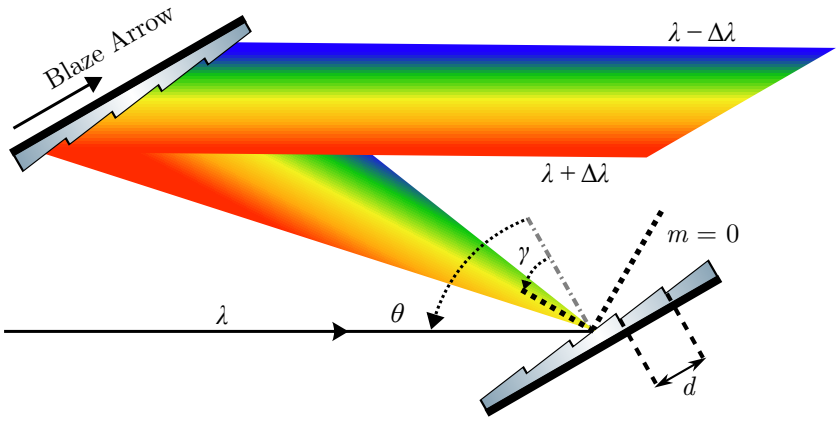


Figure 2.7: Optical layout of a simple two-parallel-grating system used to introduce negative GDD. The rainbow fill illustrates the frequency-dependent time delay introduced to the pulse. d is the groove spacing and m denotes the diffraction order – $m = 0$ corresponds to the non-dispersive (zeroth-order) reflection. The angles of incidence (θ) and diffraction (γ) are measured with respect to the grating normal. These parameters are used in the grating equation to determine the condition for constructive interference. The blaze arrow refers to the orientation of the blaze angle on a diffraction grating, specifically the direction in which the grating grooves are tilted to maximise diffraction efficiency for a particular wavelength and diffraction order.

Also shown in the figure is a 'blaze angle' – a feature often used in reflective gratings. The groove profile is etched, or mechanically machined, into the substrate of the grating.

This is then coated with a high-reflectivity surface. The blaze angle refers to a deliberate asymmetry in the groove shape, designed to maximise diffraction efficiency for a specific wavelength and diffraction order. This 'blazing' is wavelength-dependent and is usually optimised so that the majority of the incident light is steered preferentially into one order (usually the first), enhancing efficiency. Blazing also helps to reduce spectral overlap by increasing angular separation from neighbouring diffraction orders. The arrow points in the direction of the facet tilt (steep side) of the grooves.

Although Figure 2.7 depicts a scenario that does introduce a temporal chirp, it also produces a spatial one. This is something that can be problematic in a CPA system unless compensated for in some way. In real laser systems, this can be mitigated by double-passing the setup – retroreflecting the beam with a near-normal incidence mirror or roof prism placed at the spatially chirped output, or in a Fourier plane. To spatially separate the input and output beams, a tilt can be introduced in the non-dispersive plane. When well-aligned, this geometry can compensate for the spatial chirp it introduces, while doubling the temporal chirp, although this necessarily comes at the expense of additional energy loss due to extra surface reflections. With the addition of a double pass, Figure 2.7 describes a typical Tracey layout [106]. This configuration is a very attractive one as it involves no material transmission and supports large beam diameters, helping to avoid optical damage.

With an effective grating separation G, the pulse duration per pass of a diffraction grating stretcher/compressor can be calculated analytically as [88]:

$$\Delta t = \frac{G\lambda_0\Delta\lambda}{c\ (\ d\cos\left(\gamma\lambda_0\right)\ )^2} \tag{2.47}$$

Where λ_0 is the central wavelength of a pulse with bandwidth $\Delta\lambda$. Here, we have described one specific scheme to impose a negative GDD. Other grating geometries, such as Martinez and Öffner, produce a positive chirp by inserting a telescope between the two gratings. Many laser systems use Öffner or Martinez configurations to stretch the pulse, followed by a Treacy compressor for recompression, as this both eliminates the need for an imaging system in the compression stage, and lends itself to larger beam sizes, minimal unwanted material dispersion and Kerr effects.

2.5 Optical Parametric Chirped Pulse Amplification (OPCPA)

Optical parametric chirped pulse amplification (OPCPA) refers to the use of a specific type of amplifier in chirped pulse amplification. Conceptually, they are the same as 'simple' CPA systems; however, they use a different method of amplification – an OPA is used in place of the standard energy storage gain medium within the amplification chain. Amongst the advantages of OPA systems is that using a chirped seed within the OPA

process provides a further degree of freedom with regard to the spectral management of the pulse.

In Section 2.2.2, we claimed that nonlinear mixing processes such as OPA can preserve the broad bandwidth of a seed pulse, thereby enabling compression to ultrashort durations. What we omitted from this claim was that this capability is largely made possible by OPCPA technology. Previously, we highlighted that to achieve the phase-matching condition $\Delta k = 0$ (whereby, for efficient energy transfer, phase matching is required between the interacting pump, signal, and idler waves), we could simply exploit either material birefringence or turn to periodic poling. While this remains true, such approaches inherently provide phase matching only over a limited spectral bandwidth, since phase matching is wavelength dependent. As a result, only a narrow portion of the pulse spectrum is efficiently amplified. To understand how we can achieve a broad phase-matching plateau for all wavelengths across the bandwidth, we expand the wavevector mismatch $\Delta k(\omega)$ as a Taylor series about the central frequency ω_0:

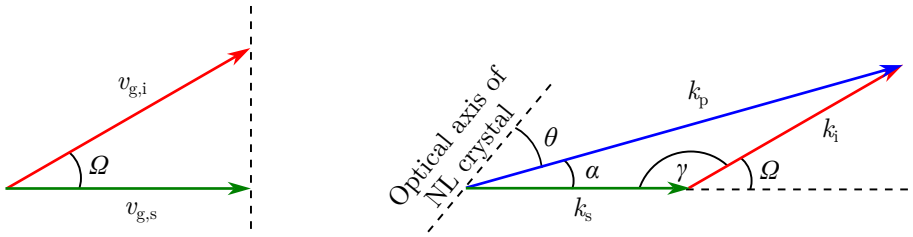
$$\Delta k(\omega) = \Delta k_0 + \frac{\mathrm{d}\Delta k}{\mathrm{d}\omega}\Big|_0(\omega - \omega_0) + \frac{1}{2}\frac{\mathrm{d}^2\Delta k}{\mathrm{d}\omega^2}\Big|_0(\omega - \omega_0)^2 + \dots \tag{2.48}$$

For narrowband pulses where $\omega \ll 1$, it is sufficient to satisfy the phase-matching condition at the central frequency, i.e., $\Delta k_0 = 0$. However, for broadband pulses around a central frequency ω_0, this condition alone is not sufficient. In this case, the phase mismatch must remain small over a wide frequency range, which additionally requires that the first-order frequency dependence of the phase mismatch vanishes, i.e., $\mathrm{d}\Delta k/\mathrm{d}\omega = 0$. This condition is equivalent to group-velocity matching.

While phase matching ensures the conservation of momentum between the interacting waves and is determined by their phase velocities, it does not imply that the pulse envelopes propagate at the same speed through the crystal: in dispersive media, the pump, signal, and idler generally have different group velocities. As a result, even when perfect phase matching is achieved at a single frequency, different spectral components of a broadband pulse walk off from each other during propagation, limiting the effective interaction length and therefore the achievable gain bandwidth. In a collinear phase-matching geometry, the pump and signal propagate along the same direction, and phase matching is controlled by a single angular degree of freedom (the orientation of the crystal optical axis). In this case, it is generally not possible to simultaneously satisfy both $\Delta k_0 = 0$ and $\mathrm{d}\Delta k/\mathrm{d}\omega = 0$, resulting in a limited gain bandwidth.

To overcome this limitation, OPCPA systems commonly introduce a non-collinear geometry between the pump and seed pulses, meaning that the phase matching conditions are determined by two angular degrees of freedom rather than one. By introducing a non-collinear angle α between the pump and seed beams, we gain an additional degree of freedom, allowing independent control of the wavevector components parallel and perpendicular to the pump direction. By choosing an appropriate α, the projection of

the signal group velocity along the pump direction can be matched to the pump group velocity. This minimises temporal walk-off between the pulses and ensures that the first-order frequency dependence of the phase mismatch vanishes. Consequently, phase matching ($\Delta k = 0$) is maintained over a much broader spectral range, enabling more efficient amplification of ultra-broadband pulses. This geometry is shown in Fig. 2.8 and is referred to as non-collinear phase matching.



(a) (b)

Figure 2.8: Configuration of group velocities and wavevectors for non-collinear phase-matching in OPAs for group velocity v_g (a) and wavevector k (b) matching. The larger group velocity of the idler wave leads to a separation of the idler and signal pulses through the nonlinear crystal, meaning that the bandwidth available for amplification is limited. This effect can be minimised by the introduction of a non-collinear angle between the pump and seed beams, allowing efficient phase matching to take place over a broader range of frequencies. The pump, seed and idler pulses are shown in blue, green and red, with the respective 'p', 's' and 'i' subscripts.

For further enhancement of the gain bandwidth, the second-order wavevector term $\mathrm{d}^2\Delta k/\mathrm{d}\omega^2$, associated with group velocity dispersion (GVD) mismatch, should also be minimised. This can be achieved by introducing angular dispersion, in which different frequency components propagate at slightly different non-collinear angles, enabling near-achromatic phase matching. In practice, this corresponds to the use of a chirped seed pulse, something which is naturally fulfilled in OPCPA systems.

OPCPA was first proposed in 1992 by A. Dubietis *et al.* [109] and later extended in 1997 by I. N. Ross *et al.* [110] to include analysis on beam quality, energy, contrast and bandwidth. Since then, OPCPA has led the way in achieving the high quality performance demanded of modern-day, high-power, short-pulse laser systems, the most notable of these include the Vulcan Nd:glass laser at the Central Laser Facility in Didcot, UK [111] (the facility which, incidentally, first pioneered the OPCPA technique) and the front end of the Omega EP laser at the University of Rochester [112]. Today, OPCPA technology underpins major UK laser upgrades currently in progress: both the Vulcan 20-20 project (aiming for 400 J in 20 fs) and the EPAC laser system at the CLF are being designed around OPCPA architectures.

2.6 Carrier Envelope Phase (CEP)

In Section 2.4, we showed how the frequency-dependent phase $\phi(\omega)$ can be described by a Taylor series expansion about a central frequency ω_0 (see Equation (2.43)). In this expansion, the zeroth-order term, $\phi_0 = k(\omega)z$, represents the absolute phase of the

pulse. That is, it determines the offset between the peak of the carrier wave electric field and the peak of the pulse envelope. This absolute phase, commonly referred to as the carrier envelope phase (CEP), is of limited importance when describing many-cycle pulses (where here we define many-cycle to mean longer than ~ 4 optical cycles). However, it becomes increasingly significant in the few-cycle regime, where the duration of the pulse envelope spans only a few periods of the carrier wave, and changes in CEP result in physical effects driven by differences in the detailed electric field evolution.

Figure 2.9 provides a visual representation of the Chimera mid-infrared few-cycle (30 fs at 3700 nm) field amplitude when varying CEP. As expected in the few-cycle regime, the waveform is sensitive to the absolute phase: varying the CEP causes noticeable changes in the temporal position of the field maxima relative to the envelope. Special cases of this occur at $\phi_0 = n\frac{\pi}{2}$, where the pulse is described as a 'cosine' pulse when n is even, and a 'sine' pulse when odd. Examples of cosine and sine pulses at $\phi_0 = 0$ and $\phi_0 = \pi/2$ are plotted in Figure 2.9 (b) and (c).

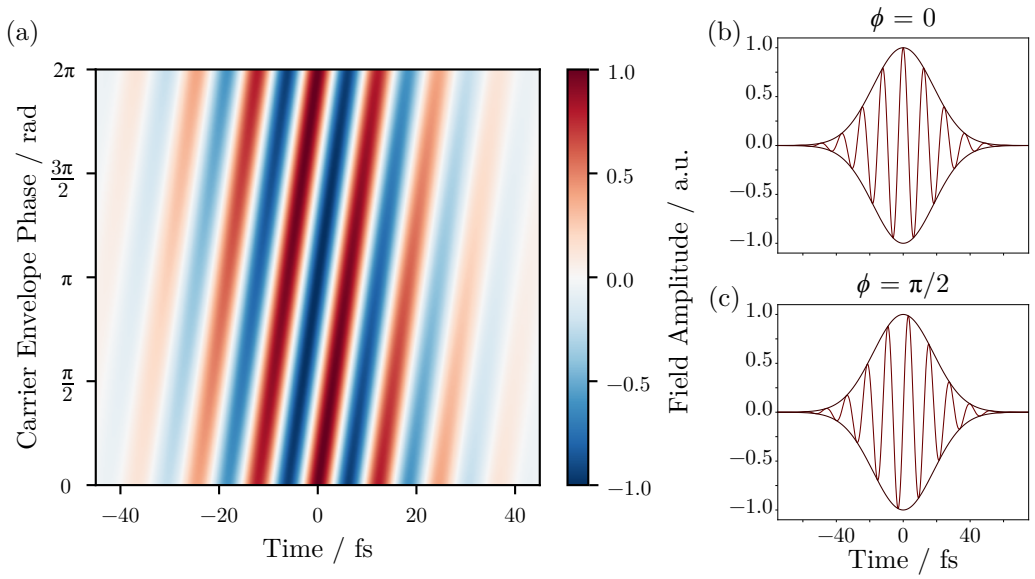


Figure 2.9: (a) The predicted electric field of the Chimera mid-infrared few-cycle laser pulse (30 fs at 3700 nm) as a function of both time and carrier envelope phase (CEP), where each horizontal slice corresponds to the electric field waveform for a given value of CEP. The plots in (b) and (c) show special cases of cosine and sine pulses, where $\phi_0 = 0$ and $\phi_0 = \pi/2$.

To further illustrate the significance of CEP in the few-cycle regime, we examine both the impact of varying the duration of a 3700 nm pulse and compare this to a pulse of the same 30 fs FWHM duration in the many-cycle regime at 800 nm. Figure 2.10 highlights how the electric field structure changes, both with changes in CEP and with pulse duration. As we might expect, the sensitivity to CEP becomes more pronounced as the number of cycles in the pulse decreases.

Ultrashort (and therefore few-cycle) pulses are of increasing interest within the community of high-intensity science. Notably, advances in the generation and measurement of attosecond pulses were recognised with the 2023 Nobel Prize in physics.

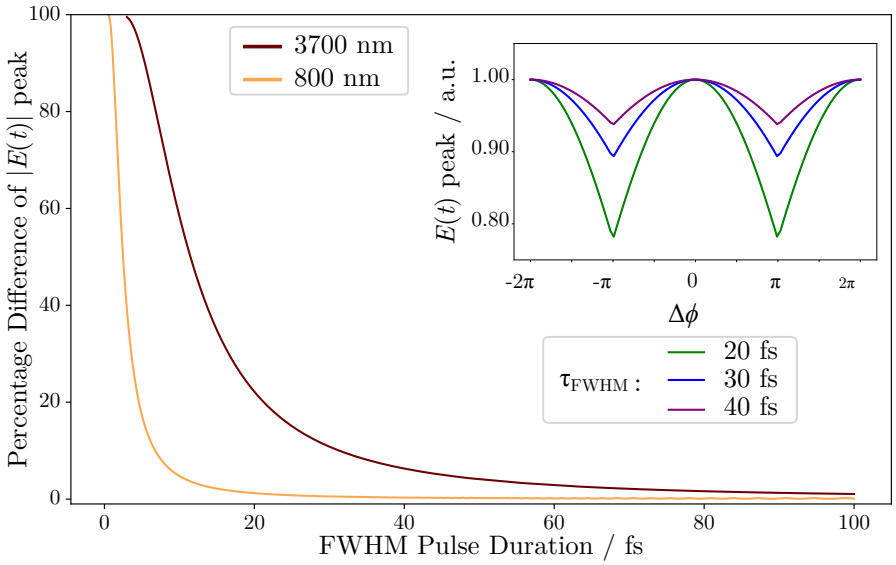


Figure 2.10: The electric field amplitude difference between the central strong peak and the next strongest peak as a function of FWHM pulse duration for two different wavelengths, 3700 nm and 800 nm. The inset shows the value of the peak electric field as a function of ϕ_0 for a 3700 nm wave at three different pulse durations.

Ultimately, though, the achievable pulse duration remains fundamentally constrained by the pulse's transform limit, with the shortest possible duration corresponding to a so-called monocycle pulse. This once more emphasises the importance of the laser's central wavelength. Despite running with a relatively long pulse duration compared to many 800 nm technologies (e.g. in our own system, we run with a notional theoretical limit of $\sim$ 30 fs at 3700 nm, and $\sim$7 fs at 800 nm), it still lies well within the few-cycle regime due to the longer central wavelength. While this may have potential benefits for many experiments requiring high intensities (such as high harmonic generation and laser wakefield acceleration experiments), the disadvantage we could face is that small changes in CEP, which would be accumulated even by the laser pulse propagating (and dispersing) through air, could have substantial shot-to-shot effects resulting from the significant variability in the field structure if not controlled.

2.7 Optical Damage

Even with the use of CPA in high-intensity laser systems, we are, of course, still susceptible to optical damage due to nonlinear effects. These invariably are due to high laser intensities close to, or beyond, the damage thresholds of bulk materials or optical coatings.

2.7.1 Self-focusing

One of the primary contributors to laser-induced damage is self-focusing, a consequence of the optical Kerr effect arising from the third-order nonlinear susceptibility $\chi^{(3)}$. Although the magnitude of $\chi^{(3)}$ varies across materials, all dielectric media exhibit some non-

zero third-order nonlinearity. This implies that under sufficiently high intensities, all will exhibit a nonlinear response – most notably, an intensity-dependent change in the refractive index:

$$n = n_0 + \frac{3\chi^{(3)}}{4n_0^2\epsilon_0 c} I = n_0 + n_2 I \tag{2.49}$$

Where n_0 is the low intensity material refractive index and n_2, defined as the nonlinear refractive index, is a material property which can be used to quantify the strength of the Kerr effect. The consequence of this intensity dependence is that when an intense laser beam passes with a local maxima (for example, a Gaussian or spatially 'noisy' beam) through a nonlinear material, it experiences a higher optical intensity along the beam axis than along the wings and effectively creates an on-axis lens that focuses the beam. This self-focusing in turn causes an effective spatial modulation of the nonlinear refractive index, the reduction of beam radius and thus the further increase of the intensity on axis. If this effect is strong enough, it can cause a total collapse of the beam and catastrophic optical damage, before it is ultimately limited by diffraction or plasma defocusing.

There is a threshold beyond which self-focusing effects become significant. This is known as the critical power for self-focusing P_{crit} and dictates the point of balance between linear diffraction and nonlinear self-focusing effects. Dependent only on laser wavelength and the nonlinear refractive index, it is given by:

$$P_{\text{crit}} = \alpha\left(\frac{\lambda^2}{4\pi n_0 n_2}\right) \tag{2.50}$$

Where α is a numerical factor dependent on the particular laser beam profile and λ is the central laser wavelength. In theory, then, it is feasible to design the possibility of any dramatic self-focusing events out of a system, although ultimately the CPA technique remains the favoured option for the safe, long-term operation of a high-power laser.

Other smaller-scale self-focussing events may still occur, for example, seeded by spatial noise from diffraction patterns caused by dust in the beam, potentially leading to hot-spots or filamentation within a system [**113**, **114**]. These are more challenging to suppress, and usually, preventative measures such as vacuum spatial filtering (VSF) are included in large systems to smooth out any small-scale intensity profile non-uniformities before they have a chance to evolve.

2.7.2 B-Integral

The B-integral is an intensity-dependent measure of the nonlinearity of light, where B stands for 'break up'. Specifically, it calculates the accumulation of nonlinear phase

through a material [115, 116]. The B-integral accumulated through a multi-element system of length L is given by the equation:

$$B = \frac{2\pi}{\lambda} \int_0^L n_2 I(z) \mathrm{d}z \tag{2.51}$$

Where $I(z)$ is the optical intensity along the beam axis z. Generally, it is desirable for the B-integral to be < 1 radians [115], although B-integrals between 1 and 3 radians are typically considered to be 'safe'. A B-integral below 1 is acceptable for reliably good temporal beam compression and a tight focus; a B-integral below π (~ 3) avoids catastrophic optical damage [115–117]. In practice, we aim to keep B $\leqslant 1$ to avoid the more subtle effects which may affect beam divergence. A 'large' B-integral would mean one > 5 radians, such that there is a significant accumulation of nonlinear phase, a high risk of self-focusing and the potential for considerable optical damage. Large B-integrals were historically accounted for by increasing the beam size in order to reduce intensities incident through a system; however, in practice, this is expensive and often impractical for very large intensities at the multi-terawatt to petawatt level. Nowadays, CPA is used to mitigate these effects.

2.8 Kerr Lens Modelocking (KLM)

Inadvertently discovered in 1990 [118], Kerr lens modelocking (KLM), sometimes referred to as 'self-modelocking' or 'magic modelocking', is a technique used to support ultra-short pulses in the femtosecond regime. It is a third-order nonlinear effect, relying on the Kerr nonlinearity (with positive $\chi^{(3)}$). Although this is usually an effect that might cause some concern within a laser system, it can, as in the case of a laser oscillator, also be exploited to support the formation of pulses well within the few-cycle regime [88, 119].

Many different gain media may be used in a Kerr lens modelocked laser source. Here, however, we will use Titanium doped Sapphire (Ti:Al_2O_3, commonly referred to as Ti:Sa), as this is an extremely high-bandwidth material [120] and the gain medium used in the Chimera laser oscillator.

A simplified cavity design of a typical KLM oscillator is shown in Figure 2.11. A gain medium, Ti:Sa, is inserted between two focusing optics with focal lengths f_1 and f_2. The cavity is contained by two end mirrors EM1 and EM2, with lens to mirror distances of L_1 and L_2, where $L_1 \leq L_2$. We refer to a stability range δ, within which lasing may occur, where δ is a free parameter contained within the distance between the focusing optics: $f_1 + f_2 + \delta$. This does not include any material thickness of the Ti:Sa crystal itself.

The stability requirements of the oscillator are confined within four limits of δ, within the bounds $\delta_0 < \delta < \delta_1$ and $\delta_2 < \delta < \delta_3$. Note that including the material thickness of the Ti:Sa will shift all limits by a factor of $L(1 - 1/n)$, where L is the material thickness and n is the refractive index, and real-world systems may use curved mirrors in place

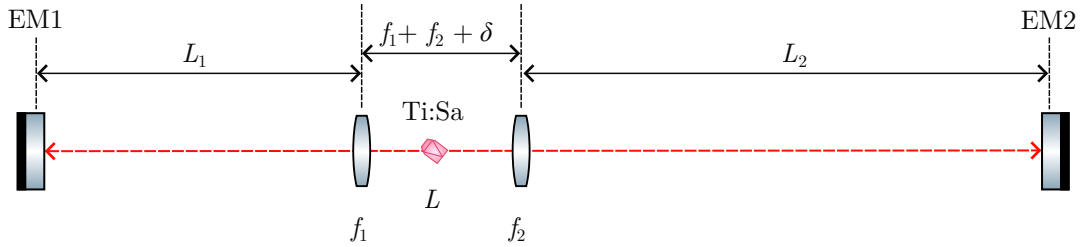


Figure 2.11: A schematic cavity representation of a typical Kerr lens modelocking oscillator. EM1 and EM2 depict the cavity end mirrors, L_1 and L_2 are the lens-to-mirror distances of each cavity arm, and f_1 and f_2 are the respective focal lengths of the cavity focusing optics. L is the material thickness of the Ti:Sa gain medium.

of refractive focusing optics. The four stability limits and their corresponding cavity modes can be visualised using geometrical optics, where the distance between the two focusing mirrors dictates either a plane or point focus at each of the end mirrors. Kerr lens modelocking is usually achieved in the vicinity of either δ_1 or δ_2 [119]. This is shown in Figure 2.12.

The purpose of KLM is to induce pulsed operation by ensuring that short, intense pulses experience lower cavity loss – or equivalently higher effective gain – than the continuous-wave (CW) mode. Neither the Ti:Sa crystal nor the cavity itself intrinsically favours pulsed operation; this preference is created by introducing a hard, soft, or 'virtual' aperture into the cavity. Owing to the Kerr effect, high-intensity pulses undergo stronger self-focusing than CW light, giving them a tighter spatial mode inside the gain medium. This spatial difference is converted into a temporal selection mechanism: the pulsed beam passes through the aperture with less loss, while the CW beam is clipped.

Modelocking initiation usually arises from noise fluctuations within the cavity. A spike in intensity can achieve sufficient peak power to experience Kerr-induced self-focusing, thereby passing the aperture with reduced loss relative to the CW, growing with every round trip and eventually seeding the development of a stable pulse train. In practice, modelocking is normally aided by a deliberate perturbation, such as a mechanical disturbance of the cavity.

Pulse shaping and stabilisation are governed by the interplay of two effects: Kerr-lensing in the spatial domain and group velocity dispersion in the time domain. While Kerr self-focusing preferentially confines high-intensity pulses, cavity dispersion tends to stretch them temporally, counteracting ultrashort pulse formation. To maintain femtosecond durations, intra-cavity dispersion compensation is required, typically implemented through prism pairs, chirped mirrors, wedged plates, or a combination thereof. Stable KLM operation, therefore, requires balancing these two mechanisms.

In the Chimera Ti:Sa oscillator, the space between the two curved mirrors (the so-called stability range) is manually adjusted so that the CW mode is deliberately made slightly broader at the crystal than the pulsed mode. As a result, CW light suffers an increased diffraction loss, while the pulsed beam, experiencing Kerr-induced self-focusing,

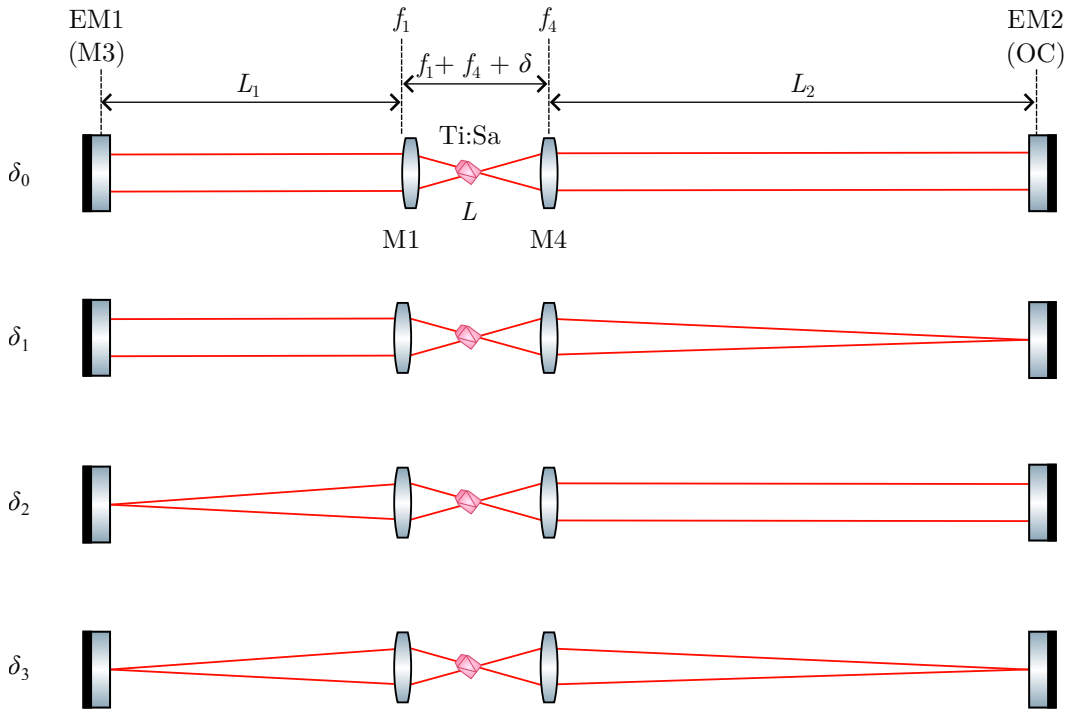


Figure 2.12: A geometrical representation of the four stability limits and corresponding cavity modes present within a Kerr lens modelocking cavity. The limits δ_0, δ_1, δ_2 and δ_3 represent the plane-plane, plane-point, point-plane and plane-plane modes, respectively. Stable cavity operation occurs within the bounds $\delta_0 < \delta < \delta_1$ and $\delta_2 < \delta < \delta_3$, usually very close to either δ_1 or δ_2. Here, EM1 and EM2 depict the cavity end mirrors, L_1 and L_2 are the lens-to-mirror distances of each cavity arm, and f_1 and f_4 are the respective focal lengths of the cavity focusing optics, M1 and M4. L is the material thickness of the Ti:Sa gain medium.

remains well confined. This arrangement acts as a virtual hard aperture inside the Ti:Sa crystal. Modelocking is then initiated when a noise spike or a perturbation, such as a small cavity length scan, provides the initial high-intensity fluctuation.

Finally, while CW and pulsed modes share the same collinear alignment, stable operation is obtained only when the cavity is slightly detuned from the optimum CW configuration. The most reliable working point is typically near the δ_1 stability position, where the CW focus is located at EM1 and a plane wavefront at EM2. Although modelocking can also be achieved near δ_2, the δ_1 regime is generally more stable when using a virtual hard aperture, as is common in KLM designs [119].

A crystal shift within such a cavity is typically required on a 6-monthly average period. This is due to the reliance on the optical Kerr effect, meaning that there will be gradual degradation of the Ti:Sa crystal over time. Most commercial systems of this type are nowadays manufactured to take this damage mechanism into account, with the crystal mounted upon a motorised stage and set to move on the order of a few microns within a 24-hour period.

3 Theory – Laser Wakefield Acceleration

This chapter develops the theoretical foundation for laser wakefield acceleration (LWFA). We begin by studying the motion of a single charged particle in an electromagnetic field. In the approximation of many laser cycles, this leads to the definition of the ponderomotive force – a time-averaged force experienced by a particle in a spatially varying oscillating field. Building on concepts introduced in Chapter 2, we extend the treatment of nonlinear optics to underdense plasmas and analyse how electromagnetic waves propagate in such media. This will bring us to how a laser wakefield is formed, along with the mechanisms that allow the self-injection and acceleration of electrons. We also address the limitations of the LWFA scheme and introduce the phenomenological framework of laser wakefield acceleration in the 3D nonlinear regime, developed by Lu *et al.* (2007) [74], which will be referenced in later chapters. The final section will introduce the mechanism of the direct acceleration of electrons by a laser field.

3.1 Single Particle Motion

For a single particle with charge q, velocity $\mathbf{v}$, mass m and momentum $\mathbf{p}$ in the presence of an electromagnetic field, its relativistic motion may be described by the Lorentz force as:

$$\mathbf{F} = \frac{\mathrm{d}}{\mathrm{d}t}\mathbf{p} = q(\mathbf{E} + \mathbf{v} \times \mathbf{B}) \tag{3.1}$$

The magnetic and electric field magnitudes being related (in vacuum) by $|\mathbf{B}| = |\mathbf{E}|/c$, at low velocities, particle dynamics are dominated by forces resulting from the electric field. They change substantially only when the particle approaches relativistic speeds, at which point the magnetic field term becomes comparable. For a particle moving at sub-relativistic velocities (i.e. $|\mathbf{v}| \ll c$), the magnetic component contributes negligibly to the force and the Lorentz equation can be simplified by neglecting the cross product on the right-hand side. For a linearly polarised plane wave electric field of the form

$\mathbf{E} = \mathbf{E}_0\cos(\omega_\mathrm{L}t$ - $\mathbf{kx} + \varphi_\mathrm{L})$, where ω_L is the laser frequency, $\mathbf{k}$ the wavevector and arbitrary phase φ_L, we can therefore reduce Equation (3.1) to:

$$m\frac{\mathrm{d}}{\mathrm{d}t}\mathbf{v} = q\mathbf{E}_0(\cos\left(\omega_\mathrm{L}t - \mathbf{k}\cdot\mathbf{x} + \varphi_\mathrm{L}\right)) \tag{3.2}$$

In this field, an electron ($q = -e$, $m = m_\mathrm{e}$) will experience a time-varying force leading to an oscillatory motion. Integrating the force with respect to time reveals the electron's velocity; an oscillatory velocity referred to as the quiver velocity:

$$\mathbf{v}_\mathrm{e} = \frac{e\mathbf{E}_0}{m_\mathrm{e}\omega_\mathrm{L}}\sin\left(\omega_\mathrm{L}t - \mathbf{k}\cdot\mathbf{x} + \varphi_\mathrm{L}\right) \tag{3.3}$$

This expression informs us that the described electron can reach a maximum speed of $v_\mathrm{e} = e|\mathbf{E}_0|/(m_\mathrm{e}\omega_\mathrm{L})$, determined by the magnitude of the electric field that surrounds it – i.e. how strongly the electron responds to the oscillating field. It is at this stage helpful to define the parameter a_0 – the normalised peak quiver momentum of a particle moving within an electric field:

$$a_0 \equiv \frac{e|\mathbf{E}_0|}{m_\mathrm{e}\omega_\mathrm{L}c} \tag{3.4}$$

Alternatively, a_0 is often defined through the field's vector potential $\mathbf{A}$, which relates to the electric field via $\mathbf{E} = -\frac{\partial}{\partial t}\mathbf{A} - \nabla\phi$. Neglecting the scalar potential ϕ in a plane wave, the peak of the normalised vector potential $\mathbf{a}$ is:

$$\mathbf{a} = \frac{e\mathbf{A}}{m_\mathrm{e}c} \tag{3.5}$$

Such that $a_0 = |\mathbf{a}|_\mathrm{max}$. This parameter a_0 allows us to distinguish between the non-relativistic or linear ($a_0 \ll 1$), and relativistic or nonlinear ($a_0 \gtrsim 1$) regimes of particle motion.

If the particle in question is travelling with velocities approaching the speed of light, the magnetic component of the Lorentz force can no longer be neglected. The definition of a_0 remains a useful one; however, the particle's motion is affected as a result of relativistic effects coming into play. In this relativistic regime, both the electric and magnetic fields play comparable roles in shaping the particle's dynamics, and the full Lorentz force equation (Equation (3.1)) must be retained.

3.1.1 The Ponderomotive Force

Before progressing further, it is helpful to consider a normalised version of the Lorentz force. Following [121, 122], we start by introducing the vector and scalar potentials ($\mathbf{A}$ and ϕ), of $\mathbf{E}$ and $\mathbf{B}$ such that Maxwell's Equations (Equations (2.1) - (2.4)), along with the equation for the Lorentz force, can be reduced to two equations:

$$\mathbf{E} = -\frac{\partial}{\partial t}\mathbf{A} - \nabla\phi \tag{3.6}$$

$$\mathbf{B} = \nabla \times \mathbf{A} \tag{3.7}$$

Using these potentials, we can update the Lorentz force to:

$$\frac{\mathrm{d}}{\mathrm{d}t}\mathbf{p} = q\left[-\frac{\partial}{\partial t}\mathbf{A} - \nabla\phi + \mathbf{v} \times \nabla \times \mathbf{A}\right] \tag{3.8}$$

With the vector triple product $\mathbf{A} \times (\nabla \times \mathbf{B}) = \nabla\ (\mathbf{A} \cdot \mathbf{B})$ - $(\mathbf{A} \cdot \nabla)\mathbf{B}$ and the relationship between the ordinary and partial differentials of time, (also known as the convective derivative), such that:

$$\frac{\mathrm{d}}{\mathrm{d}t} = \frac{\partial}{\partial t} + \frac{\mathrm{d}\mathbf{r}}{\mathrm{d}t} \cdot \frac{\partial}{\partial \mathbf{r}} = \frac{\partial}{\partial t} + \mathbf{v} \cdot \nabla \tag{3.9}$$

We can rewrite Equation (3.8) as:

$$\frac{\mathrm{d}}{\mathrm{d}t}\mathbf{p} = q\left[-\frac{\mathrm{d}\mathbf{A}}{\mathrm{d}t} + \mathbf{v} \cdot \nabla\mathbf{A} - \nabla\phi + \mathbf{v} \times \nabla \times \mathbf{A}\right] \tag{3.10}$$

$$\frac{\mathrm{d}}{\mathrm{d}t}(\mathbf{p} + q\mathbf{A}) = q\nabla(\mathbf{v} \cdot \mathbf{A}) - q\nabla\phi \tag{3.11}$$

Multiplying both sides by $1/(m_{\mathrm{e,i}}c)$, we can introduce the normalised quantities $\mathbf{a} = -q_{\mathrm{e,i}}\mathbf{A}/(m_{\mathrm{e,i}}c)$, $\phi \to -q_{\mathrm{e,i}}\phi/(m_{\mathrm{e,i}}c^2)$, $\mathbf{v} \to \beta = \mathbf{v}/c$ and $\mathbf{p} \to \mathbf{p}/(m_{\mathrm{e,i}}c)$ to find:

$$\underbrace{\frac{\mathrm{d}}{\mathrm{d}t}(\mathbf{p} - \mathbf{a})}_{①} = \underbrace{c\nabla\phi}_{②} - \underbrace{(c\nabla\mathbf{a}) \cdot \mathbf{v}}_{③} \tag{3.12}$$

We may now identify the individual terms in this equation. Here, ① refers to the rate of change of canonical momentum ($\mathbf{u}$), ② to the electrostatic force due to the

distribution of charge, and ③ to the coupling term between the electromagnetic field and current density. Using the definition of canonical momentum $\mathbf{u} = \mathbf{p} - \mathbf{a}$ and the further definition $\mathbf{p} = \gamma\beta = \gamma\mathbf{v}/c$, (where, in normalised units, the electron Lorentz factor $\gamma = \sqrt{1+p^2} = \sqrt{1+u^2+a^2+2\mathbf{u}\cdot\mathbf{a}}$), we can rewrite this as:

$$\frac{\mathrm{d}}{\mathrm{d}t}\mathbf{u} = c\nabla\phi - (c\nabla\mathbf{a})\cdot\frac{\mathbf{u}+\mathbf{a}}{\gamma} \tag{3.13}$$

$$\frac{\mathrm{d}}{\mathrm{d}t}\mathbf{u} = -\frac{1}{\gamma}c\nabla\frac{\mathbf{a}^2}{2} - c\nabla\mathbf{a}\cdot\frac{\mathbf{u}}{\gamma} + c\nabla\phi \tag{3.14}$$

Given the oscillatory motion of a particle in an electromagnetic field as described by the Lorentz force, it is useful to consider the time-averaged response over an arbitrary period ΔT. For this, we find:

$$\frac{1}{\Delta T}\int_T^{T+\Delta T}\mathrm{d}t\left[\frac{\mathrm{d}}{\mathrm{d}t}\mathbf{u} = -\frac{1}{\gamma}c\nabla\frac{\mathbf{a}^2}{2} - c\nabla\mathbf{a}\cdot\frac{\mathbf{u}}{\gamma} + c\nabla\phi\right] \tag{3.15}$$

Using angled brackets $\langle\rangle$ to denote time averaging, we finish with:

$$\left\langle\frac{\mathrm{d}}{\mathrm{d}t}\mathbf{u}\right\rangle = \underbrace{-\left\langle\frac{1}{\gamma}c\nabla\frac{\mathbf{a}^2}{2}\right\rangle}_{④} + \underbrace{c\nabla\langle\mathbf{a}\rangle\cdot\frac{\mathbf{u}}{\gamma}\Big|_T^{T+\Delta T}}_{⑤} + \underbrace{\int_T^{T+\Delta T}\mathrm{d}t\, c\nabla\langle\mathbf{a}\rangle\cdot\frac{\partial}{\partial t}\left(\frac{\mathbf{u}}{\gamma}\right)}_{⑥} \tag{3.16}$$

The standard approach at this stage is to consider the time-averaged response over a single oscillation $\Delta T = 2\pi/\omega$, and introduce the slowly varying envelope approximation (SVEA). This approximation takes advantage of the fact that a high-frequency pulse with $\omega = |\mathbf{k}|c$ has a spatiotemporal envelope of duration $\tau \gg 1/\omega$ and transverse extent $w \gg 1/|\mathbf{k}|$, implying that the envelope evolves over much longer time and length scales than the rapid carrier oscillations. Under this approximation, we may write $\langle\mathbf{a}\rangle \simeq 0$, causing the second and third terms in Equation (3.16), ⑤ and ⑥, to vanish. We are then left only with the first term ④: it is this which defines the nonlinear ponderomotive force acting on a single particle:

$$\mathbf{F}_\mathrm{p} \equiv \left\langle\frac{\mathrm{d}}{\mathrm{d}t}\mathbf{u}\right\rangle \simeq -\left\langle\frac{1}{\gamma}c\nabla\frac{\mathbf{a}^2}{2}\right\rangle \tag{3.17}$$

Since this force depends on $\mathbf{a}^2 = (-q_{\mathrm{e,i}}\mathbf{A}/(m_{\mathrm{e,i}}c))^2$, it is independent of the sign of the charged particle, meaning that the ponderomotive force expels all charged particles

from regions of high field intensity, no matter their charge. It does, however, have an inverse dependence on mass, meaning that although ions will indeed feel some force $\mathbf{F}_{\mathrm{p,i}}$, due to their significantly larger mass compared to the mass of the electrons, they feel a force over much longer timescales. This occurs on the order of an ion plasma period, $\tau_{\mathrm{i}} \sim \omega_{\mathrm{p,i}}^{-1}$, with the ion and electron plasma periods related by $\tau_{\mathrm{i}} = \sqrt{(m_{\mathrm{i}}/m_{\mathrm{e}})}\,\tau_{\mathrm{e}}$, meaning that the ion motion can be treated as negligible compared the motion of the electrons. The consequence of this is that if a high-intensity electromagnetic wave, for example, a short-pulse laser, propagates through a plasma, an overall repulsive force acts on the electrons, displacing them out of regions of high laser intensity, while the ions remain effectively stationary. The resulting charge separation sets up an electrostatic restoring force, exciting a plasma wave in the wake of the laser pulse. Although the initial electron motion driven by the transverse laser field is itself transverse, this expulsion occurs only while the laser pulse is present. Consequently, electrons return towards the axis at different longitudinal phases, generating a longitudinal density modulation that, according to Gauss's law, necessarily produces a longitudinal electrostatic field, within which particles can be trapped and accelerated. This mechanism underpins laser wakefield acceleration (LWFA). We will return to this in more detail in Section 3.3.

If, however, we are operating in the few-cycle regime, the SVEA no longer applies, and the simplification breaks down. The ponderomotive force, by definition, involves a time average over fast oscillations [**123**], making it nonsensical to extend it across all terms in Equation (3.16). If we may not use the pondermotive force to aid our understanding of laser wakefield formation, and predict and interpret behaviours and results, we bring upon ourselves an additional host of challenges. The ponderomotive force in its original definition remains the force that drives wakefield acceleration; however, when in a regime where the SVEA breaks down, it will no longer be the sole contributor to electron motion, nor can it necessarily be relied upon to interpret any results.

Looking at the full expression of the time-averaged force, (Equation (3.16)), we can try to understand the physical meaning of terms ⑤ and ⑥. For an electron interacting with a laser pulse, term ⑤ can be understood to represent a change in energy due to the spatial gradient of the laser's envelope ($\nabla\langle\mathbf{a}\rangle$). Physically, this describes the energy gained (or lost) as the electron enters or exits a laser pulse field with a non-uniform envelope – that is, when the strength of the laser field is changing in space. Under the SVEA, where the laser envelope is symmetric and evolves gradually, this term averages out to zero over time. In the few-cycle regime, the electron can experience asymmetry in the envelope, entering and exiting at different intensities, and the term becomes non-negligible. We refer then to this term as an 'envelope push' – it results not from the high-frequency oscillations of the carrier wave, but from the shape and asymmetry of the laser envelope itself.

Looking now at term ⑥, we can interpret this as the work done on an electron due to its own acceleration while moving through a non-uniform laser field envelope. Unlike term ⑤, which accounts for energy gained or lost as the electron enters or exits the laser

field, term ⑥ reflects the energy exchange while the electron is being accelerated within the spatially and temporally varying envelope. We can refer to this as a non-adiabatic envelope term, as it captures internal dynamics that arise when the envelope changes too rapidly for the electron motion to remain symmetric or adiabatic. Under the SVEA, this term again is negligible due to the symmetric nature of the pulse and electron trajectories. In the few-cycle regime, where the envelope changes rapidly, the electron's acceleration can become asymmetric, leading to a net energy gain or loss.

The impact of these effects on electron dynamics is still an open area of research. There is growing interest in how few-cycle pulses influence LWFA, especially in the context of scaling down laser energy to millijoule levels while increasing repetition rates to kilohertz in 800 nm systems. As discussed in Section 2.6, carrier envelope phase (CEP) plays a significant role in few-cycle pulse dynamics. In this regime, the peak of the electric field no longer necessarily aligns symmetrically with the pulse envelope, breaking the assumptions of the SVEA and enhancing non-adiabatic terms like ⑤ and ⑥. It is unsurprising, therefore, that this influence extends to the evolution of laser-driven plasma wakes. Most theoretical treatments to date have focused on incorporating the CEP into LWFA models [66, 124–127]. The first theoretical model addressing this was introduced in 2009 by Nerush and Kostyukov [126], who showed that although momentum gain from the ponderomotive force is independent of CEP, it can significantly affect the symmetry of the plasma bubble. This does not, however, mean that electron momentum will not be affected: transverse oscillations of the bubble due to CEP effects consequently cause asymmetry at the moment of self-injection of the electrons to be accelerated. Since the final electron energy correlates with injection conditions and output angle, CEP-induced distortions can degrade the beam quality, leading to irregular energy spectra or reduced acceleration efficiency [127].

We note too that under the SVEA, the ponderomotive force has no requirements in the polarisation of the radiation field. This is not the case when the SVEA ceases to be valid. While term ④ retains polarisation independence, terms ⑤ and ⑥ explicitly depend on the laser polarisation through the vector potential $\langle \mathbf{a} \rangle$, which is aligned with the polarisation axis. For linear polarisation, this introduces a strong CEP sensitivity along the laser's transverse axis, whereas for circular or elliptical polarisation, the rotational symmetry suppresses CEP dependence. The authors in [127] experimentally confirmed these effects by comparing CEP-driven spectral changes across different laser polarisations.

Although in this chapter we will continue to define LWFA with reference to the ponderomotive force, we should proceed with some caution.

3.1.2 Relativistic Motion

With the full derivation shown in Appendix B.1, in a stationary frame and assuming the particle starts from rest, its momentum as a result of interacting with a plane, monochromatic electromagnetic wave evolves as:

$$p_x = a_0 \sin\theta \tag{3.18}$$

$$p_z = \frac{a_0^2}{2} \sin^2\theta \tag{3.19}$$

Where we are using the vector potential $\mathbf{a} = a_x\hat{\mathbf{x}} = a_0 \sin\theta\hat{\mathbf{x}}$, with $\theta = kz - \omega t$ and momentum $\mathbf{p}$ normalised as before. As a function of θ, this results in spatial coordinates of:

$$x = a_0 \cos\theta \tag{3.20}$$

$$z = \frac{a_0^2}{4}\left(\theta - \frac{1}{2}\sin 2\theta\right) \tag{3.21}$$

With these solutions plotted in Figure 3.1, we see that the effects of travelling at relativistic velocities, or indeed the effect of including the magnetic field term, are that the particle will drift along the propagation direction of the wave with a cycle-average drift velocity $v_\mathrm{d}/c = \langle p_z\rangle/\langle\gamma\rangle = a_0^2/(4 + a_0^2)$. In the drift frame, we see the electrons process in the well-known figure-of-8 motion. This drift velocity term highlights the fundamental difference between relativistic and non-relativistic schemes, and suggests the possibility that particles may be accelerated via this mechanism (although the Lawson-Woodward Theorem disputes this for the case of a simple plane wave [128–130]).

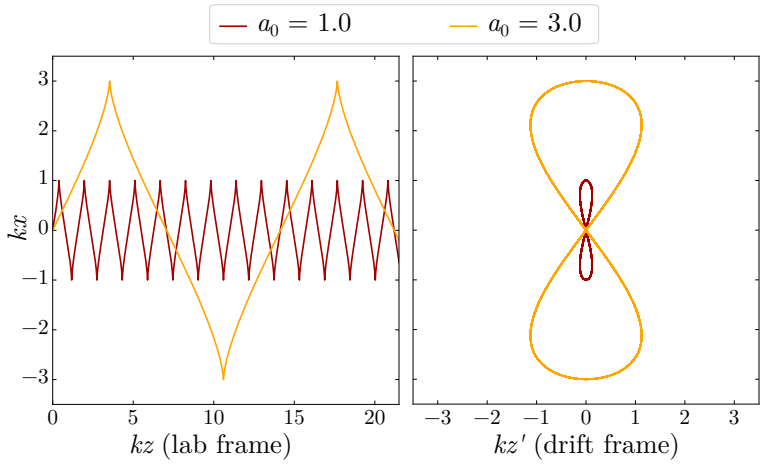


Figure 3.1: The trajectory of an electron while in the presence of an infinite, linearly polarised, monochromatic plane wave with field strength a_0 and wavenumber k in the lab (left) and drift (right) frames. The drift frame moves at the cycle-average drift velocity of the electron, tracing out a characteristic figure-of-8 motion.

3.2 Nonlinear Optics in Underdense Plasmas

Moving on from single-electron descriptions, we now extend the formalism introduced in Chapter 2, where we discussed the propagation of electromagnetic waves in vacuum, to consider the collective effects of an underdense plasma when perturbed by electromagnetic fields. This will lead to an expression for the plasma refractive index which, for intense radiation such as is present in a short-pulse laser, leads to nonlinear effects such as relativistic self-focussing and pulse compression.

3.2.1 EM Waves in Plasmas

Here we define a plasma to be a quasi-neutral ensemble of charged particles which exhibits collective electromagnetic behaviour [131]. Although many domains of interaction exist, dependent on the presence of EM waves or externally applied fields, here we will restrict ourselves to cold, low-density, neutral media (with no externally applied fields), whereupon electron plasma waves may oscillate if excited by a propagating EM wave.

Starting from the wave equation (Equation (2.9)), we note that since a plasma consists of free charges, not bound ones, the material polarisation vector $\mathbf{P}$ is taken to be zero. This simplifies Equation (2.5) for the electric displacement field to $\mathbf{D} = \epsilon_0 \mathbf{E}$. For a quasi-neutral plasma with charge density $\rho = 0$ and current density $\mathbf{J} = -en_\mathrm{e}\mathbf{v}$ (and assuming that the plasma reacts sufficiently fast to remain quasi-neutral – i.e. to maintain $\nabla \cdot \mathbf{E} = 0$), we can define a wave equation for an electric field $\mathbf{E}$ propagating through a plasma:

$$\nabla^2 \mathbf{E} = -\mu_0 e n_\mathrm{e} \frac{\partial \mathbf{v}}{\partial t} + \frac{1}{c^2}\frac{\partial^2 \mathbf{E}}{\partial t^2} \tag{3.22}$$

Using the simplified (non-relativistic) Lorentz Equation (Equation (3.2)), we may relate the acceleration of electrons due to an electric field via $\partial \mathbf{v}/\partial t \approx \mathrm{d}\mathbf{v}/\mathrm{d}t = -e\mathbf{E}/m_\mathrm{e}$, such that:

$$\nabla^2 \mathbf{E} = \frac{\mu_0 e^2 n_\mathrm{e}}{m_\mathrm{e}} \mathbf{E} + \frac{1}{c^2}\frac{\partial^2 \mathbf{E}}{\partial t^2} \tag{3.23}$$

Substituting a plane wave solution $\mathbf{E} = \mathbf{E}_0 \exp\left[i(\mathbf{k}\cdot\mathbf{r} - \omega t)\right]$ into the wave equation gives:

$$-k^2 \mathbf{E} = -\frac{\mu_0 e^2 n_\mathrm{e}}{m_\mathrm{e}} \mathbf{E} - \frac{\omega^2}{c^2}\mathbf{E} \tag{3.24}$$

Which simplifies to the dispersion relation for electromagnetic waves in a cold plasma:

$$\omega^2 = \omega_\mathrm{p}^2 + c^2 k^2 \tag{3.25}$$

Here ω_p is the plasma frequency, defined as:

$$\omega_\mathrm{p} \equiv \sqrt{\frac{n_\mathrm{e} e^2}{\epsilon_0 m_\mathrm{e}}} \tag{3.26}$$

This frequency characterises the natural oscillation of electrons in the plasma and sets a lower bound for wave propagation. For $\omega < \omega_\mathrm{p}$, waves are evanescent and can only penetrate the plasma up to a length of c/ω_p; a length scale defined as the skin depth. For $\omega > \omega_\mathrm{p}$, transmission is allowed. From the dispersion relation, we can define the phase velocity v_p and group velocity v_g of the wave as:

$$v_\mathrm{p} = \frac{\omega}{k} = \frac{c}{\sqrt{1 - \frac{\omega_\mathrm{p}^2}{\omega^2}}} \tag{3.27}$$

$$v_\mathrm{g} = \frac{\partial \omega}{\partial k} = c\sqrt{1 - \frac{\omega_\mathrm{p}^2}{\omega^2}} \tag{3.28}$$

These relations highlight the dispersive nature of plasmas – the group velocity is always less than c, preserving causality, and decreases as the wave frequency approaches the plasma frequency, whereas the phase velocity is always greater than c. Finally, the plasma refractive index, defined as $\eta = c/v_\mathrm{p}$, takes the form:

$$\eta = \sqrt{1 - \frac{\omega_\mathrm{p}^2}{\omega^2}} = \sqrt{1 - \frac{n_\mathrm{e}}{n_\mathrm{c}}} \tag{3.29}$$

Where n_c is the critical plasma density for a given wave frequency, defined by $\omega = \omega_\mathrm{p}$. Above this density, wave propagation is forbidden, and the plasma becomes reflective to the incident radiation. For $n_\mathrm{e} < n_\mathrm{c}$, the plasma is considered to be 'underdense' – this regime is the one in which LWFA can occur. From Equation (3.29) we see that in contrast to conventional dielectric media, the plasma refractive index decreases with increasing electron density. Consequently, regions of lower plasma density correspond to regions of higher refractive index, allowing for optical guiding in plasmas with a transverse density profile with an on-axis minimum. This enables a plasma channel, capable of guiding a laser pulse over distances greatly exceeding its vacuum Rayleigh length, to be formed. In a uniform plasma, where the refractive index is constant, no such guiding occurs, and the laser diffracts unless additional nonlinear or external guiding mechanisms are present.

We note that the phase velocity of the plasma wave is determined by the group velocity of the driver that excites it – in the case of LWFA, this would be the laser pulse. Accordingly, the relativistic Lorentz factor associated with the plasma wave, γ_p, is given by:

$$\gamma_\mathrm{p} = \frac{1}{\sqrt{1-(\frac{v_\mathrm{g}}{c})^2}} \approx \frac{\omega}{\omega_\mathrm{p}} \tag{3.30}$$

This approximation holds in the limit where the laser frequency ω is much greater than the plasma frequency ω_p, which is typically the case in underdense plasmas used for LWFA.

The instantaneous Lorentz factor of a quivering electron is given as $\gamma(t) = (1+a^2(t))^{\frac{1}{2}}$. In a weakly relativistic scheme where $a_0 \ll 1$, it can be approximated to $\gamma(t) \approx 1 + \frac{1}{2}a^2(t)$. This relativistic correction enters the refractive index relation through the modified plasma frequency, $\omega_\mathrm{p}^2 \rightarrow \frac{\omega_\mathrm{p}^2}{\gamma}$. For weakly relativistic pumps, small density perturbations (δn) and slowly varying frequency shifts ($\delta\omega$), the index of refraction becomes $\eta = (1-\omega_\mathrm{p}^2/\langle\gamma\rangle\omega^2)^{\frac{1}{2}}$ and, using a Taylor expansion (ignoring products between perturbations), can be expanded as [132]:

$$\eta = \left(1 - \frac{1}{2}\frac{\omega_\mathrm{p}^2}{\omega^2}\left(1 + \frac{\delta n}{n} - \frac{\langle a^2\rangle}{2} - 2\frac{\delta\omega}{\omega}\right)\right) \tag{3.31}$$

This expression highlights how intensity (via $\langle a^2\rangle$), density modulations, and frequency shifts influence the propagation of the laser in a plasma, underpinning nonlinear optical effects such as relativistic self-focusing and temporal pulse compression. These effects are the respective results of transverse and longitudinal gradients in the refractive index.

3.2.2 Relativistic Self-Focussing

In the case of a highly intense ($a_0 \geq 1$) laser pulse, the normalised vector potential can no longer be approximated as a weakly relativistic plane wave. Instead, its spatial variation across the transverse profile of the beam becomes significant. Assuming a laser pulse approximately Gaussian in shape, there will be a variation of the normalised vector potential a_0 across the pulse form, peaking on-axis. This causes the refractive index to be largest on-axis and to decrease radially outward, generating a refractive index gradient that acts like a focusing lens. This phenomenon is a third-order ($\chi^{(3)}$) nonlinear optical effect, analogous to Kerr lensing in solid-state media, and is the same effect that underpins Kerr-lens modelocking, as described in Chapter 2. In plasmas, the equivalent effect is known as relativistic self-focusing, originating from the relativistic mass increase of electrons in the high-intensity laser field. Spatial variations in the local refractive

index can also arise from the expulsion of electrons from regions of high laser intensity – a ponderomotive effect. These mechanisms can contribute to and enhance the relativistic self-focusing of the laser pulse [133].

Self-focusing effects counteract the natural diffraction experienced by a laser propagating in a medium, meaning that the pulse can be guided and sustain high intensities over many Rayleigh lengths. The condition for self-focusing can be expressed as a threshold on the beam waist w_0 [132]:

$$\frac{a_0^2}{32} w_0^2 \frac{\omega_\mathrm{p}^2}{c^2} > 1 \tag{3.32}$$

This can be translated into a form of a critical laser power P_c:

$$P_\mathrm{c} = \frac{8\pi\epsilon_0 m_\mathrm{e}^2 c^5}{e^2} \frac{\omega^2}{\omega_\mathrm{p}^2} \tag{3.33}$$

$$\approx 17.4 \frac{\omega^2}{\omega_\mathrm{p}^2} \text{ GW} \tag{3.34}$$

A laser pulse with a power P, self-focusses for $P > P_\mathrm{c}$, self-guides for $P = P_\mathrm{c}$ and diffracts for $P < P_\mathrm{c}$. This guiding behaviour can also be exploited experimentally. For example, Thomas *et al.* [134] demonstrated that short pulses focused using high F number optics (i.e., with slow convergence) can be self-guided over many Rayleigh lengths within a plasma with minimal pulse energy loss. Self-guiding is a critical enabling mechanism in laser wakefield acceleration (LWFA), allowing the laser pulse to maintain sufficient intensity to drive a wake over the full dephasing length – and thus enabling electron acceleration to multi-GeV energies (see Section 3.3.4).

3.2.3 Pulse Compression and Photon Acceleration

If there exists a plasma frequency gradient arising from a gradient in the refractive index, then there will necessarily be a corresponding gradient in the group v_g (and phase v_p) velocity of the laser pulse (see Equations (3.28) and (3.27)). Consider two points x_1 and x_2 in the pulse envelope, separated by a distance L, and moving with group velocities $v_{\mathrm{g}1}$ and $v_{\mathrm{g}2}$. After some time Δt, the separation will change by:

$$\Delta L = (v_{\mathrm{g}2} - v_{\mathrm{g}1})\Delta t = \Delta v_\mathrm{g} \Delta t = c\Delta\eta\Delta t \tag{3.35}$$

Approximating Δv_g as $\approx \frac{\partial v_\mathrm{g}}{\partial x} L$, the fractional rate of change of the pulse length can be written as [132]:

$$\frac{1}{L}\frac{\partial L}{\partial t} = \frac{\partial v_\mathrm{g}}{\partial x} = c\frac{\partial \eta}{\partial x} \tag{3.36}$$

Since the refractive index increases with a_0 (see Equation (3.31)), this implies that the front of the pulse, where intensity is rising and a_0 is increasing, experiences slower group velocity, while the back of the pulse, where intensity is decreasing, experiences a faster group velocity. This longitudinal refractive index gradient leads to temporal pulse compression, as the rear of the pulse catches up with the front.

Pulse compression is driven by longitudinal rather than transverse refractive index gradients. These can originate from any of the three terms in Equation (3.31). As these gradients exist in both space and time, their impact is not limited to the temporal compression of the pulse envelope – they also induce spatially varying frequency shifts. This phenomenon, in which photons gain or lose energy as they propagate through plasma density gradients, is known as photon acceleration and is the spectral counterpart to temporal pulse compression.

Similarly to the case of pulse compression, spatial gradients in the refractive index also lead to variation in the local phase velocity of the laser pulse, and hence in its wavelength. We can consider two wavefronts at positions x_1, x_2, separated by λ and moving with velocities $v_{\mathrm{p}1}$ and $v_{\mathrm{p}2}$. The difference in their speeds leads to a change in their separation over time:

$$\Delta\lambda = \Delta v_\mathrm{p}\Delta t \approx \frac{\partial v_\mathrm{p}}{\partial x}\lambda\Delta t \tag{3.37}$$

Using $v_\mathrm{p} = c/\eta$ (Equation (3.27)), the rate of change of wavelength is then:

$$\frac{\partial \lambda}{\partial t} = \lambda\frac{\partial v_\mathrm{p}}{\partial x} = -\frac{c\lambda}{\eta^2}\frac{\partial \eta}{\partial x} \tag{3.38}$$

This describes photon acceleration: when the refractive index increases along the direction of propagation ($\partial\eta/\partial x > 0$), the phase velocity decreases, and the wavelength shortens, leading to an effective increase in photon energy. In the plasma (or co-moving) frame, this process can be interpreted as a shift in the photon's frequency. Differentiating $\omega = ck/\eta$ with respect to time, assuming fixed k gives:

$$\frac{\partial \omega}{\partial t} = \frac{c\omega}{\eta^2}\frac{\partial \eta}{\partial x} \tag{3.39}$$

Connecting the temporal frequency shift to the spatial variation of the refractive index. For an increase (decrease) in refractive index, the frequency of light decreases (increases) with time – there is a red (blue) shift [135, 136].

3.3 Laser Wakefield Acceleration (LWFA)

This section introduces the laser wakefield acceleration (LWFA) mechanism, including the nonlinear 'bubble' or 'blowout' regime, common electron injection techniques (with emphasis on self-injection), key energy gain limitations, and the widely used Lu scaling model [74] for predicting electron beam parameters.

To generate a laser wakefield, a high intensity ($> 10^{18}$ Wcm^{-2}), short pulse (≤ 1 ps[1]) laser is injected into a cold, low density ($\sim 10^{16}$ - 10^{19} cm^{-3}) neutral medium, which is thereupon ionised into a plasma by the leading edge of the laser. Light pressure created by the pulse, known as the ponderomotive force, expels electrons in the plasma away from regions of high intensity. This electron displacement leaves behind immobile (on the relevant timescales), positively charged ions, creating a restoring electrostatic field. As the electrons are pulled back toward equilibrium, they overshoot, and a 'wake' of longitudinal plasma oscillations, known as a plasma wakefield, is created behind the laser pulse. The frequency of this oscillation is referred to as the plasma frequency, given by Equation (3.26). According to Gauss's Law, the charge imbalance caused by the oscillating density wave will necessarily create large electric fields within which particles, if successfully trapped, may be accelerated.

3.3.1 Plasma Wave Formation

The simplest way to derive the generation of plasma waves is in the one-dimensional linear regime, where the laser intensity is weak ($a_0 \ll 1$), and electron motion is non-relativistic. For a 1D system where all electrons are displaced by a small distance x relative to the stationary ion background, the resulting charge separation produces a restoring field. Using Gauss's law, this field is

$$E_x = -\frac{en_\mathrm{e}}{\epsilon_0}x \tag{3.40}$$

Where the negative sign indicates the field acts to restore the electrons to equilibrium. Applying the Lorentz force for $v \ll c$ (or indeed simply Newton's second law!), the equation of motion is therefore:

$$\frac{\mathrm{d}^2x}{\mathrm{d}t^2} = -\omega_\mathrm{p}^2 x \tag{3.41}$$

[1] Of course, these days picosecond timescales are no longer considered to be very short!

Where ω_p is the plasma frequency, $\omega_p^2 = e^2 n_e/(m_e \epsilon_0)$. This is the equation for a simple harmonic oscillator, with general solution:

$$x(t) = A \sin(\omega_p t - \phi) \tag{3.42}$$

Where ϕ is a phase constant. These local oscillations of electrons at frequency ω_p correspond, collectively, to a longitudinal electrostatic plasma wave left in the wake of the laser pulse. For such a wakefield to form efficiently, the group velocity of the laser pulse must approximately match the plasma wave phase velocity, with $k_p = \omega_p / v_g$.

Since this derivation assumes small displacements, it implicitly bounds the maximum sustainable electric field. Beyond this, the plasma wave can break – that is, electron trajectories can cross, and coherence is lost. The corresponding maximum longitudinal field in a cold, non-relativistic plasma, commonly known as the cold non-relativistic wave-breaking limit[2], is given by [137]:

$$E_0 = \frac{m_e \omega_p c}{e} \tag{3.43}$$

For a plasma with electron density 10^{18} cm^{-3}, this corresponds to a theoretical maximum field strength of ~ 100 GV/m.

To go beyond this 1D linear approximation, a more complete description in three dimensions can be obtained by using the cold fluid model, which couples the Poisson equation, the continuity equation, and the fluid momentum equation. We will not include this derivation here; it can be found in full in [138, 139].

For laser intensities $a_0 > 1$, the quiver motion of electrons becomes relativistic, and the plasma response enters a highly nonlinear regime. In this case, the plasma can sustain even higher electric fields, described (in 1D) by the relativistic cold wave-breaking limit [140]:

$$E_{WB} = \sqrt{2}(\gamma_p - 1)^{\frac{1}{2}} E_0 \tag{3.44}$$

Where the relativistic gamma factor γ_p associated with the phase velocity of the plasma wave is given as $\gamma_p = (1 - v_p^2/c^2)^{-\frac{1}{2}}$. In the context of laser-driven plasma waves,

[2] In his paper, Dawson actually defines the cold non-relativistic wave-breaking limit specifically for a finite phase velocity $v_p \neq c$. Although the expression we quote is indeed the one commonly used, it is naïve and incorrect to use $v_p = c$ – if the plasma wave moves at exactly c, then in the cold, non-relativistic limit, no electrons could ever be trapped. It would require an infinitely strong field for the electrons to catch up with the plasma wave [137].

where the wake phase velocity is approximately equal to the laser group velocity in an underdense plasma, one may approximate $\gamma_\mathrm{p} \approx \gamma_\mathrm{g} \approx \sqrt{1+\frac{a_0^2}{2}}$ such that [141]:

$$E_\mathrm{WB} = E_0 \frac{a_0^2/2}{\sqrt{1+a_0^2/2}} \tag{3.45}$$

For a full 3D treatment, a non-perturbative, numerical treatment is needed, such that the fields are computed self-consistently with the evolving particle distribution. This is a task promoted to kinetic simulations such as particle-in-cell (PIC) methods (see Section 4.2).

At the point of wavebreaking, the plasma wave becomes steep enough that background electrons can become trapped in the accelerating phase of the wake. This process enables self-injection and the subsequent acceleration of electrons.

3.3.2 Electron Injection

To be accelerated by the steep longitudinal electric fields of a wakefield, electrons must first be injected into the accelerating phase of the plasma wave. Several experimental injection schemes have been developed to achieve this, including self-injection, ionisation injection [70, 71], and density down-ramp injection [67]. The first scheme to be observed experimentally [142], self-injection, is also the simplest to implement as it arises naturally when the wakefield amplitude becomes large enough for wavebreaking to occur.

Described by Thomas in 2010 [143] by analysing electron trajectories in the wakefield potential within the bubble regime (introduced in the next section, Section 3.3.3), the threshold field strength required for self-injection to take place is given by:

$$a_0 \gtrsim \ln\left(2\gamma_\mathrm{p}^2\right) - 1 \tag{3.46}$$

Where γ_p is the Lorentz factor associated with the wake's phase velocity. This threshold was experimentally validated by Mangles *et al.* [144], who further expressed the condition in terms of the laser power P and the fraction of energy α contained within the full-width-at-half-maximum (FWHM) focal spot:

$$\alpha\frac{P}{P_\mathrm{c}} > \frac{1}{16}\left[\ln\left(\frac{2n_\mathrm{c}}{3n_\mathrm{e}}\right) - 1\right]^{\frac{1}{3}} \tag{3.47}$$

Where P_c is the critical power for relativistic self-focusing (see Equation (3.33)), n_c is the critical density, and n_e is the plasma density.

While self-injection is the only injection scheme considered in this thesis, ionisation injection is now more widely used in LWFA experiments. This preference stems from its effectiveness at lower laser intensities and its ability to produce higher charge beams with improved quality and reduced divergence [71, 145]. In contrast, self-injection is often considered a more stochastic process, which can result in reduced control over the beam quality.

With mid-infrared (MIR) laser drivers, it is not yet clear which injection scheme should be favoured. This uncertainty stems from the different scaling of electric field strength with laser wavelength and intensity: the peak electric field of a laser scales as $E \propto \sqrt{I}/\lambda$, meaning that for a given intensity, longer wavelengths (such as in the MIR) produce lower peak electric fields. This has important consequences for injection thresholds, particularly for schemes such as ionisation injection, which rely on field-induced tunnelling [146].

3.3.3 The Bubble Regime

The bubble, or 'blowout' regime, is a highly nonlinear acceleration regime that occurs when a high-intensity, short-pulse laser with normalised vector potential $a_0 \gtrsim 1$ propagates through an underdense plasma. The bubble regime is so-named because when a high-intensity driver is directed at a plasma, the electrons in the plasma are entirely expelled from the axial region, creating a bubble-like cavity in the wake of the laser pulse. In this regime, the ponderomotive force of the laser is strong enough to expel nearly all electrons from the central axis of propagation, forming an ion cavity – or 'bubble' – devoid of electrons, trailing behind the laser pulse. The expelled electrons are pushed radially forward and outward and then pulled back toward the axis, forming a dense sheath that wraps around the rear of the bubble in a quasi-spherical trajectory. At the rear, the steep density spike is accompanied by wavebreaking, leading to self-injection of electrons into the bubble cavity. This location also coincides with the strongest accelerating fields. Under optimal conditions, the fields inside the bubble become quasi-linear and uniform, providing excellent conditions for accelerating injected electrons with monoenergetic spectra. This was first reported in simulation in 2002 by Pukhov and Meyer-ter-Vehn [147], 2 years before being shown experimentally.

Figure 3.2b shows the approximately spherical shape of the bubble regime cavity, formed by a laser with initial $a_0 = 2.9$, compared to a weakly nonlinear plasma wakefield driven by a laser with $a_0 = 0.9$ in Figure 3.2a. The figures are both overlaid with the longitudinal electric field E_x felt by the electrons. In the $a_0 = 0.9$ case, the field is close to being sinusoidal in shape – the wave does not break, and no electron injection is observed. In the $a_0 = 2.9$ case, the field takes a characteristic sawtooth-like profile, with steep gradients at the density spike at the back of the bubble corresponding to the region of maximum accelerating fields. Here, the fields exceed the wavebreaking limit, leading to self-injection of electrons, which can be seen accumulating at the rear of the bubble.

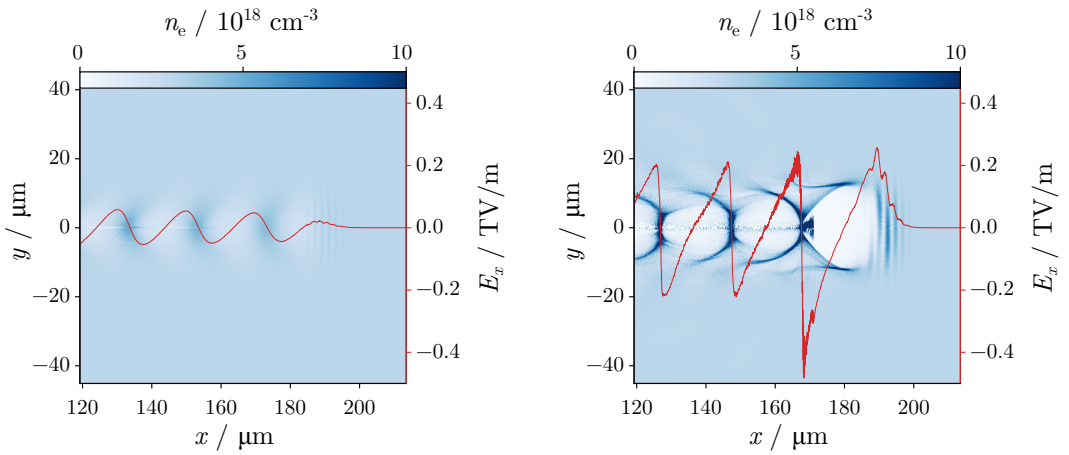


Figure 3.2: An electron density plot from an AM cylindrical simulation of LWFA using the particle-in-cell code Smilei. The plots show a laser ($\lambda = 3700$ µm, $\tau_{\mathrm{FWHM}} = 30$ fs) propagating left-to-right across the simulation window after being focussed to a spot size of $d_{\mathrm{FWHM}} = 13.6$ µm in a plasma with electron density $n_e = 2.8 \times 10^{18}$ cm^{-3}. The longitudinal electric field E_x is overlaid in red. The snapshot is taken roughly 0.77 ps into the simulation. (a) A modestly relativistic regime of LWFA (driving laser starts with $a_0 = 0.9$). The laser pulse excites a weakly relativistic wake. The plasma modulation and resulting longitudinal electric field are observed, but the bubble is not fully cavitated, and no self-injection (or acceleration) occurs. (b) A highly relativistic regime of LWFA (driving laser starts with $a_0 = 2.9$). The laser pulse excites a strongly relativistic wake, with a fully developed bubble cavity within which the electron density drops close to zero. The electric field becomes highly nonlinear, exceeding the wavebreaking threshold, triggering electron self-injection into the bubble cavity, visible at the rear of the structure.

Several scaling laws exist for laser intensity, spot size, and pulse duration to optimise bubble formation and electron injection while minimising energy losses. We explore these in more detail in Section 3.3.5, but broadly, the best performance is typically achieved when the focal spot size is comparable to or larger than the plasma wavelength ($w_0 \gtrsim \lambda_p$). Tightly focused beams tend to suffer from diffraction, leading to beam breakup and poor shot-to-shot reproducibility [134]. In contrast, 'matching' the laser spot size to the blowout bubble radius ($w_m = w_0 \sim R_b$) allows the pulse to self-guide without significant distortion of the pulse profile over the interaction length, leading to more stable and reproducible wakefield conditions [74]. While the matched condition ultimately gives the most stable configuration, starting with a slightly larger spot size ($w_0 > w_m$) allows the pulse to evolve naturally through self-focusing and compression, with the stabilising effect of self-focusing driving the pulse towards the optimal matched profile.

3.3.4 Limitations on Electron Energy Gain

Laser wakefield acceleration (LWFA), particularly with the phenomenon of electron self-injection, is often regarded as one of the more fortuitous outcomes in plasma physics, where several physical effects align in a way that favours efficient particle acceleration. The ponderomotive force exemplifies this. Given its origin in oscillating electric fields, one might expect a dependence on laser polarisation. However, as can be seen in

Equation (3.17), the ponderomotive force exhibits no such dependence: it is insensitive to the laser's polarisation and behaves identically for both linear and circular polarisations, provided the energy density remains the same. Another advantageous feature is the onset of relativistic self-guiding. This naturally extends the interaction length between the laser and plasma, allowing electrons co-propagating with the laser pulse to gain energy over distances far exceeding the Rayleigh range. Despite these beneficial properties, several physical limitations constrain the maximum energy gain achievable in LWFA. The most important of these are: laser diffraction, electron dephasing, and pump depletion.

Laser Diffraction

As a laser propagates through any medium, it undergoes diffraction, causing the beam to expand transversely. The characteristic length over which diffraction becomes significant is known as the Rayleigh range, defined as the distance from the beam waist w_0 to the point where the beam radius has increased by a factor of $\sqrt{2}$. In the context of laser wakefield acceleration, diffraction can limit the effective interaction length between the laser and plasma. It is possible, however, to mitigate diffraction through mechanisms such as plasma waveguiding [30], or by exploiting relativistic self-focusing, which, under the right conditions, leads to self-guiding of the laser pulse and the associated wakefield [148].

Electron Dephasing

Dephasing occurs when accelerated electrons catch up to the plasma wave. As they move forward within the wake, they eventually enter regions where the electric field reverses, causing deceleration and degrading beam quality. In relativistic laser-plasma interactions, electrons are accelerated to velocities approaching the speed of light, $v \approx c$, while the phase velocity of the plasma wave is slightly less, $v_\mathrm{p} < c$. As a result, the electrons inevitably overtake the wake, a process which occurs over the dephasing length, L_d. This sets a fundamental limit on the acceleration distance in LWFA. The linear and nonlinear dephasing length scales (in 3D) are calculated by [74]:

$$L_\mathrm{d} = \frac{n_\mathrm{c}}{n_\mathrm{e}} \frac{\lambda_\mathrm{p}}{2\pi} \begin{cases} 1 & \text{for } a_0 \ll 1 \\ \frac{4}{3}\sqrt{a_0} & \text{for } a_0 \gg 1 \end{cases} \tag{3.48}$$

Where λ_p is the plasma wavelength, given by $\lambda_\mathrm{p} = 2\pi c\sqrt{\frac{\varepsilon_0 m_\mathrm{e}}{e^2 n_\mathrm{e}}}$, n_e is the electron density and n_c is the critical density.

To produce monoenergetic electron beams with low energy spread, the acceleration length must remain shorter than, or approximately equal to, the dephasing length. Once electrons begin to dephase, the beam experiences broadening in both energy and divergence. It is therefore generally desirable to design plasma stages such that the interaction length does not exceed L_d [29].

Pump Depletion

Pump depletion is a natural consequence of 'donating' laser energy to the plasma. As the pulse drives the plasma wave, it is continually losing energy. Since the laser energy is finite, there will necessarily be a point at which the laser can no longer support accelerating gradients, thereby limiting the energy gain of the captured electrons. The point at which this has a negative effect on electron energy gain is called the depletion length, L_{pd} – this is the distance at which the laser pulse has lost a significant fraction of its initial energy. In 3D, the linear and nonlinear pump depletion lengths can be estimated as [74]:

$$L_{\mathrm{pd}} = \frac{n_{\mathrm{c}}}{n_{\mathrm{e}}}\lambda_{\mathrm{p}} \begin{cases} \frac{1}{a_0^2} & \text{for } a_0 \ll 1 \\ 1 & \text{for } a_0 \gg 1 \end{cases} \tag{3.49}$$

It is interesting here to note that in the 3D nonlinear regime, the pump depletion length becomes independent of the laser intensity. This was first identified by Decker *et al.* in 1996 [149], who showed that in the strongly nonlinear regime, energy loss is dominated not by distributed absorption, but by etching of the laser's leading edge. In this regime, the front of the laser pulse is continuously depleted as it excites the plasma wave, while the tail of the pulse propagates through the nearly evacuated cavity. As a result, the laser erodes from the front backwards: energy is lost primarily from the front, leading to a shortening of the pulse. The rate at which this occurs is known as the etching velocity, given by $\beta_{\mathrm{etch}} = n_{\mathrm{e}}/n_{\mathrm{c}}$. Both the width of the density spike at the back of the bubble and the rate of energy loss carry the same a_0 dependence, which causes the a_0 terms to cancel out in the depletion length. Consequently, in the bubble regime, the pump depletion length becomes effectively intensity-independent, governed solely by the plasma density and wavelength.

For a range of laser strengths a_0 and plasma densities n_{e}, the curves of constant dephasing and pump depletion lengths are shown in Figure 3.3, calculated for a laser with wavelength $\lambda = 3700$ nm using Equations (3.48) and (3.49) for $a_0 \gg 1$. The figure illustrates the strong dependence of both limits on plasma density. Notably, the nonlinear pump depletion length exhibits no dependence on a_0 in this regime, while the dephasing length shows only a weak a_0 scaling.

3.3.5 Lu Scaling Model

In 2007, a phenomenological framework of laser wakefield acceleration in the 3D nonlinear blowout regime was set out by Lu *et al.* [74]. This pioneering paper established a set of scaling laws, which the vast majority of experiments still use to help design and inform their setups and data. Here we discuss the main results.

A fundamental result is the empirical relationship between the laser's normalised vector potential a_0 and the radius R of the plasma bubble it drives:

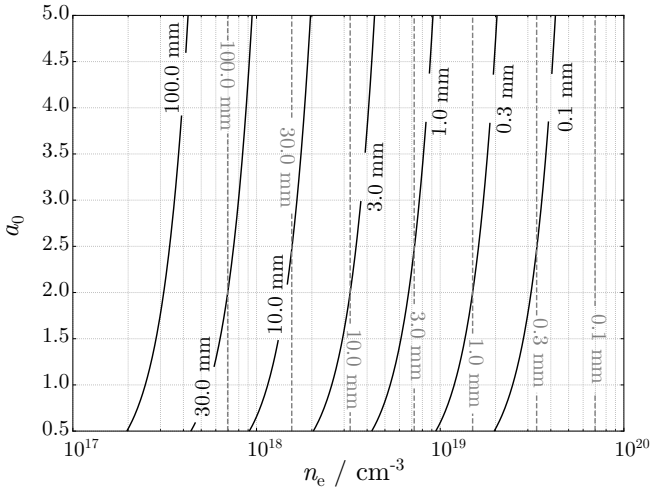


Figure 3.3: The dephasing (black, solid) and depletion (grey, dashed) lengths plotted as functions of a_0 and plasma density n_e.

$$k_p R \simeq 2\sqrt{a_0} \tag{3.50}$$

From this, the concept of a matched spot size is introduced – the focal spot size w_0 that ensures stable guiding of the laser through the plasma. Optimal guiding occurs when the laser spot size matches the blowout radius, $k_p w_0 \sim k_p R$.

The characteristic dephasing and pump depletion lengths in this regime are also derived, as given in Equations (3.48) and (3.49).

Most significantly for the results presented in this thesis, Lu *et al.* derive a scaling law for the maximum energy W_{max} attained by accelerated electrons:

$$W_{max}/m_e c^2 = \frac{2}{3} a_0 \frac{\omega_0^2}{\omega_p^2} = \frac{2}{3} a_0 \frac{n_c}{n_e} \tag{3.51}$$

An obvious limitation of the Lu model is that it does not account for any evolution of the laser pulse – it assumes an initially matched spot size. As discussed in Sections 3.2.2 and 3.2.3, as a laser propagates through an underdense plasma, it will inevitably undergo both spatial and temporal evolution. Such evolution – including self-focusing or pulse compression – leads to an increase in the laser's a_0, making it likely that Lu's original energy scaling in Equation (3.51) underestimates the actual energies available to accelerated electrons.

This limitation was addressed in 2020 by Bloom *et al.* [**150**], who proposed an updated model that incorporates laser pulse evolution during propagation. Assuming that the

pulse evolves toward the matched spot size, the maximum value of a_0 can be estimated as $a_{\max} \approx 2[\alpha E_{\mathrm{L}}/(\tau_{\mathrm{f}} P_{\mathrm{c}})]^{1/3}$, where α is the fraction of laser energy E_{L} trapped within the bubble, P_{c} is the laser critical power and τ_{f} is the final predicted pulse duration after pulse compression. Here τ_{f} is approximated as $\tau_{\mathrm{f}} \approx \tau_0 - n_{\mathrm{e}} l/(2cn_{\mathrm{c}})$, where the pulse evolves over a propagation length l. The maximum propagation length is limited either by the pump depletion length ($L_{\mathrm{pd}} = c\tau_0 n_{\mathrm{c}}/n_{\mathrm{e}}$), or, for $L_{\mathrm{target}} < L_{\mathrm{pd}}$, the length of the target. The revised maximum energy of the accelerated electrons then becomes:

$$W'_{\max}/m_{\mathrm{e}}c^2 = \frac{2}{3} a_{\max} \frac{n_{\mathrm{c}}}{n_{\mathrm{e}}} \tag{3.52}$$

3.4 Direct Laser Acceleration

Direct laser acceleration (DLA) is a mechanism that can lead to significant additional energy gain by electrons in the bubble regime of a laser wakefield accelerator. Unlike standard LWFA, where electrons gain energy primarily from the strong longitudinal electric fields in the plasma wake driven by the laser's ponderomotive force, DLA involves a direct transfer of energy from the laser field itself to the electrons. This process is often described as a 'direct $\mathbf{v} \times \mathbf{B}$ push' [151]. In near or fully blown-out regimes, electrons trapped in the first plasma wave period (or 'bucket') oscillate transversely in the plane of laser polarisation. If the frequency of these oscillations resonates with the laser field experienced by the electrons, a coherent transfer of energy can occur between the laser's transverse electric field $\mathbf{E}_\perp$ and the electron's transverse velocity $\mathbf{v}_\perp$, increasing the transverse momentum. This transverse momentum can then be converted into a longitudinal momentum via the $\mathbf{v} \times \mathbf{B}$ force associated with the laser field. The energy gained by electrons due to the transverse electric field can be expressed as [152]:

$$W_\perp = -e \int_0^t \mathbf{E}_\perp \cdot \mathbf{v}_\perp dt' \tag{3.53}$$

Where $\mathbf{v}_\perp$ is the transverse velocity of the electron. Since the dominant contribution to $\mathbf{E}_\perp$ arises from the laser field, this term quantifies the DLA component of the total energy gain in the LWFA regime.

For effective energy gain via DLA, electrons must remain within the transverse extent of the laser field long enough to accumulate energy, and the frequency of their oscillation must be in resonance with the Doppler-shifted laser frequency. Such conditions are most favourable in long, guided plasma channels where the laser field can maintain overlap with the oscillating electrons over extended distances.

DLA has a characteristic signature in electron momentum space: it exhibits a forked structure in its momentum phase space. If present, this is something that would manifest itself in an electron spectrum.

4 Methods

This chapter describes the key experimental and numerical methods used in this thesis. We will explain a short-pulse diagnostic used to measure temporal durations on the Chimera laser system, as well as methods used for the detection and characterisation of its mid-infrared beam. We will also discuss the simulation method used to map out and plan a proposed laser wakefield acceleration experimental parameter space. Simulation convergence tests are performed in order to present and provide benchmarking for a consistent simulation framework for all other simulations presented in this thesis.

4.1 Laser Diagnostics

Historically, a temporal profile could be measured simply by using a device with a shutter speed shorter than the event itself. When referring more specifically to lasers, 'long' nanosecond pulses can be directly recorded with photodiodes coupled to a fast oscilloscope, and few picosecond pulses can be recorded with streak cameras, but there does not exist a camera which can capture a shorter pulse. Ultra-fast pulses (in the femto- to attosecond regime) must therefore be measured using alternative techniques; they must be measured – or 'correlated' – against the shortest thing that exists: very often, themselves.

The simplest way to achieve this in practice is by crossing a copy of the pulse with itself in a nonlinear medium, typically using a Michelson interferometer to introduce a temporal delay τ between the two pulses. The delayed pulses are then spatially and temporally overlapped in the medium – for example, a nonlinear crystal used for second-harmonic generation, or a photodiode or LED acting as a two-photon detector. The latter of these methods requires a band gap energy greater than the photon energy of the laser pulse, relying on the principle that electrons will be excited across a material bandgap if and only if two photons are absorbed simultaneously [153–155]. In both cases, the measured signal (e.g. current or light intensity) is proportional to the square of the intensity of the light directed at it. From this we can extract an autocorrelation signal $S(\tau)$, which is given by [115]:

$$S(\tau) \propto \int_{-\infty}^{\infty} \Big(|E(t) + E(t+\tau)|^2 \Big)^2 \mathrm{d}t \tag{4.1}$$

$$\propto \int_{-\infty}^{\infty} \Big(|E(t)|^2 + |E(t+\tau)|^2 + 2\Re[E(t)E(t+\tau)^*] \Big) \mathrm{d}t \tag{4.2}$$

Where $E(t)$ is the electric field of the laser pulse. Two broad types of autocorrelation exist: intensity autocorrelation and interferometric autocorrelation. If the detector has a sufficiently large bandwidth to resolve the rapidly oscillating interference term (third term in Equation (4.2)), then we have an interferometric autocorrelation and the signal as a function of delay will show fringes with spacing $\Delta\tau = 2\pi/\omega_0$, where ω_0 is the central frequency of the pulse. If the detector is too slow to resolve this term, the result is an intensity autocorrelation, which provides information only about the pulse envelope. Interferometric autocorrelations are sensitive to the carrier envelope phase, and their signals typically exhibit fringes with both positive and negative lobes due to the oscillating electric field. In contrast, intensity autocorrelations are always positive and symmetric, since the detector integrates over the fast oscillations and responds only to the average intensity. Both types of autocorrelation can be implemented in scanning or single-shot configurations. In scanning systems, the delay τ is varied using a movable delay line to reconstruct the autocorrelation trace over time. In a single-shot system, the time delay is encoded spatially across a beam profile and detected simultaneously using a camera.

Although a relatively simple technique to implement in the laboratory, crucial information about the shape of the pulse envelope can not be extracted when using this method – in order to make a quantitative estimate about the pulse duration we detect, we must assume some pulse envelope, meaning that any phase information is lost. In order to extract both the amplitude and the phase of an ultrashort pulse, we must also measure its spectrum. The first technique to do this, the FROG (frequency-resolved optical gating), was developed by Kane and Trebino in 1993 [156].

In this section, we will describe one method – TOAD (Third Order Autocorrelator Diagnostic) or XTOAD (Extended Third Order Autocorrelator Diagnostic) – used to characterise the OPA pump pulse of the Chimera laser using techniques that can also provide information about the pulse shape and contrast. Many other such pulse characterisation tools exist and are in regular experimental use worldwide, for example FROG [156], SPIDER [157] and DScan [158, 159] (and all their variations), however these diagnostics have not yet been applied to the work carried out in this thesis and have therefore been omitted.

As mid-infrared detectors are much less well developed compared to their visible/near-infrared counterparts, we also describe methods used to detect and characterise our mid-infrared pulse.

4.1.1 Third Order Autocorrelator Diagnostic (TOAD)

Ultrashort pulse diagnostics typically rely on some instance of nonlinearity. As nonlinear processes are intensity-dependent, they are highly sensitive to the spatial and temporal characteristics of a pulse, allowing critical pulse information, such as duration and spectral phase, to potentially be recovered. Most autocorrelation-based diagnostics techniques use second-order processes such as two-photon absorption or second harmonic generation (SHG). FROG, SPIDER and DScan all use SHG as a tool to characterise their input short pulses, with additional steps used to extract spectral phase. The reason for this is that SHG is adaptable for a very large range of wavelengths and, compared to higher-order processes, is typically very robust. SHG is also relatively efficient, even with low input intensities, and so can generate stronger signals with simpler experimental setups and less demanding detection requirements.

Despite its considerable advantages, there are two major disadvantages when using SHG to diagnose short pulses. Firstly, scattered light from the SHG crystal can significantly reduce the measurable contrast and dynamic range. Secondly, SHG autocorrelation inherently produces time-symmetric traces, as the technique relies on overlapping two identical replicas of the pulse. The consequence of this is that it is not possible to distinguish whether temporal structural features in the pulse wings occur before or after the main portion of the pulse. This ambiguity can be problematic in high-intensity laser-plasma experiments, particularly with fragile solid targets, where pre-pulses or pedestal structures arriving before the main pulse can initiate unwanted interactions or damage.

In order to produce an asymmetric, high contrast trace to enable the characterisation of a high dynamic range short pulse, we may turn to a third-order autocorrelator. First designed at Imperial College in 1993 [160], the Third Order Autocorrelator Diagnostic (TOAD) was proposed as a very high dynamic range ($\sim 10^{11}$) pulse characterisation technique. An upgraded design was later developed at Imperial College (described extensively in [116]), which was used to characterise a range of modelocked laser oscillator sources and high-power, compressed laser pulses from a selection of major UK laser laboratories [161]. Commercial third-order autocorrelation devices are also available – for example, the SEQUOIA (Amplitude Technologies). The SEQUOIA advertises a noise floor limited dynamic range of $> 10^9$ [162].

The Imperial College TOAD is a custom-built 1 µm diagnostic with a temporal scan window of 2 ns. It includes an 'extended' (XTOAD) mode, with the addition of a flip-in cross-polarised wave generation (XPW) stage as a way to further enhance the ability to discriminate against diagnostic artefacts. This is done by providing an additional nonlinear 'pulse cleaning' step for a second harmonic probe beam. Its development was partially supported by AWE contract 30242504 "Advanced technologies for high-contrast short-pulse operation on the Orion Laser Facility" and is routinely used to diagnose the Cerberus laser system at Imperial College across a dynamic range of $> 1 : 10^{11}$. XTOAD

mode does not apply to the measurements presented in this thesis but is described in detail in [163]. The optical layout of the Imperial College TOAD is shown in Figure 4.1.

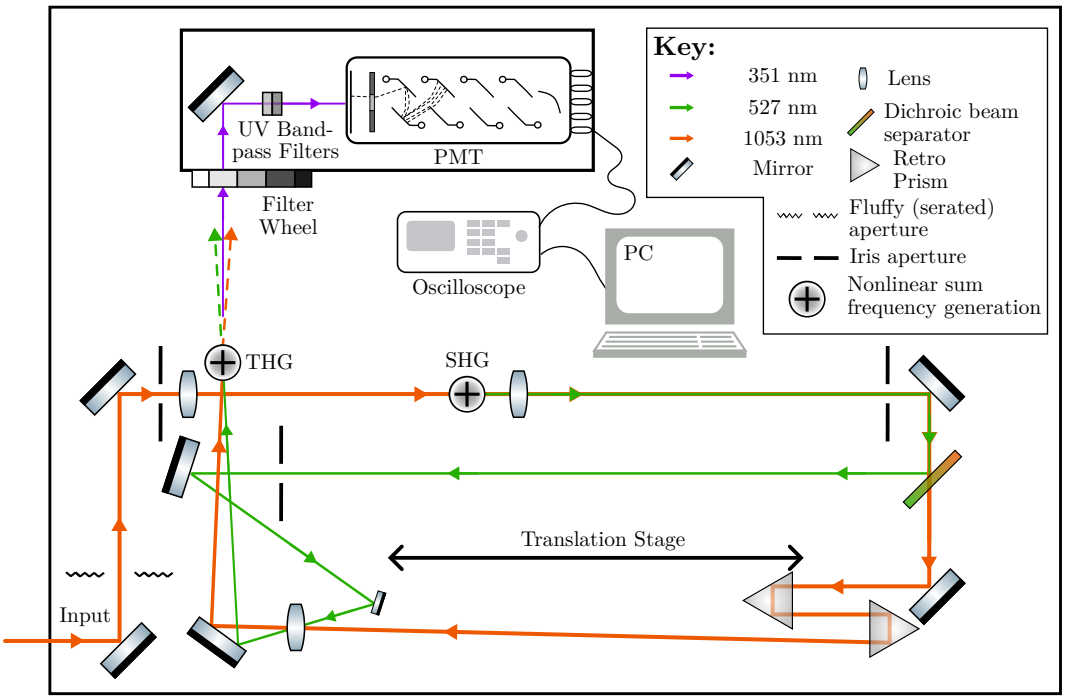


Figure 4.1: Optical layout of the Imperial College Third Order Autocorrelator Diagnostic (TOAD). The TOAD is a two-stage system, whereby the second harmonic light generated in the first nonlinear SHG stage is cross-correlated with the fundamental in the second stage. This produces third-harmonic light, which can be measured as a function of time delay between the fundamental and SHG probe. The high-gain photomultiplier tube (PMT) is contained in a light-tight housing, inside which the third harmonic signal is spectrally and spatially filtered to provide a very low background measurement. Both the filter wheel and the translation stage can be operated manually or through a computer interface, with software developed by Dr Nicholas Stuart [116].

The basic TOAD/XTOAD principle of operation is in the non-collinear cross-correlation of the laser fundamental with its second harmonic (rather than itself) to produce a third harmonic signal by way of sum-frequency generation. In the third order frequency mixing process, the second harmonic signal acts as a probe to the fundamental, made possible due to its shorter pulser duration (of order $1/\sqrt{2}$), allowing a temporal intensity measurement as a function of delay. The position encoded motorised translation stage (Thorlabs LTS300/M) has 30 cm of available travel, with a temporal step size of down to 5 fs. This provides high-resolution measurements over a 2 ns scan window (when the translation stage is double-passed). Although introducing an extra degree of complexity from a second-order correlator, as it becomes a two-stage rather than single-stage process, it allows an inherently background-free, time-asymmetric autocorrelation measurement: there is no third harmonic signal when either the second harmonic or fundamental pulses are not temporally overlapped, or if they are individually incident inside the mixing crystal. Beam paths for the fundamental, second and third harmonic beams are also angularly separated after the third harmonic generation crystal, allowing for robust discrimination of the 'real' signal against other background processes.

The third harmonic signal is detected using a light-tight, spectrally filtered, high-gain photomultiplier tube (PMT, ET-Enterprises 9124B) and recorded on a Tektronix Oscilloscope. The Imperial TOAD can be operated manually or through an automated computer interface – both the filter wheel and the translation stage may be controlled using software developed by Dr Nicholas Stuart [116]. As this is a third-order process and the autocorrelation signal scales as the cube of the electric field intensity, a Gaussian deconvolution factor of $1/\sqrt{3}$ is required to retrieve a pulse duration.

Although designed primarily to monitor laser contrast, the data we present in this thesis uses the TOAD only as a pulse duration diagnostic used to characterise the compression and temporal profile of the 1 µm Chimera OPA pump beam.

4.1.2 Mid-Infrared Diagnostics

Thanks to the wide availability of affordable and convenient diagnostics, beam profiling and spectral characterisation of laser systems operating in the visible and near-infrared is often relatively straightforward. For routine beam monitoring and alignment of more common systems (in the 400 - 1000 nm wavelength range), standard silicon-based cameras are used, as they are widely accessible and often cost less than £20. Similarly, most compact spectrometers, also relying on silicon detectors, are most readily available in the visible-near-infrared wavelength range. For wavelengths above ∼1100 nm, however, silicon becomes transparent and can no longer be used as a direct single-photon detector. This makes beam profiling and spectral characterisation of the mid-infrared (MIR) Chimera laser beam (centred at 3700 nm) much more challenging, particularly as there is a limited number of sensor materials and technologies that are sensitive to the MIR spectral range.

In addition to its wavelength, another important factor is the system's repetition rate. High-repetition-rate systems (in the kilohertz range) are generally much easier to characterise, both spatially and spectrally, due to their higher average power. In the MIR stage of the Chimera laser, we enter into a low energy (microjoule - millijoule), low repetition rate (10 - 100 Hz) regime. This results in a low average power, presenting significant challenges for both beam profiling and spectral detection, particularly when needing to discriminate against a 70 MHz co-propagating unamplified oscillator background. We now describe the MIR spectrometer and beam profiler used to address some of these challenges on the Chimera laser system.

Mid-Infrared Spectrometer

Although the general principles behind near- and mid-infrared spectrometers are the same (i.e. dispersing light via a grating or prism and detecting it using a suitable sensor), the practical challenges involved in MIR detection make these systems far more specialised and significantly more expensive. One of the main limitations arises from material transparency: regular glass and fused silica are opaque in the MIR range, requiring the use of more costly materials such as calcium fluoride (CaF_2) and zinc selenide (ZnSe) for

lenses and windows. The manufacture of gratings also becomes problematic, as standard gratings, optimised for use in the near-infrared, are not appropriate for high-quality MIR characterisation. Instead, custom gratings made from MIR-compatible materials, with non-standard line spacing, are often required.

At one stage, the system included a home-built scanning spectrograph with a resolution of ~ 50 nm, consisting of a 300 l/mm grating mounted on a rotation stage and imaged onto a single-element PbSe photodiode detector with a 2.4 µm longpass filter. Although no longer in operation – its components having since been repurposed elsewhere in the system – this spectrometer was originally designed and used by Dr. Nicholas Stuart to collect the MIR spectral data presented in Figure 5.20 (Chapter 5).

The MIR spectrometer currently used in the Chimera system, also built by Dr Nicholas Stuart, consists of a commercial spectrometer body retrofitted with a custom MIR grating and mount, coupled to a Pb:ZnSe linear detector array. This is the spectrometer used for the MIR spectral data presented in Figure 5.25. Although the Pb:ZnSe detector operates as a photon detector rather than a thermal one, MIR measurements are nonetheless susceptible to significant ambient thermal background radiation. To discriminate the laser signal from this background, an optical chopper wheel is employed to modulate the incident laser signal, enabling internal normalisation against the laboratory thermal background.

Mid-Infrared Beam Profiler

The requirements for a mid-infrared beam profiler are to provide fast and convenient 2D spatial detection, enabling detailed characterisation of the beam's shape and intensity distribution. Due to the developmental nature of mid-infrared systems and their relatively limited use, the choice of suitable camera systems is extremely narrow. This means that MIR beam profilers are often incredibly specialised. On the Chimera laser system, we use a pyroelectric detector (Spiricon, Pyrocam IIIHR) for this purpose. Pyroelectric detectors are made from materials, such as $LiTaO_3$ (Lithium Tantalate) or DLaTGS (Deuterated L-Alanine-doped Triglycine Sulfate), that exhibit pyroelectricity – the generation of an electric charge in response to a change in temperature. This temperature variation across the detector surface creates a voltage difference, which is then processed to reconstruct the beam's spatial profile and intensity distribution. Although the minimum detectable average power of the camera (> 1 mW) is larger than that generated on the MIR seed generation alone, the use of external triggering allows reliable detection and characterisation of the MIR beam at all subsequent stages in the system.

4.2 Numerical Methods

When attempting to model a plasma, there must necessarily be a trade-off between accuracy, generality and computational cost. This stems from the inherent complexity of plasma systems: no single numerical tool can fully capture all relevant physics across

the wide range of spatial and temporal scales involved. Depending on the regime of interest, plasma behaviour spans from sub-femtosecond time scales (e.g., electron plasma oscillations) to seconds of confinement (e.g., magnetic or inertial fusion), and from nanometre-scale features like the Debye length to metre-scale structures in astrophysical jets.

To optimise computational cost, a choice must be made based on the dominant physics. Simulation tools are therefore specialised to capture different effects: collisionless plasmas require kinetic models; collisional plasmas are often treated using fluid (hydro) approximations; radiatively dominated plasmas require radiation-hydrodynamics, and the presence of strong magnetic fields calls for magnetohydrodynamic (MHD) models. In the case of laser wakefield acceleration (LWFA), where the plasma is underdense, relativistic, and largely collisionless, a kinetic approach is necessary, and thus Particle-in-Cell (PIC) simulations are the tool we turn to.

4.2.1 Particle-in-Cell (PIC) Simulations

Due to the limited ability to well-diagnose LWFA experiments, both as a consequence of their physical size ($\sim$ microns) and the small timescales ($\sim$ picoseconds - femtoseconds) over which they occur, experimental access to fast transient phenomena such as electron injection, beam evolution and laser self-focussing and depletion, as well as the spatial mapping of the plasma bubble and electron bunch trajectory remains extremely limited. When conducting such an experiment, conventional diagnostics such as electron and X-ray spectrometers and back-scattered laser beam profilers are designed to capture post-interaction signatures, meaning that information on physical processes happening within the interaction itself is largely inferred. To this end, numerical simulation, particularly kinetic particle-in-cell codes, is an essential tool in laser-plasma physics. Indeed, since being implemented to support the very first proposal of the plasma acceleration scheme by Tajima and Dawson in 1979 [28], they have been applied extensively to diagnose, explore, and understand many pioneering theoretical and empirical advances. These include the first experimental results of monoenergetic electron beams in 2004 [29–31], seminal work on deriving scaling laws by Lu *et al.* in 2007 [74] (see Section 3.3.5), Bayesian optimisation to improve the beam quality and stability of the laser-plasma accelerator (first in simulation, then in experiment) [164], and in exploring as yet well-chartered regimes such as positron acceleration [165].

To construct such a simulation, one must consider directly solving Maxwell's equations for every simulated particle (a real experiment will contain particle densities of $\mathcal{O}(10^{24})$ m^{-3}), at each position of interest within a 6-dimensional phase space. Laser pulses must be able to freely and nonlinearly evolve, with considerations such as pulse-compression and self-focussing, and there must be a capability to correctly simulate and resolve relativistic, non-Maxwellian particle beams. All this must be simulated over millimetre length scales and picosecond durations. At maximum fidelity, this is, unsurprisingly, a hugely computationally expensive feat. To make it operationally feasible,

the particle-in-cell (PIC) algorithm was developed. Rather than directly evolving the electromagnetic fields of both the laser and the plasma, they are instead calculated at the vertices of a predefined grid, thus reducing the computational load. To reduce the load even further, large ensembles of real particles can be discretised and represented as 'macro-' or 'quasi-particles', owing to their collective behaviour. It is these two computational simplifications which earned the algorithm its moniker: *particle-in-cell*.

4.2.2 The PIC Algorithm

Before the PIC loop can be implemented, the simulation must first be initialised with particle starting positions, preliminary values for charges, current densities and fields, and boundary conditions at the edges of the simulation domain. This can be done as a restart of a previous simulation or with a user-defined input file. Thus, with user-prescribed spatial profiles for number density, particles per cell, mean velocity and temperature for each simulated particle species at a time $t = 0$, quasi-particles can be loaded into each cell, and the initial electromagnetic fields can be calculated over the whole simulation grid.

PIC codes often use the Finite-Difference Time-Domain (FDTD) Method as a way to reduce computational expense when solving Maxwell's Equations [166]. This method defines a staggered simulation grid, meaning the electric and magnetic field components are defined at offset spatial locations rather than co-located at the same grid points – a popular choice being the Yee grid [167]. To ensure that particle motion is simultaneously updated with the fields, PIC codes use leapfrog particle pusher schemes (such as the Boris [168] and Vay [169] pushers) to update particle positions and momenta whilst ensuring energy conservation. The particle momenta themselves are calculated from the Lorentz force, with the fields calculated using Maxwell's equations in the previous step.

The general PIC computation loop, from a timestep n to a timestep $n + 1$, is then as follows:

1. Gather and interpolate electromagnetic fields at particle positions.
2. Push all particles forward & compute new particle velocities and positions.
3. Project new charge and current densities onto the grid.
4. Solve Maxwell's equations and compute the new electromagnetic fields on the grid.

A drawback of solvers like FDTD is their tendency to suffer from unphysical numerical dispersion. This becomes especially concerning when simulating relativistic particles as such dispersion can lead to numerical Cherenkov radiation, resulting in unphysical superluminal particle velocities [170]. This is particularly problematic when, as is the case for LWFA, the electromagnetic waves propagate along the main grid axes [171]. There is unfortunately no fail-safe solution for Cherenkov instabilities, but there are some ways that we can attempt to mitigate their effects. These include low-pass filters

on the deposited currents (suppressing spurious high frequencies) and improved force interpolation techniques between the grid and particle positions.

The PIC code Smilei [172] was used for all LWFA simulations presented in this thesis. Smilei is an open-source and collaborative code, designed for supercomputer performance, and co-developed by experts in high-power computing and physicists using the code to simulate their experiments. The force interpolation method available to implement in Smilei is called B-TIS3 (B-translated interpolation scheme) and, although this does not remove the Cherenkov radiation, it does minimise its effects upon the simulated macroparticles [173].

Although modelling in 3-dimensions is the most accurate way to capture the physics of plasma accelerators, it is a challenging task and comes at a vast computational expense. 2D simulations offer faster runtimes and can provide good qualitative insights; however, they often suffer from numerical inaccuracies such as reduced focusing forces in slab geometries, with focusing fields underestimated by as much as a factor of two relative to 3D simulations [174–176]. An alternative approach is to decompose physical quantities, such as the electromagnetic fields and plasma density, into a series of azimuthal Fourier modes. This approach is far less computationally expensive than using 3D geometries, as each azimuthal mode (AM), of which an idealised cylindrically symmetric laser wakefield accelerator typically requires only two, represents a 2D radial slab. By including only the lowest-order modes, quasi-3D simulations can be performed that closely reproduce the key 3D physics of a cylindrically symmetric LWFA. Smilei has the capability to run in both Cartesian and cylindrical (AM) geometries [174, 177]. All simulations presented in this thesis are in an AM cylindrical geometry.

4.2.3 Resolution Tests

Before using or analysing any simulation, it is essential to test the balance of the physics we need to capture and resolve against computational cost and run time. A key starting point is the choice of longitudinal resolution (dx). This parameter significantly affects both numerical Cherenkov radiation and the ability to accurately capture fine-scale features of the plasma wake, such as the sharp density gradients at the rear of the bubble. These features are particularly critical in the self-injection process, where small numerical inaccuracies in the electric fields or bubble curvature can lead to qualitative differences in the number and timing of injected electrons.

All resolution test simulations were performed using the parameters representative of those used in the full simulation campaign. We simulate a laser wavelength $\lambda = 3.7$ µm, pulse energy $E_{\mathrm{L}} = 25$ mJ and pulse duration $\tau_{\mathrm{FWHM}} = 30$ fs, focused to a FWHM intensity spot size of $d_{\mathrm{FWHM}} = 13.6$ µm. The simulated plasma density is 2.8×10^{18} cm^{-3}.

Simulations to ascertain the required resolution were performed at 25, 50, 100, 250 and 500 grid points per laser wavelength. These were all simulated using 2 azimuthal modes and 10 particles per cell. Panels (a)-(d) in Figure 4.2 show the axial plasma

density and longitudinal electric field, E_x for this resolution scan at the moment just before injection, approximately $t = 0.57$ ps into the simulation. It should be noted that when adjusting the longitudinal resolution in this simulation, the temporal resolution is necessarily affected as a result of the Courant–Friedrichs–Lewy (CFL) condition [178]. This convergence condition is a numerical requirement which relates the spatial and temporal grids, meaning that the information can travel at most one grid cell per timestep. In AM cylindrial geometries, CFL-limited timestep Δt must satisfy:

$$\Delta t \leq \frac{1}{c}\left(\frac{1}{\Delta r^2} + \frac{1}{\Delta x^2}\right)^{-1/2} \tag{4.3}$$

Where r and x are the respective radial and longitudinal resolutions. As a result, it is not possible to extract perfectly synchronised snapshots across simulations: features will have propagated slightly different distances and may be at different stages of evolution. To compensate, we select the snapshot closest to 0.57 ps in each simulation and shift features by $c \times \Delta t$, where c is the vacuum speed of light and Δt is the difference between the desired time and the actual snapshot. The largest correction applied corresponds to $\Delta t = 0.04$ ps. We do not account for any possible variation in the evolution of features due to this timing offset.

On the scale of the full simulation window (Figure 4.2 (a) and (c)), we observe some variability in both the axial density and longitudinal fields. Changes in longitudinal resolution appear to have notably affected the structure of the second wakefield bucket. The region we are most interested in, however, is within the range of the first wakefield bucket, as it is here that the highest energy electrons will be initially captured and accelerated. In the region of $x = 110$ µm, the only significant deviation is observed in the lowest-resolution case, which is spatially offset from the others.

Zooming in on the high-density spike and the region of strongest longitudinal electric fields reveals more pronounced changes with increasing resolution. We notice a substantial change in the peak density – an increase of roughly 50% between the 100 points per wavelength and the others. The width, however, sees a corresponding decrease, implying that the total charge remains constant. We also see shifts in the locations of the field minima and density maxima. Excluding the outlier at 25 grid points per wavelength (indicating that this resolution is insufficient), these differences are on the order of hundreds of nanometres – roughly 4% of the laser wavelength, and $< 0.8\%$ of the plasma wavelength. Such small deviations are unlikely to have a significant effect on the strength of the accelerating field and could potentially even be attributed to the uncertainty resulting from the change in temporal resolution.

Another important parameter is the number of azimuthal modes (AM) we are simulating. This is of particular importance for us as the Chimera laser parameters we are using to broadly define the region of interest in our simulations are in the so-called

few-cycle regime, with a pulse duration of approximately 2.3 optical cycles. This means that there is a very strong likelihood that we will depart from the perfect cylindrical symmetry described by the 2 standard azimuthal modes. Here we test 2, 3, 4 and 5 azimuthal modes, (all simulated at a resolution of 100 grid points per laser wavelength and with particles per cell (ppc) = (AM-1) $\times$ 10), the results of which are shown in panels (e)-(h) of Figure 4.2. We note that this AM to ppc relationship is one advised by the Smilei developers when operating in azimuthal mode geometry. At the scale of the full simulation window, we see no obvious changes. Even at the 1 µm scale, the variations in peak position and magnitude are very minor, with variations of the order of 2% of the laser wavelength, and $< 0.6\%$ of the plasma wavelength.

Finally, we test the number of particles per cell (ppc), running simulations for 10, 20, 40, 80 and 160 ppc (all simulated at a resolution of 100 grid points per laser wavelength and 2 azimuthal modes). These results are presented in panels (i)-(l) of Figure 4.2. Similarly to the case of the azimuthal modes, there are no obvious deviations, even at the 1 µm scale, indicating that even 10 particles per cell is sufficient for our simulations.

Rather than speculating on the importance of peak positions, it is crucial to ensure that these small variations do not translate to unreasonable variation in the physics we are trying to simulate. We therefore also compare the evolution of accelerated electrons across these resolution scans (both the final energy spectrum and as an evolution of the 90^{th} percentile longitudinal momentum through time), the results of which are shown in Figure 4.3.

Here again, we see the largest deviation in longitudinal resolution, and little variation in the scans of the number of azimuthal modes and particles per cell. At no point do the longitudinal resolution scan results appear to converge – for each increase in resolution, we seem to produce both higher charge and higher energy electrons by the end of the simulation. This is slightly unexpected (typically, you would expect the effects of Cherenkov radiation to manifest themselves more substantially at lower resolutions), and we might start to wonder if it is possible to be too well resolved.

PIC simulations are based on an assumption of collective behaviour, achieved through statistical sampling from a distribution of macro-particles. This averaging is only valid when each macro-particle represents a sufficiently large number of real particles, typically on the order of $10^6 - 10^9$ [179]. We can try to approximate this ratio by calculating the volume per cell (in a cylindrical geometry as $V = 2\pi r \Delta r \Delta x$) and multiplying this by our simulated plasma density. In our highest resolution case, with 500 grid points per laser wavelength, this yields approximately 10^5 real particles per cell. Given 10 macro-particles per cell, the ratio of real particles to macro-particles is $\sim 10^4 : 1$, suggesting we may be approaching a regime in which individual particle effects and numerical artefacts could become significant. At lower resolutions, such as 250 and 100 grid points per laser wavelength, this ratio increases to $10^5 : 1$, potentially offering a more stable statistical basis.

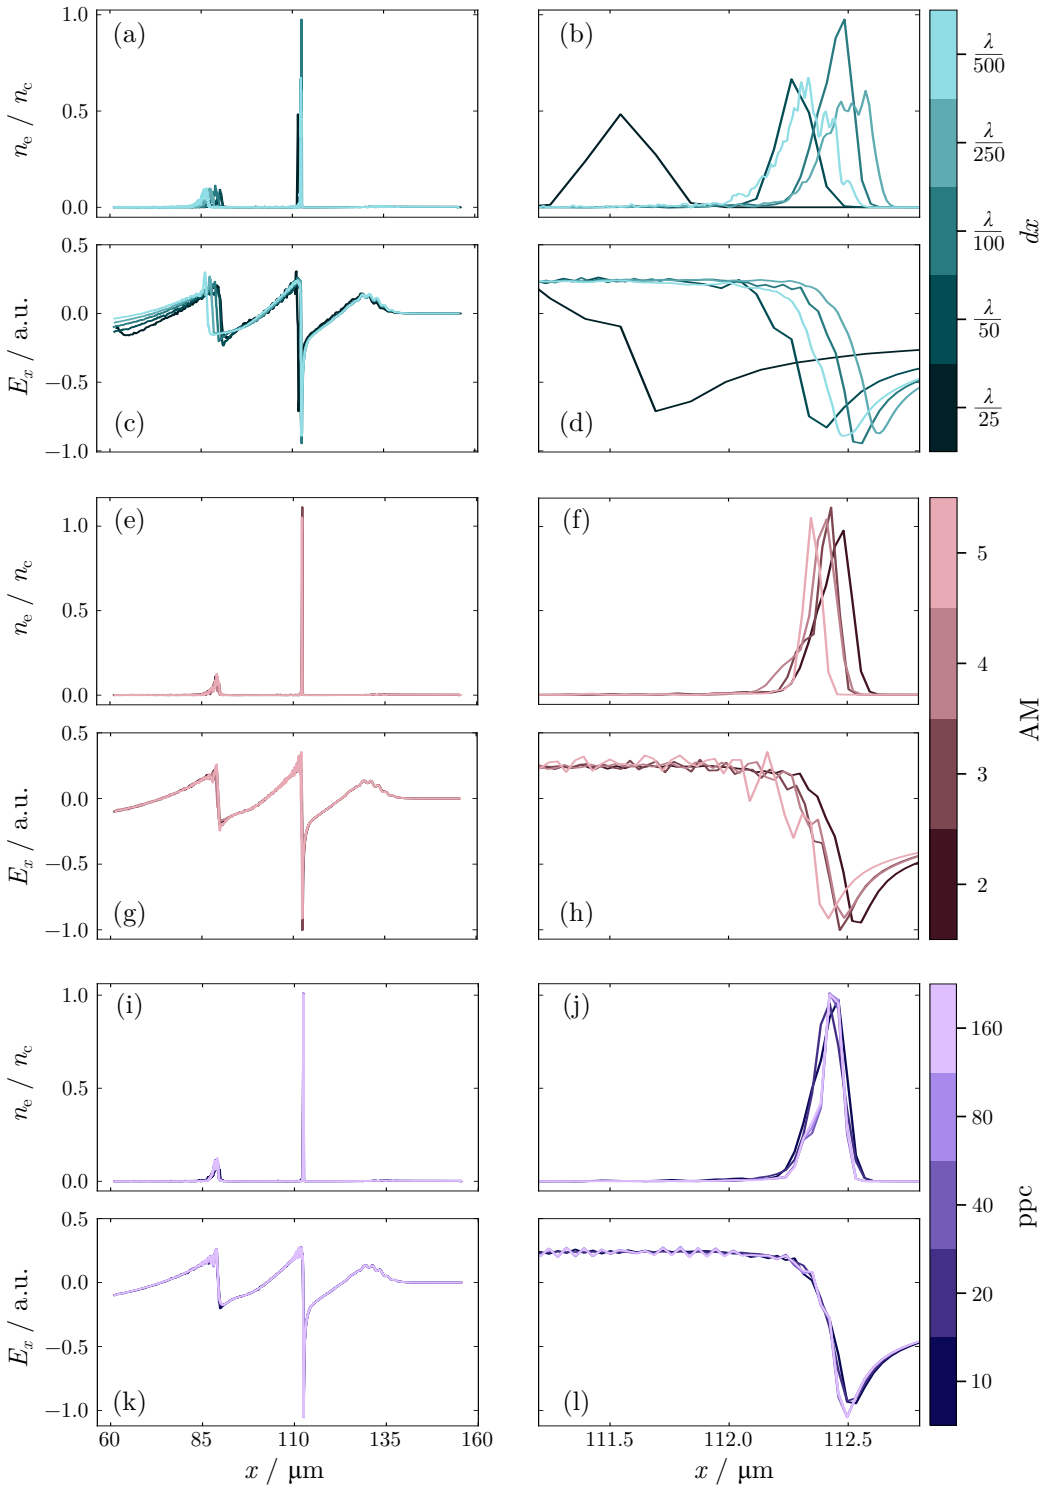


Figure 4.2: Axial plasma density and longitudinal electric field plotted for simulations run using the PIC code Smilei, with: (a)-(d) varying longitudinal resolution (dx), (e)-(h) azimuthal modes (AM) and (i)-(l) particles per cell (ppc). In each case, the left column shows the full simulation window, while the right column zooms into a region of interest. For the plasma density, this region highlights the high-density spike formed by electrons in the laser wake. For the longitudinal electric field, the zoomed region shows the strong accelerating field; when this exceeds the wakebreaking limit, electrons may be injected into the bubble cavity. The laser enters the simulation box from the left and propagates to the right. On the scale of the full window, the scans have little effect on the plasma density or field, though small differences appear in the zoomed region, particularly in the case of the longitudinal resolution test.

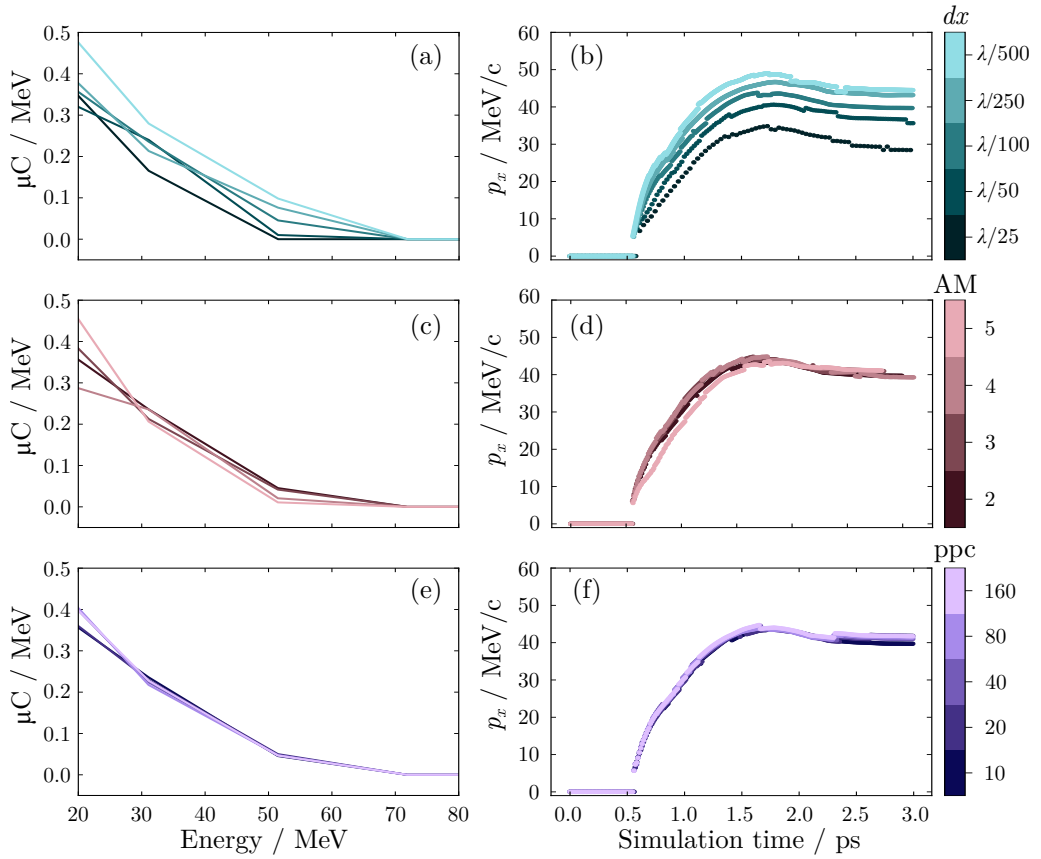


Figure 4.3: Final energy spectrum and evolution of the 90$^{\text{th}}$ percentile of all electron longitudinal momenta through time for: (a)-(b) varying longitudinal resolution (dx), (c)-(d) azimuthal modes (AM) and (e)-(f) particles per cell (ppc). The energy spectrum is obtained using the Smilei ParticleBinning diagnostic (where particle energies below 20 MeV are not considered), and the p_x momentum track is created using the TrackParticles diagnostic, which follows a random selection of individual particles throughout the course of the simulation. The AM and ppc scans show little difference, both converging early. The spectra from the longitudinal resolution tests increase in energy with increasing resolution.

Whether or not this is a real effect will have to be confirmed or refuted when we are able to conduct our first laser wakefield experiment driven by our mid-infrared laser. It is worth noting, however, that these more computationally expensive, high-resolution regimes are rarely explored in conventional PIC simulations. This is because simulating longer-wavelength drivers reduces computational costs due to the scaling of key parameters (such as resolution) with the laser wavelength. It is typical for simulations running with 800 nm driver lasers to run with only 25 grid points per laser wavelength and $\leq$ 10 particles per cell [**172**, **180**].

So far, we have considered the robustness of the captured physics relative to grid spacing, azimuthal modes and particles per cell. We now consider the computational costs. While increasing the longitudinal resolution does, as anticipated, impact the results of the simulation, it comes at a vast computational expense. These simulations were all performed using the CX1 computing cluster at Imperial College London, with simulations run across 256 CPUs. For each increase in dx presented here, the wall time

taken to run the simulation increased by 200% at each step. The lowest resolution ($\lambda/25$) took 23 minutes to run, while the highest resolution ($\lambda/500$) took nearly 35 hours, requiring a restart after surpassing the normal run time allowance (24 hours) distributed to regular users of the Imperial supercomputer. This kind of run time is untenable for the large parameter space we wish to explore. Increasing the azimuthal modes and ppc result in similar trends, although they are not quite so dramatic. By increasing the number of modes, we increase the run time by roughly 60% for each additional mode. Increasing ppc extends run time by roughly 30% per $\times$ 2 increase in particles.

As the number of modes and particles per cell appear to have very little effect on the physics we are simulating, we choose to run all future simulations using AM = 2 and ppc = 10. We choose a longitudinal resolution of 100 grid points per laser wavelength. This we feel offers a reasonable run time (of $\sim$3.5 hours per simulation) so that large parameter scans may still be undertaken, without significant impact to the accuracy of the physics output. The simulation parameters chosen for all further simulations presented in this thesis are summarised in Table 4.1.

Number of azimuthal modes	2
Particles per cell	10
Resolution ($dx \times dr$)	$\lambda/100 \times \lambda/20$

Table 4.1: The simulation parameters chosen for all following simulation results.

5 The Chimera Multi-Beam Mid-Infrared OPCPA Laser System at Imperial College London

Named after the multi-species mythological creature – "who was not a human being, but a goddess, for she had the head of a lion and the tail of a serpent, while her body was that of a goat, and she breathed forth flames of fire" [181] – the Chimera laser system is a high-power, short-pulse, multi-beam, multi-wavelength optical parametric chirped pulse amplification (OPCPA) system currently under active development at Imperial College London following its initial commissioning in 2016. Co-funded by Defence and Science Technology Laboratory (DSTL), EPSRC, EP/N018680/1 and the US Air Force Office of Scientific Research (AFOSR, FA9550-16-1-0013), Chimera was commissioned as part of the first of a series of multi-national UK-US MURI projects. Based on a MOPA (master oscillator, power amplifier) configuration, it is a complex system with a 0.8 µm few-cycle Ti:Sa front end used to seed a high-energy chain of Nd:YLF pump amplifiers as well as an 0.8 µm OPCPA stage. The proposed beamlines are the signal (3.7 µm), idler (1.5 µm) and front-end oscillator (0.8 µm), the first two of which are created during the OPCPA process. Using difference frequency generation (DFG), we initially generate a low-energy seed of mid-infrared wavelengths of 2 - 4.5 µm with a full-width at half-maximum (FWHM) bandwidth of $\sim$ 1 µm. 3.7 µm central wavelength pulses are further amplified by a series of KTA-based optical parametric amplification (OPA) gain stages, pumped at 1 µm. Of the three proposed beamlines, only the 3.7 µm signal beam is currently operational, with a current (uncompressed) energy record at 3.7 µm of 31 mJ at $\sim$ 100 fs and 128 mJ at 1.5 µm. It has an eventual 3.7 µm energy goal of $>$ 30 mJ at $\sim$ 30 fs. The author has served as the sole operator, user, and developer of the Chimera laser system since October 2021, overseeing its refinement and upgrade. The system was originally constructed by Dr. Nicholas Stuart (now at RAL), who initiated its development by reconfiguring a previous 0.8 µm OPCPA system [117].

There exist several ultra-broadband MIR systems worldwide; however, due to a lack of classical gain-storage media with broad bandwidths, OPCPA MIR generation and amplification systems all experience the same issues of poor energy upscaling capabilities. As such, our main competitors are relatively specific. We judge this based on wavelength range, pulse duration and pulse energy. The systems most comparable to our own are presented in Table 5.1.

We focus here on OPCPA systems operating in the 3 - 4 µm range. Reported pulse energies span from sub-microjoules to a few tens of millijoules, with pulse durations in the tens of femtoseconds or longer. For reference, a single-cycle pulse at 3.7 µm corresponds to 12.3 fs, meaning that these durations lie firmly in the few-cycle regime. Beyond pulse energy and duration, it is of interest to compare how MIR pulses are generated and compressed across systems: since the field remains relatively immature, no clear 'best practice' has yet emerged. These aspects offer insight not only into current performance but also into the potential higher-energy scalability and long-term reliability of each approach. With the exception of two, all comparable systems generate MIR light by extracting the idler from a final OPA stage. Our approach is distinctive in that we generate an initial low-energy, high-bandwidth MIR seed via difference frequency generation, pumped by an 800 nm OPCPA sub-system. More unusually still, Chimera employs a custom-designed temporal stretcher and compression architecture: an Öffner stretcher and Treacy grating compressor, both built with bespoke 320 l/mm gratings manufactured by Plymouth Grating Laboratory.

Despite the differences in MIR generation and pulse management, the systems at Vienna and the University of Maryland are the closest benchmarks to Chimera in terms of wavelength and uncompressed parameters. The system at RIKEN, Japan, provides the closest comparison with regard to pulse energy, duration and MIR generation, although it differs in its central wavelength. These comparisons position Chimera as a leading contender in the MIR field, with the potential, once fully compressed, to surpass the current record intensities reported in this spectral range. Chimera should, in principle, be scalable to even higher MIR energies, ultimately limited by the availability of large-aperture KTA crystals. Its 1 µm pump chain is, by design, extendable to the 10 - 50 J level through the addition of further Nd:YLF or Nd:Glass flashlamp-pumped rods, albeit at the cost of reduced repetition rate.

In this chapter, we will detail the main building blocks of the Chimera laser system, covering the overall design, the parameters required for stable day-to-day operation, investigations into the performance of individual components, and its first commissioning shots for a MIR-driven solid-target high harmonic generation (HHG) experiment. This includes the characterisation and optimisation of the Femtolasers Rainbow oscillator, the regenerative and flashlamp pumped amplifiers, individual OPA stages, and diagnostic systems to aid routine alignment – a particular challenge due to the lack of low-cost MIR camera systems. Silicon cameras, which are routinely used for near-infrared alignment, are ineffective in the MIR, since the 1.12 eV bandgap of silicon cannot be matched by a

System ID	Central Wavelength	Energy	Pulse Duration	MIR Generation Method	MIR Stretch / Compression Method
Chimera Laser System, Imperial College London, UK	3.7 µm	31 mJ (uncompressed)	100 fs (uncompressed)	DFG generated seed (MgO:PPLN) amplified by 3-stage MIR OPA (KTA)	Öffner stretch, grating compressor, 320 l/mm
TU Wien, Austria [182–184]	3.9 µm	21 mJ → 19.7 mJ	94 fs → 30 fs	Idler from KTA OPA	Self-compression in bulk YAG
University of Maryland, USA [185, 186]	3.9 µm	25 mJ	87 fs	Idler from KTA OPA	Grating compressor, 240 l/mm
RIKEN, Japan [187]	3.3 µm	31 mJ → 21 mJ	5 ps → 70 fs	DFG generated seed amplified by 2-stage MIR DC-OPA ($MgO:LiNbO_3$)	Bulk stretch in 140 mm Si, compression in 4-pass $60\times60\times100$ mm^{-3} & $60\times60\times50$ mm^{-3} CaF_2
Chinese Academy of Sciences Shanghai, China [188]	4 µm	11.8 mJ → 2.6 mJ	105 fs → 21.5 fs	Idler from KTA OPA	Compressed via HCF and grating compressor
TU Wien, Austria [189]	3.2 µm	5 mJ → 2.5 mJ	80 fs → 22 fs	Idler from KTA OPA	Compressed using HCF and CaF_2
Max-Born Institute, Germany [190]	3.1 µm	1.1 mJ	58 fs	Idler from MgO:LNB	Chirped mirror compressor
University of Duisburg-Essen, Germany [191]	3 µm	430 µJ	65 fs	Idler generated in final stage using thick $LiIO_3$ crystal	Bulk compression in sapphire
ETH, Switzerland [192–194]	3.4 µm	∼12 µJ	∼41.6 fs	Idler from APPLN OPA	Bulk compression in Si
ELI-ALPS, Hungary [195, 196]	3.2 µm	∼0.6 µJ	∼38 fs	Idler from DFG using MgO:PPLN and $LiNbO_3$	Bulk compression in 20 mm Si

Table 5.1: Summary of MIR OPCPA systems comparable to the Chimera laser system. The arrows in the energy and pulse duration columns indicate parameters before and after pulse compression.

single MIR photon. We also place Chimera in the context of the broader field, discussing methods used both in the literature and within our system for generating, stretching, and compressing broadband, high-power mid-infrared pulses that are well matched to drive high-intensity MIR experiments such as laser wakefield acceleration.

5.1 System Overview

A schematic diagram showing the building blocks of the Chimera laser system is shown in Figure 5.1. All wavelengths of the OPCPA Chimera laser system originate from a single, few-cycle, absolute phase-lockable master oscillator, centred at 800 nm, meaning that both signal and pump lines are seeded from a single source for robust passive optical lock and low timing jitter. This avoids the need for external high-speed, low-jitter electronic synchronisation or indeed a second oscillator. With roughly 2 nJ ($\sim$ 150 mW at 70 MHz) at the oscillator output, the pulse is first stretched using a pair of chirped mirrors before being split into its initial pump and seed lines. Using a single-mode, polarisation maintaining photonic crystal fiber (NL-PM-750, with core diameter 1.8 µm), roughly 60% of the 800 nm pulse is directly wavelength shifted to 1053 nm, stretched again and sent into two consecutive multi-pass, diode-pumped Nd:YLF gain amplifiers (Regen and Ring Booster), reaching first $\sim$ 2.7 mJ and then $\sim$ 50 mJ. At this stage, the repetition rate of the system is reduced to 100 Hz. These components together describe the front end of the Chimera laser system.

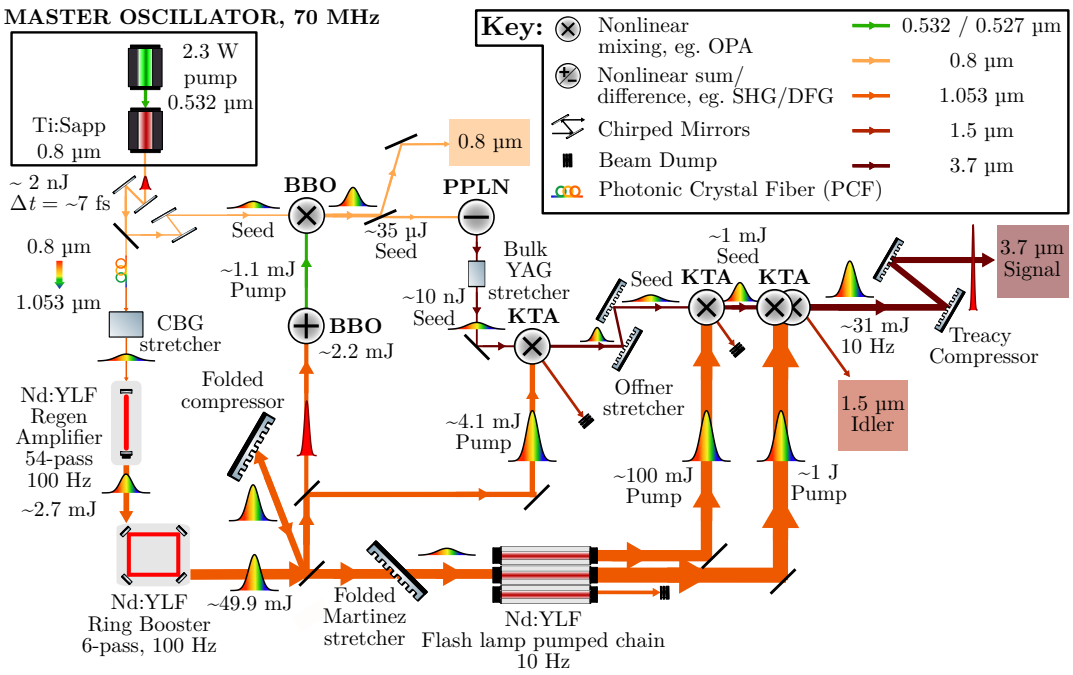


Figure 5.1: A simplified component diagram of the Chimera laser system (not to scale). The rainbow fill on the displayed pulses is used to illustrate the process of chirping the pulse within the CPA process – its wavelength components have been stretched in time. The arrow sizes are presented as a guide for the intended energy distribution of the laser pulse.

The remaining 40% of the 800 nm oscillator pulse is directed into a small 800 nm OPCPA sub-system as the seed for an OPA. This first OPA stage (using a BBO nonlinear crystal) is pumped using a small fraction of the 1053 nm Nd:YLF Ring Booster output, which is recompressed to $\sim$ 8 ps and frequency doubled to 527 nm. This provides roughly 1.1 mJ of 527 nm light into the OPA and the 800 nm seed, after being stretched through material dispersion in a fused silica glass block (thickness 50 mm), is amplified to $\sim$ 70 µJ. Currently, 50% of this light is directed into a diagnostics line, whereby the 800 nm OPA output may be monitored. This line is also where we may, in future, upgrade the system to include a second 800 nm OPA. The other 50% is directed into a difference frequency generation stage (using a magnesium oxide doped periodically poled lithium niobate crystal) from which the 3700 nm central bandwidth light is generated.

Difference frequency generation (DFG) is a common method of producing 'unusual' wavelengths of light, for example, in the mid-infrared. Creating high-intensity ultrashort mid-infrared light presents certain technical challenges: at present, a very limited number of materials are known to naturally lase in the MIR (these include certain fiber, CO_2 and a few solid-state lasers). More significantly, however, is that to produce an ultrashort pulse, you necessarily require a sufficiently broad bandwidth, as given by the time-bandwidth product (Equation (2.42)). The broad wavelength range that this requires is not one in which typical gain storage materials are able to operate, particularly in the MIR, and it is this which is the main limiting factor to the production of ultrashort MIR light. DFG is an elegant solution to this problem, as nonlinear crystals can provide the broad bandwidths required for an ultrashort pulse. Available bandwidth will depend on the transparency and phase matching criteria of the nonlinear crystal, as well as the crystal length and the refractive index the pump light experiences in the crystal.

At the time of writing, we believe that our system is unique in the way we generate, stretch, and compress our MIR pulse. In the Chimera laser system, we use MgO:PPLN (magnesium oxide doped periodically poled lithium niobate) – an artificially engineered, highly efficient, nonlinear, quasi-phasematching crystal – to generate our MIR seed light through an unseeded DFG process, equivalent to spontaneous parametric down-conversion. This provides us with roughly 10 nJ of MIR seed, centred at 3700 nm, with over 1000 nm of available (temperature-dependent FWHM) bandwidth. As neither the signal nor idler fields are externally seeded, however, the generated MIR pulse originates from vacuum fluctuations and therefore does not possess a deterministic phase relationship with the 800 nm pump pulse. While energy and momentum conservation strictly define the generated frequencies, the absolute optical phase of the MIR field will be random from shot to shot. Compared to other reported MIR systems, this method is unusual, likely due to the later challenge of amplifying the full bandwidth of our MIR light. All bar two (RIKEN, Japan [187] and ELI-ALPS, Hungary [195, 196]) reported systems worldwide use a MIR idler produced in an OPA processes, and typically this is done in the final stage of their systems [182–186, 188–196].

5.2 The Femtolasers Rainbow Oscillator

The master oscillator that generates the initial laser pulse in the Chimera laser system is a Kerr-lens modelocked Ti:Sapphire (Ti:Sa) oscillator (Rainbow, Femtosource GmbH). Ti:Sa systems are widely regarded as the 'workhorses' of short-pulse laser oscillators: they have the capability to realise the shortest cycle (sub two-cycle) pulses for the highest comparative peak powers and their bandwidth spans over an octave of available frequency ranges in the near-infrared (NIR) [119]. A schematic optical diagram of the Chimera oscillator is shown in Figure 5.2.

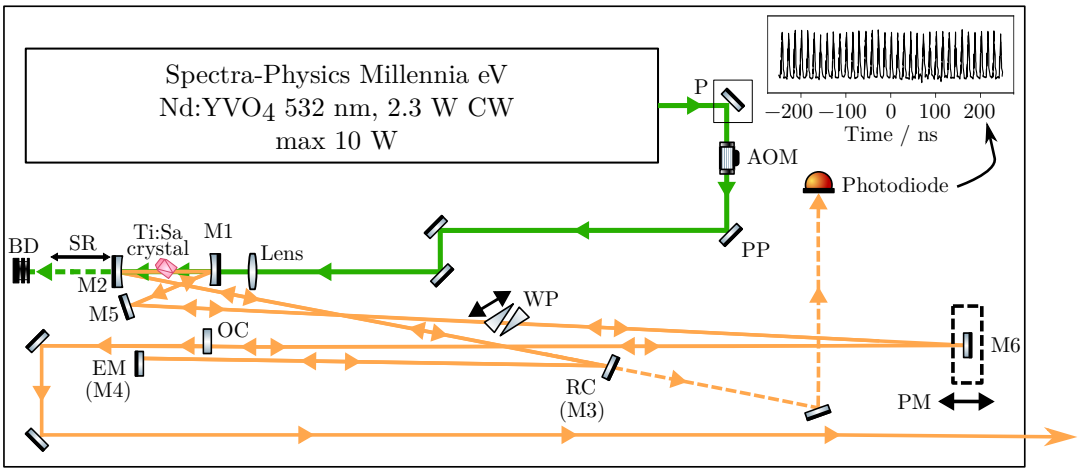


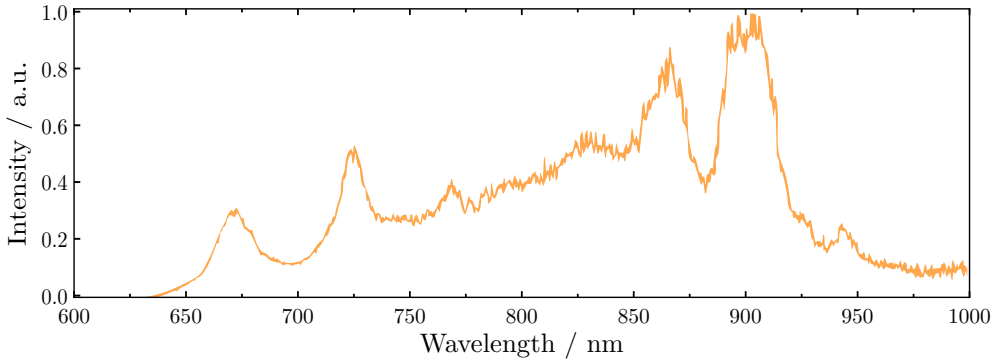
Figure 5.2: The Femtosource Laser cavity. P: Periscope; AOM: Acoustic Optical Modulator; PP: Pump Pointing; Ti:Sa: Titanium:Sapphire; SR: Stability range; BD: Beam Dump; RC: Rear Cavity Mirror; EM: End Mirror; WP: Wedge Plates; PM: Picomotor-driven Mirror; OC: Optical Coupler. The orange and green lines represent the 800 nm signal and the 532 nm pump beams, respectively. The modelocked oscillator pulse train (at 70 MHz) collected by the photodiode is also included, with the y (voltage) axis removed for clarity. Diagram not to scale.

The oscillator system has undergone several upgrades since initial commission in 2006, the most recent change being on 8th August 2024 with the replacement of the gain medium – a thin ($\sim$ 2 mm) Ti:Sa crystal cut at Brewster's angle. The Ti:Sa is optically pumped by a continuous wave (CW) 532 nm beam: a frequency-doubled, Q-switched, diode-pumped laser (Coherent Verdi V5), with a solid-state gain medium Neodymium-doped Yttrium Orthovanadate (Nd:YVO_4) and a maximum power output of 10 W. Both the oscillator and pump laser are housed in a single water-cooled unit with a shared baseplate for optimised stability and temperature control. Prior to its replacement, the Ti:Sa crystal was pumped at 2.3 W, with a pulsed power output of 90 mW at 800 nm and FWHM bandwidth of $\Delta\lambda = \sim 130$ nm. This translates to a transform-limited pulse duration of $\Delta t = \sim 7$ fs. As of the crystal replacement, we see routine pulsed power output of $\sim$ 150 mW at 800 nm, whilst retaining an approximate FWHM bandwidth and transform-limited pulse duration of $\Delta\lambda = \sim 130$ nm and $\Delta t = \sim 7$ fs. The repetition rate of the system (70 MHz) is locked to an external rubidium atomic clock set to match the Cerberus laser system, a hybrid OPCPA-Nd:Glass multi-beam, multi-target area system with which Chimera shares a lab. This synchronisation allows for potential

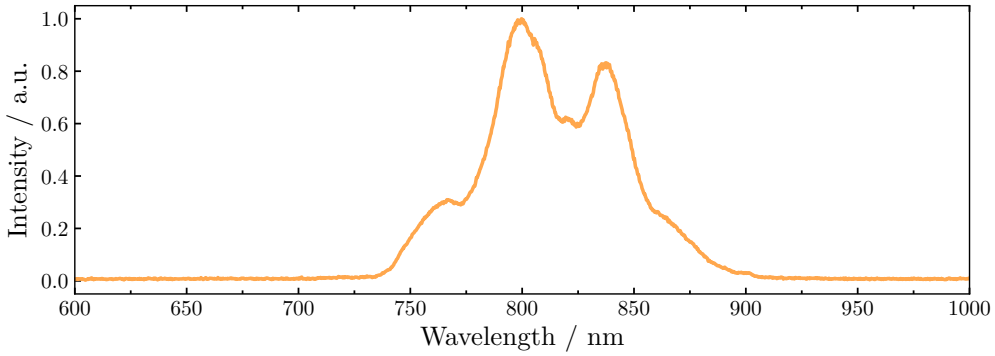
coupling or phase-locking between the two systems if desirable in a future upgrade or experiment – for example, a two-laser system pump-probe geometry, or the pumping of a large MIR OPA gain stage with a 1053 nm beam from Cerberus. A carrier envelope phase (CEP – see Section 2.6) stabilisation unit was in place in earlier versions of the Chimera oscillator (see [117]). Although we retain the capability for CEP control, it is not currently in operation – this is a candidate for a future upgrade when starting to deliver few-cycle, high-intensity light into an experiment; however, it is unclear if CEP control of the 800 nm seed will translate into a similar absolute phase lock of the MIR seed generated by DFG.

The oscillator cavity is configured with a typical X-shaped design. Modelocking is initiated by manually perturbing the cavity End Mirror (labelled EM/M4 in Figure 5.2). Given that we rely on the optical Kerr effect to create our pulsed laser beam, gradual degradation of the Ti:Sa crystal is expected. This means that over time, the crystal will need to be shifted in order to maintain optimal operation. The curved mirrors M1 and M2, with respective focal lengths f_1 and f_2, focus the laser mode within the cavity and the distance between them defines both the spatial mode and the so-called 'stability range (δ)' – the range of distances over which the cavity is stable (see Section 2.8). This can be adjusted externally, and its extrema (confined within four limits, δ_{0-3}) are representative of the cavity's stability requirements. Modelocking may not take place at either δ_0 or δ_3: it is found close to either δ_1 or δ_2, with the most robust pulsed operation found in the range of δ_1. As the cavity is configured using a virtual hard aperture (a geometrical or optical configuration to achieve mode-selective behaviour, as opposed to a physical obstruction), this is expected [119]. Although modelocking can be achieved across the full range of the δ_2 stability region, operation in this regime is typically unstable, and once modelocking is achieved, it regularly falls out of its pulsed operation. If achieved, the pulsed mode at δ_2 is often accompanied by the appearance of a narrow CW spike in the spectrum. This is usually the result of CW power exceeding pulsed power, allowing a residual CW mode to co-propagate with the pulse as it circulates in the cavity [119]. The measured spectra at the output of the Chimera oscillator in 2013, as well as in August 2024 (at both stability positions following the replacement of the Ti:Sa crystal), are shown in Figure 5.3.

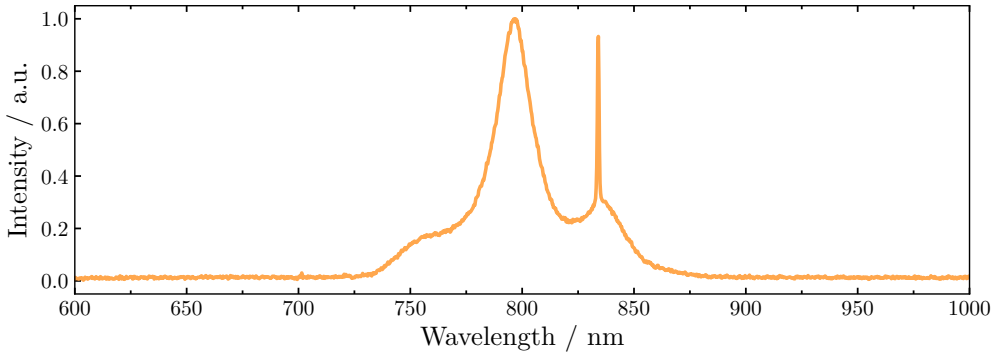
Stable and consistent oscillator performance is crucial as it dictates the behaviour of all subsequent seed generation and amplification stages. Output power and spectrum are the most straightforward parameters to measure when monitoring system performance, as these are near CW values, while characterising the pulse duration itself requires a more complex diagnostic, such as autocorrelation. With increasing pump power, we observe only a small increase in the output modelocked power. This is accompanied by a slight broadening of the spectral bandwidth (see Figure 5.4a), likely due to increased nonlinear phase accumulation via the Kerr effect. However, this advantage comes at the cost of reduced conversion efficiency (Figure 5.4b). Indeed, Kerr lensing strength is known to have a maximum point of effectiveness [119], and operating at higher pump powers



(a) Chimera oscillator spectrum taken in 2013 by Katalin Mecseki [117]. This spectrum is the output of the Femtolasers Rainbow oscillator with carrier envelope phase (CEP) stabilisation and corresponds to a 6.2 fs transform-limited pulse. The pulses leaving the oscillator are actually $\sim$ 15 fs due to accumulated material dispersion in the CEP detection unit.



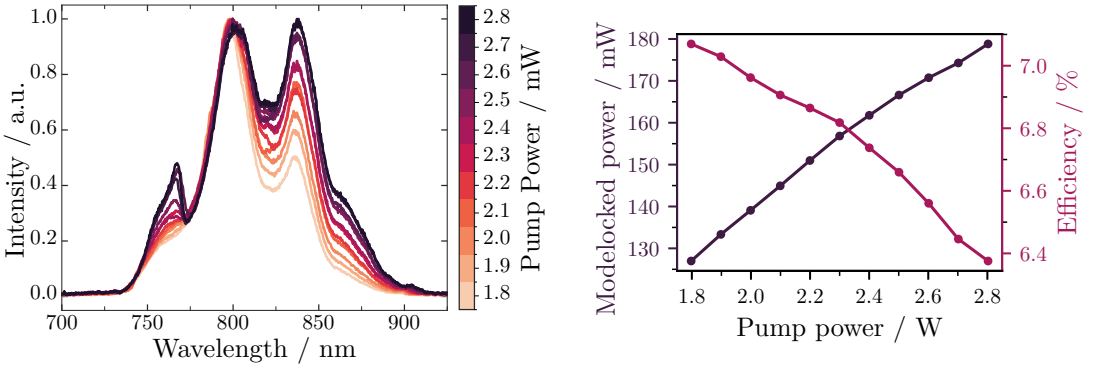
(b) Chimera oscillator spectrum when operating close to stability bound δ_1, taken on 28/08/24 by Annabel Gunn after replacing the Ti:Sa crystal. This corresponds to a $\sim$ 7 fs (Gaussian FWHM) transform-limited pulse. This spectrum is the output of the Femtolasers Rainbow oscillator with no CEP stabilisation. There is therefore no accumulated material dispersion due to the CEP detection unit.



(c) Chimera oscillator spectrum when operating close to stability bound δ_2, taken on 28/08/24 by Annabel Gunn after replacing the Ti:Sa crystal. The narrowband spike at $\sim$ 835 nm is due to the CW lasing still present within the cavity.

Figure 5.3: Comparison of Chimera oscillator spectra. The spectra were recorded using an OceanOptics HR4000 spectrometer.

can accelerate crystal degradation and reduce system lifetime. A practical compromise must therefore be reached. Pumping at 2.3 W, we obtain an output modelocked power of 156.8 mW, giving a conversion efficiency of 6.8%. Although increasing the pump power will naturally provide greater modelocked power and bandwidth, we find 2.3 W is sufficient for our purposes.



(a) (b)

Figure 5.4: Measured Chimera oscillator spectrum and pump power VS modelocked power efficiency plot.
(a) A spectrum, taken with an OceanOptics HR4000 spectrometer and overlaid at increasing pump powers (1.8 mW - 2.8 mW). We see that while the bandwidth increases with larger pump powers, the spectrum retains roughly the same shape, and this effect appears to be marginal to the overall performance of the system.
(b) A plot to detail the respective modelocked power and conversion efficiency for increasing pump powers (1.8 mW - 2.8 mW). Here, we see a decrease in conversion efficiency for increasing pump powers.

Given that this oscillator allows the optical locking of all downstream pump and seed pulses (ensuring that their relative timing is defined purely by optical path length rather than by electronic jitter or timing delays), maintaining a stable oscillator repetition rate is essential. Chimera global timing is controlled by a rubidium atomic clock module (Stanford Research Systems FS725, Rb frequency standard) from which all active pulse switching and synchronisation signals are derived. This means that, as all timing is fixed relative to the atomic clock, any mismatch between the oscillator repetition period and the system's expected timing (as a result of thermal expansion of the cavity as it warms up) can cause the temporal gain window of an OPA to be missed.

There are two ways we account for this on the Chimera laser system. Firstly, as of November 2024, we choose to keep the oscillator on and running for several days at a time, shutting it down over the weekend or extended periods of downtime only. As the oscillator is cooled using a single closed-loop chiller, this means that once the cavity has reached thermal equilibrium in operation, it becomes largely immune to external lab conditions such as room temperature fluctuations, humidity, or variations in mains water supply. As such, there should be limited thermal drift in the oscillator cavity. Figure 5.5 presents the oscillator spectra taken over four consecutive days, showing negligible change to its spectral shape over this period. To correct for residual thermal drift and maintain long-term frequency stability, however, we also explicitly monitor the

pulse train frequency inside the cavity using a frequency counter (Agilent HP 53131A). Any deviation from our frequency set-point of 70,004,000 Hz is accounted for by using an intra-cavity mirror mounted on a piezo-actuated picomotor stage (Thorlabs PAS009), which compensates the change of frequency by adjusting the length of the cavity (see Figure 5.2).

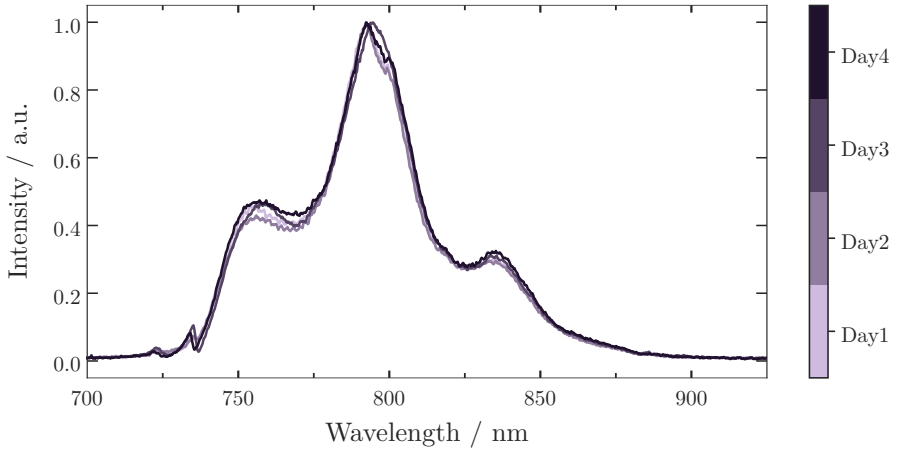


Figure 5.5: The Chimera oscillator spectrum over a period of four consecutive days (30/06/25 - 03/07/25). There is negligible change in the spectra, justifying the decision to keep the oscillator on for multi-day periods.

After being chirped (stretched), the initial oscillator pulse is split into two beamlines. 40% of the pulse is directed into a small 800 nm OPCPA sub-system which, after amplification, will become the seed from which we generate our mid-infrared (see Section 5.5.1). The other 60% is wavelength shifted, from 800 nm to 1053 nm, so that it can provide a seed for the pump-chain for all following OPA processes. This wavelength up-shifting is achieved by employing the nonlinearity of a photonic crystal fiber, as described in Section 5.3.

5.3 The Photonic Crystal Fiber

In the Chimera laser system, a nonlinear photonic crystal fiber (PCF) is used to generate a seed pulse at 1053 nm, enabling precise optical synchronisation with Nd:YLF pump sources. It is made from fused silica and is comprised of a triangular array of microscopic air-filled holes, surrounding a central solid core running through the entire length of the fiber (Figure 5.6).

When ultra-short, high-energy pulses propagate through a PCF in the vicinity of its zero-dispersion wavelength (ZDW), the optical spectrum can undergo significant broadening as a result of nonlinear effects. The large refractive index contrast between the silica core and the surrounding air-hole cladding provides strong transverse confinement of the optical mode, such that, under appropriate conditions and in the anomalous dispersion regime, it is possible to launch a soliton – a pulse which maintains its temporal profile while propagating through a nonlinear medium. As the soliton propagates, it

experiences self-phase modulation due to the Kerr nonlinearity, as well as a continuous frequency downshift arising from the intra-pulse Raman scattering (the soliton self-frequency shift). The extent of this spectral evolution depends on the fiber length and dispersion properties (which are adjusted by changing the diameter or lattice pitch of the cladding), but can also be controlled, to some degree, by adjusting the input pulse energy.

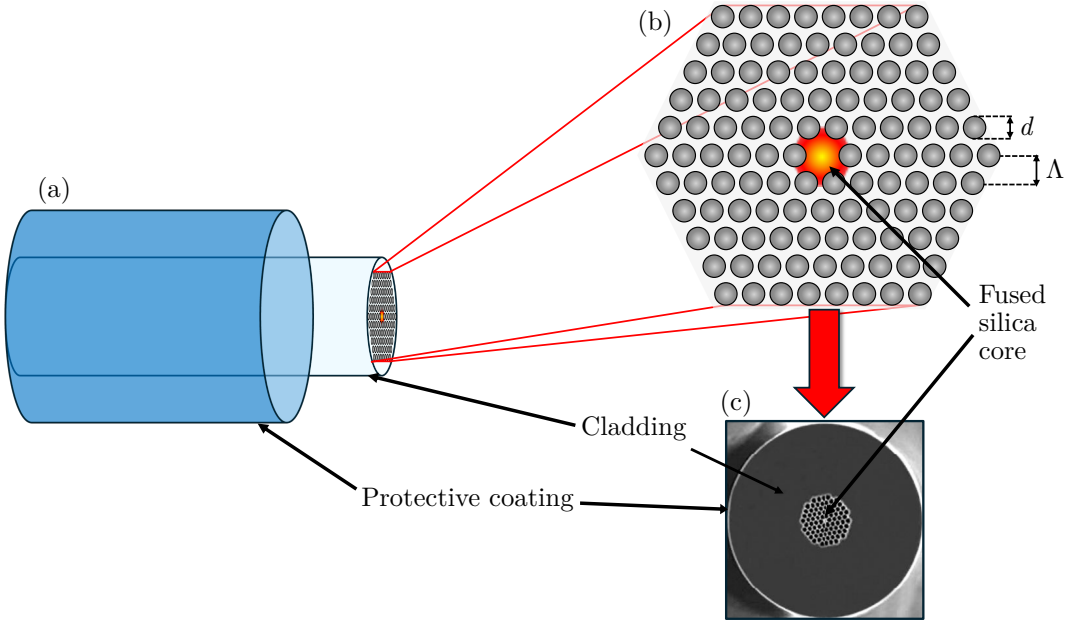


Figure 5.6: The Photonic Crystal Fiber (PCF).
(a) An image to illustrate the anatomy of a PCF.
(b) Shows the structure of the core and the surrounding triangular array of microscopic air-filled holes. The indicated variables Λ and d are the so-called lattice pitch (hole-to-hole distance) and air hole diameter parameters. For a NL-PM-750 fiber, estimated using an SEM image in [197], Λ and d are roughly 1.26 µm and 0.87 µm, respectively.
(c) A microscope picture of the end of a NL-PM-750 fiber, taken from the manufacturer's datasheet [198].

The fiber currently in use is a NL-PM-750 single-mode, polarisation-maintaining PCF, with a ZDW of 750 nm and core diameter of 1.8 µm [198]. It is seeded by a Femtolasers Rainbow oscillator (see Section 5.2) at 800 nm with a target conversion wavelength of 1053 nm. This wavelength is chosen to match the gain bandwidth of existing Nd:YLF regenerative amplifiers, as well as several stages of Nd:YLF flashlamp-pumped amplifiers further down the system.

The fiber is 37 cm long, with input light coupled in using a high numerical aperture (NA=0.4) microscope objective. Dispersion within the objective is compensated for by six bounces across a pair of chirped mirrors (Thorlabs DCMP175) placed before the objective, contributing $\phi'' = -175$ fs^2 per bounce. The length of fiber for optimum conversion was numerically modelled by Dr Katalin Mecseki [117] who, for a 34 kW (peak power), 6.2 fs input pulse, found the ideal length to be 29.5 cm. A longer-than-required length of fiber allows an extra degree of tolerance within our setup – changes in input power can be managed simply by adjusting the alignment of the input beam

into the fiber core, and if the ends need to be re-cleaved due to damage, the overall length of the fiber is still sufficient to reach our goal pump wavelength. The fiber is mounted on a high-precision manual XYZ stage, covered by a plastic enclosure to protect from dust-induced facet damage at the fiber ends. Coupling can be directly adjusted using this stage, but, for day-to-day alignment, piezo-electric adjustment of the pointing mirror allows for smoother, finer control. Poor coupling into the fiber's guided mode reduces nonlinear interaction length and intensity, resulting in inefficient conversion and reduced output at the desired wavelength. At 37 cm, we have sufficient length to reach wavelengths beyond our 1053 nm target, providing operational headroom and tolerance for fluctuations in input power, coupling, or pulse duration. An example spectrum of the PCF output, alongside the input oscillator spectrum, is shown in Figure 5.7a. This is accompanied by a photograph of the active PCF *in situ*, in Figure 5.7b.

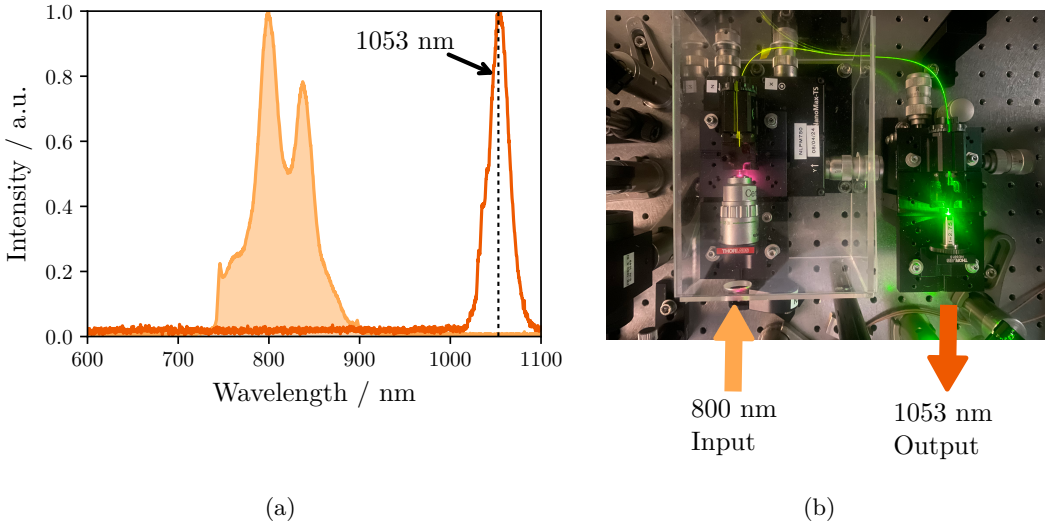


Figure 5.7: Experimental output of the Photonic Crystal Fiber.
(a) Spectra showing the experimental output of the NL-PM-750 PCF (dark orange), with the 1053 nm target signal central wavelength marked in black, and the seeding Rainbow oscillator (light orange, filled). These were both measured using an OceanOptics HR4000 spectrometer on 05/04/24. Note that the intensities of the oscillator and fiber are individually normalised.
(b) A photograph, taken with an iPhone 11, of the working NL-PM-750 PCF *in situ*. When the fiber is optimised for maximum conversion to 1053 nm, we observe significant 'apple green' leakage light. Bad conversion would be indicated with longer wavelengths of visible light. No leakage indicates the input light is not correctly coupled into the fiber core. The red light at the input is infrared light, picked up by the phone. It is not visible to the naked eye.

To ensure proper alignment of the seed pulse into subsequent amplification stages, the far-field output of the photonic crystal fiber is monitored using the residual, or 'waste', light transmitted through an eight-pass chirped Bragg grating (Optigrate CBG-1053-90-M, $10 \times 10 \times 28$ mm). This grating serves to stretch and chirp the pulse before amplification, and the rejected off-wavelength components (which are not efficiently diffracted in the reverse direction to the input beam) provide a reference for alignment diagnostics. This stretcher configuration is non-trivial and is therefore shown in Figure 5.8.

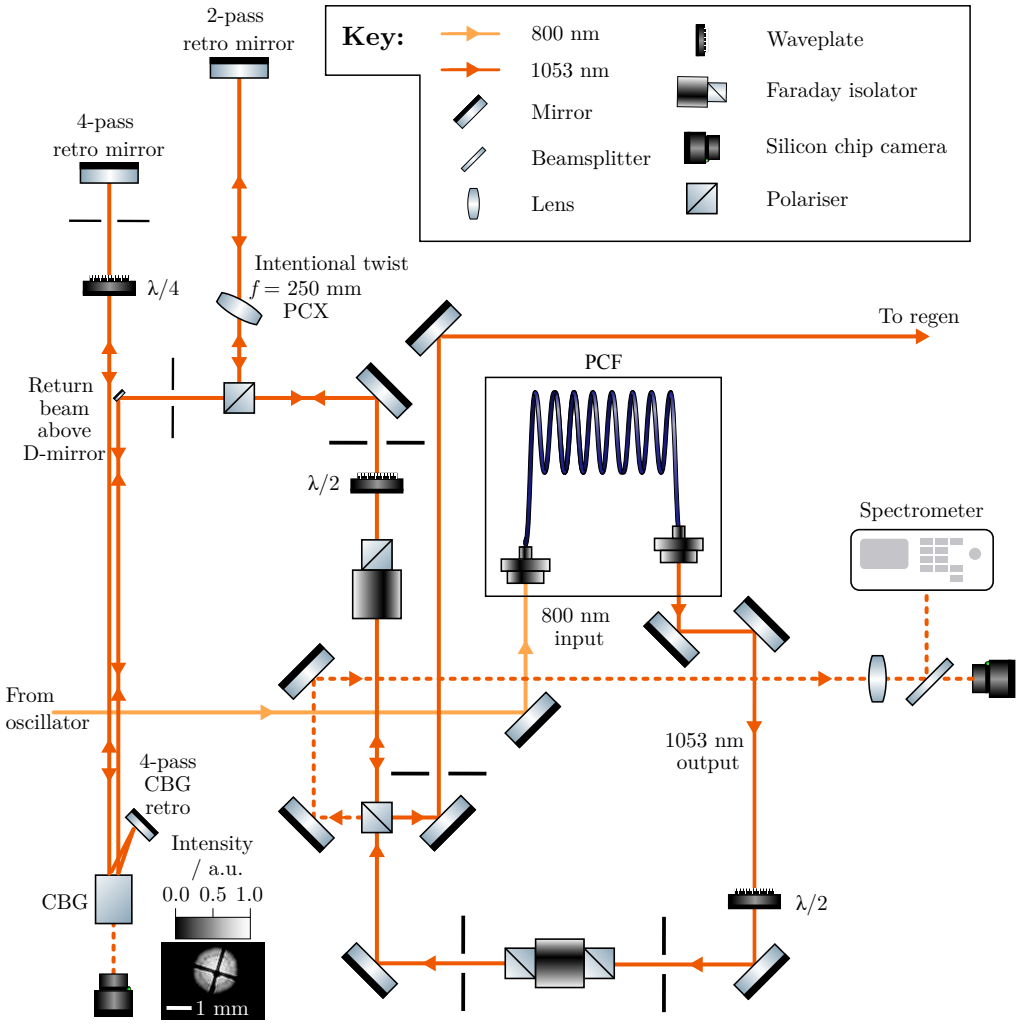


Figure 5.8: A simplified top-down view of the optical path of the 8-pass chirped Bragg grating (CBG). The horizontal offset shown in the CBG beam-path is used to illustrate a vertical offset on the real system to allow for the geometric multi-pass. Using a quarter-waveplate and polariser combination, the CBG performs 8 stretches of the 1053 nm pulse. There is an intentional $\sim 27°$ twist on a lens in the 2-pass stage. This is to correct an astigmatism collected into the beam profile before it is sent to be amplified in the regen. To aid alignment, the 'waste' light which has been transmitted through the CBG is collected on a standard silicon chip CCD camera. There is an iris aperture directly in front of the camera, as well as a crossed pair of fine tungsten wires hand-glued to the camera module mount into a cross hair to aid precise alignment. Optimal pointing of the 1053 nm seed into the regenerative amplifier is obtained when the beam from the stretcher system is centred onto the crosshairs. Diagram not to scale.

Although this stretcher system is, at first glance, straightforward, careful alignment is required in order to send the correct pulse, with the correct stretch factor, into the regen to be amplified. The CBG has a design wavelength of 1053 nm with 97% diffraction efficiency and is optimised for maximum diffraction efficiency in the reverse direction (near Littrow) at 1053 nm. To achieve the maximum stretch, the pulse interacts with the CBG 8 times, with a stretching factor of ~ 160 ps/nm per pass. As the CBG can be used in both polarisation states, this multi-pass geometry is achieved through both geometric and polarisation adjustments to the beam. This makes alignment non-trivial as

multiple passes can be difficult to overlap and challenging to individually detect. When the system has been carefully aligned through each successive pass through the CBG, day-to-day optimal alignment is achieved when the rejected (wrong wavelength/residual transmitted) beam is centred on the crosshairs of the CBG rejection camera, as shown in the Figure. Fine adjustments to the beam position can be made using the X and Y output controls of the fiber mount. This alignment simultaneously optimises both the output spectrum captured by the fiber spectrometer and ensures optimal pointing and spatial overlap of the 1053 nm seed into the regenerative amplifier. Once stretched in the CBG, the wavelength-shifted pulse is ready for amplification using traditional gain-storage amplifiers.

5.4 Gain-Storage Laser Amplifiers

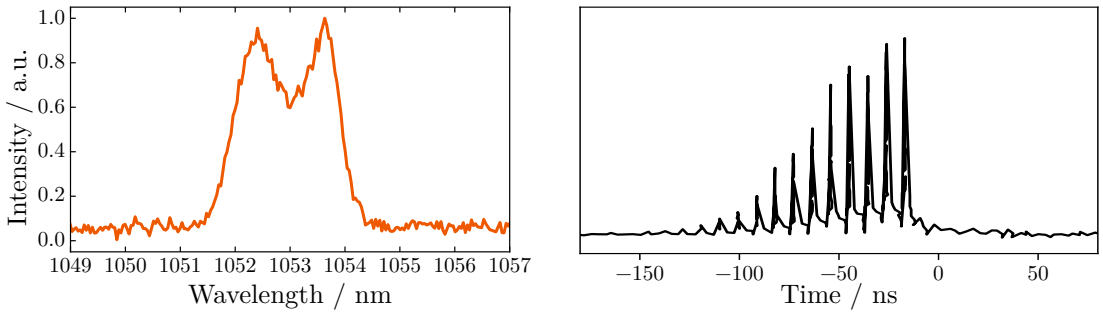
Whilst Chimera is broadly an OPCPA system, it does nevertheless use several types of conventional laser gain amplifiers to amplify pump light before it reaches an OPA stage. Two of these amplifiers are diode-pumped Nd:YLF (Neodymium-doped yttrium lithium fluoride) laser heads used within optically switched cavities to amplify propagating 1053 nm light in multiple passes. YLF is a naturally birefringent medium with common, polarisation-dependent, laser transitions at 1047 nm (for π polarisation) and 1053 nm (for σ polarisation). Operation at 1047 nm offers higher gain, but the transition is more strongly thermally dependent, making it prone to thermal lensing and less suitable for high-repetition-rate operation [199]. It is also more sensitive to pump-induced birefringence, which increases the noise floor of the amplifier. Operating at 1053 nm not only reduces these effects but also allows an extra degree of design freedom – it matches well with Nd:phosphate-glass amplifiers. This compatibility allows capacity for further amplification using Nd:glass heads, should that be desirable in a future high-energy upgrade.

The two diode pumped Nd:YLF amplifying stages consist of a 54-pass linear regenerative amplifier (abbreviated to regen) operating at 100 Hz, pumped continuously at 16.8 A, and a custom 6-pass ring regenerative amplifier (abbreviated to ring, ring regen or ring booster) made up of 5 diode modules, each operating at 100 Hz and pulse pumped at 80 A, 400 µs. These amplifiers are crucial front-end components of the system, without which none of the later nonlinear optical stages can be pumped to generate ultra-fast high-intensity MIR light.

5.4.1 The Linear Regenerative Amplifier

The linear regen is the first amplifier to interact with the 1053 nm beam. Arranged in a linear geometry, pulses are actively switched in and out of the laser cavity using a quarter-waveplate and a fast BBO Pockels cell (Leysop) to determine the number of round trips. BBO is preferred here over other crystals such as KDP for this application due to its fast switching speed, high damage threshold and low piezoelectric ringing, as well as its wide transparency range, with low absorption at 1053 nm [200]. The cavity

is 135 cm long with a round-trip time of $\sim$ 9 ns, ensuring compatibility with the high voltage supply switching of the Pockels cell driver (Bergamn BME), which has a rise time of $\sim$ 3 ns. The regenerative amplifier itself is an 83 mm long, 3 mm diameter, diode-pumped Nd: YLF module (Northrop Grumman RBA35-1C2) with 0.9% doping. The module is set into a commercial mount and is radially CW pumped by 15 sets of diode arrays, themselves built into three water-cooled diode bars. Run at 100 Hz and pumped at a CW diode current of 16.8 A, it is an extremely high gain system of $\mathcal{O}(10^9)$ and, when running well, produces a typical output energy of 2.8 mJ with 0.3% RMS (root mean square) energy stability, which is critical for the performance of subsequent cascaded nonlinear gain stages. This energy output includes an unwanted long-pulse background component of Q-switched ASE (amplified spontaneous emission) from the amplifier of $\sim$ 0.26 mJ. To achieve the optimal energy stability, the regen is driven into saturation, such that the available laser gain becomes comparable to the loss within the cavity. This is essential for the optimal performance and energy stability of the OPA stages later on in the system. The spectral profile, measured using a Thorlabs LC100/M - Smart Line Camera, is shown in Figure 5.9a and the photodiode trace showing the regen gain roll off after successive trips is shown in Figure 5.9b.



(a) (b)

Figure 5.9: Output of the 54-pass diode pumped Nd:YLF linear regenerative amplifier. (a) The spectral profile of the 1053 nm laser pulse at the linear regenerative amplifier output. Gain narrowing, accumulated over many round trips, is compensated for using a low finesse, low reflectivity intracavity etalon. It is the etalon which produces the central dip in the spectrum. (b) The linear regenerative amplifier photodiode trace collected from leakage light in the regen cavity. This trace is monitored on a fast oscilloscope. Each successive trip sees small amounts of gain until the available gain becomes comparable to the loss within the cavity, causing the gain to saturate. The y (voltage) axis is not physically meaningful and has been removed for clarity.

Due to the finite bandwidth of YLF, a pulse will experience gain narrowing when subject to such high gain in an amplifier. To maintain the bandwidth required for our ideal pulse duration (whilst maintaining a relatively long pulse duration within the amplification process to suppress unwanted self-phase modulation and other damage mechanisms), we can modulate the narrowed spectrum, either before or during the amplification process. In our case, we use a thin (500 µm), low finesse, low reflectivity intra-cavity etalon, made using fused silica and a custom 2-layer TiO_2 / MgF_2 partial (R $\sim$ 1%) anti-reflection coating. These exhibit a high enough damage threshold to be

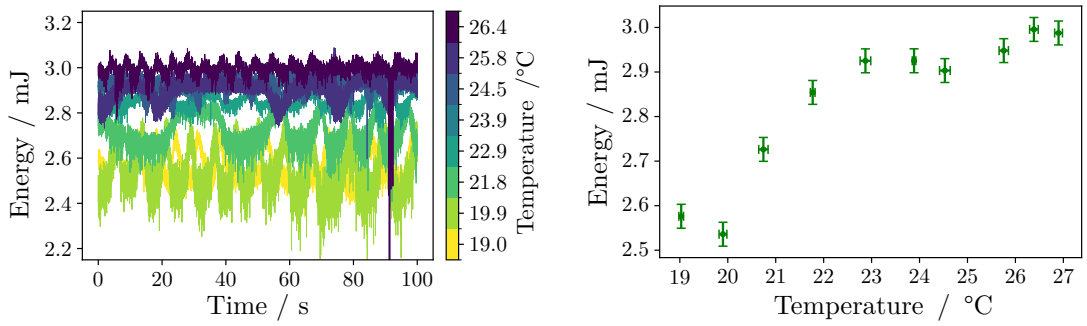
used within the amplifier cavity, allowing for spectral shaping, controlled by angling the etalon relative to the cavity axis, without too much compromise on seed energy [201].

There are multiple modes in which an etalon can be used within a system like this, depending on where the etalon transmission is placed relative to the gain profile. The authors in [202] chose to arrange the etalon transmission maximum at the gain maximum, thus reducing efficiency, and narrowing the overall bandwidth by introducing extra loss into the cavity. Described and numerically modelled in [117, 201], we take the opposite approach and instead place the etalon transmission minimum at the gain centre, thus maximising the usable bandwidth rather than minimising it. In addition to combating gain narrowing and reducing B-integral, this technique also gives much greater control over the temporal shape of the pulse, allowing us to select a more flat-top (as opposed to Gaussian) spectral shape. This is advantageous in an OPCPA system such as ours, as Gaussian pulses are more likely to result in large amounts of parametric fluorescence in an OPA process, thus reducing the amount of 'useful' pump energy that may be coupled into the seed. More crucially still, if the leading edge of a Gaussian pump pulse is not well optimised to the bandwidth of the seed, this parametric fluorescence can result in a laser pre-pulse. In contrast, a flat-top spectral profile supports a more uniform temporal gain envelope, better matching the seed bandwidth throughout the pump duration and significantly reducing the likelihood of generating an unwanted pre-pulse resulting from parametric fluorescence.

For maximum efficiency, the pump diode output wavelength must match the absorption characteristics of the chosen laser medium. This is tuned using a closed-loop chiller, whose primary function is to remove waste heat from the amplifier heads. The operating temperature recommended by the supplier is 20 - 35 °C; however, in practice, this should be specifically set so that the temperature-dependent wavelength emitted by the diode module matches the absorption peak of the gain medium [203]. This temperature dependence is one we expect: it is well known that the bandgap of the semiconductor junction narrows with increasing temperature, resulting in a redshift of the emitted wavelength [204]. The implication of this is that as the temperature and bandgap of the semiconductor junction within the module change, so too will the wavelength that it emits. For typical GaAlAs-based laser diodes, this wavelength shift is approximately 0.25 nm/°C, although the exact coefficient depends on the diode's composition, age in service, and structure [203]. This effect provides a practical method to fine-tune the diode emission to match the gain medium's absorption cross-section.

The closed-loop chiller in use is a Lytron LCS20 Liquid-to-Liquid Cooling System. It does not contain any heating elements, but rather the user must input a set-point temperature to which the chiller will cool, meaning that the warm-up time is set by the diodes depositing heat into the loop. By adjusting the set-point of the chiller, we may conduct a temperature scan in order to determine the appropriate cooling temperature and to maximise its amplifying efficiency in the system. Figure 5.10 shows data, recorded

using a GENTEC QE12SP energy meter, taken over 10,000 shots (100 s at 100 Hz) at each temperature step.



(a) (b)

Figure 5.10: Output energy of the Nd:YLF linear regenerative amplifier whilst varying diode temperature, recorded over 10,000 shots at each step using a GENTEC QE12SP energy meter. Over this scan, the diodes are driven at a constant current of 16.8 A. The temperature was measured using a temperature probe inserted into the chiller loop and recorded manually every 10 s.

(a) Time-varying output energy plotted at each temperature step. Two temperature steps have been removed for clarity – all steps appear in (b). The cyclic variation seen at all temperatures is not one which is currently understood.

(b) The mean energy output at each temperature step. The errors in y are the standard deviation in the recorded output energy. The errors in x are the standard deviation of 10 recordings of temperature at each temperature position.

From this data, we see a very clear increase in amplifier performance as the temperature of the diode module is increased. As the output energy tends towards saturation, temperature and coupling efficiency play a diminishing role. With a maximum operating current of 17.0 A through the pump-diode bars within the head and the expectation that the current will need to be increased as the diode ages, we aim to run as cautiously as is feasible to prolong the lifetime of the laser heads. With this in mind, we resolve to operate at the lowest current and lowest temperature which may produce the required output energy, and more importantly, energy stability.

Seen at every temperature step in Figure 5.10b, we observe a cyclic beating in the output energy of the regen. A relatively long periodicity points towards this having a thermal origin rather than an optical or electrical one; however, the temperature scan revealed only an optimal running temperature. More curiously still, changing the temperature of the chiller appeared to affect the periodicity of the beating, although not with any consistent trend. A reduction in oscillation amplitude is the only obvious change we see, and this is more likely to be a result of driving the regen into a more stable, saturated state. Given the high total gain of the regen ($\sim 10^9$), quite subtle changes in the single-pass gain that are difficult to measure directly may be responsible.

Once running at a set chiller temperature of 25°C, a second scan was conducted, changing the current across the diodes. Note that this scan was taken on a different date. Shown in Figure 5.11, increasing the diode current applied across the gain module not only increases the total gain, but significantly improves the % RMS energy stability. In

this instance, the periodicity of the beating stays roughly the same, but the amplitude suffers a dramatic increase with decreasing current, causing the RMS stability to rise from 0.8% at 16.9 A to 8% at 16.5 A. An RMS of less than 1% is absolutely critical to stable production of the mid-infrared.

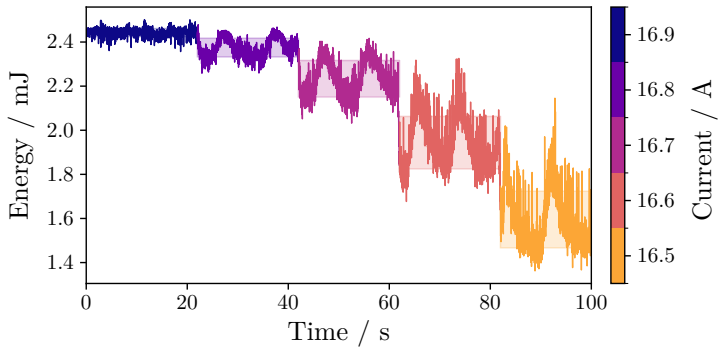


Figure 5.11: Output energy of the Nd:YLF linear regenerative amplifier whilst varying diode driving current. Each current step was recorded over 2000 shots using a GENTEC QE12SP energy meter. This scan was run at a set chiller temperature of 25°C. For decreasing driving current, the mean output energy not only decreases but also becomes significantly less stable. The periodic beating observed in previous scans increases in amplitude at each current step.

Investigations into this characteristic beating are still ongoing. The Cerberus laser system, with which Chimera shares a lab, also shares this instability in its regenerative amplifier of the same model (albeit at a much lower level due to the Cerberus regen being seeded with a significantly higher pulse). This would appear to rule out the possibility of this being the consequence of a poorly matched or low-energy seed (Cerberus being driven much harder), as well as individual chiller or electronic (e.g. Pockels cell timing) issues. It instead points towards a lab-wide issue, either in the internal lab environment or in the external university water cooling system. Despite its origin remaining unexplained, it is clear that when the regen is driven into a saturated state, either by driving the diodes at higher current or with improved coupling of the seed into the regen cavity, these fluctuations can largely be avoided.

Due to the very high gain of the amplifier, the output spatial profile of the regen is almost entirely dominated by the cavity rather than the input seed, meaning that this part of the system is robust against small spatial variations in the seed. Beam profiles for the input and output beams interacting with the regen are shown in Figure 5.12. As the regen cavity is a linear one, the seed can be mode-matched by using an adjustable telescope for optimal spatial overlap with the laser rod and coupling efficiency into the amplifier cavity.

The ring amplifier, being a lower-gain, non-cavity system, is very sensitive to the spatial quality of the input beam, and thus the output mode of the linear regen becomes important. To control this as much as possible, we use a combination of ‘fluffy’ (Gaussian profile serrated laser cut 200 µm thick steel) apertures and vacuum spatial filters (VSFs) within (and before) the ring amplification loop. These help to remove any aberrations or high-frequency hot spots within the beam – high spatial frequency structures created by

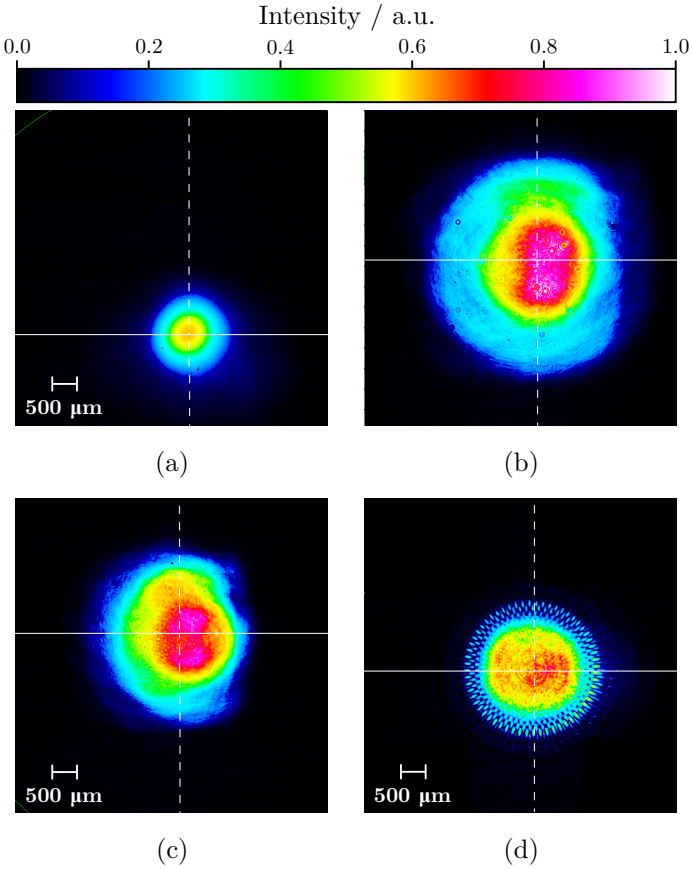


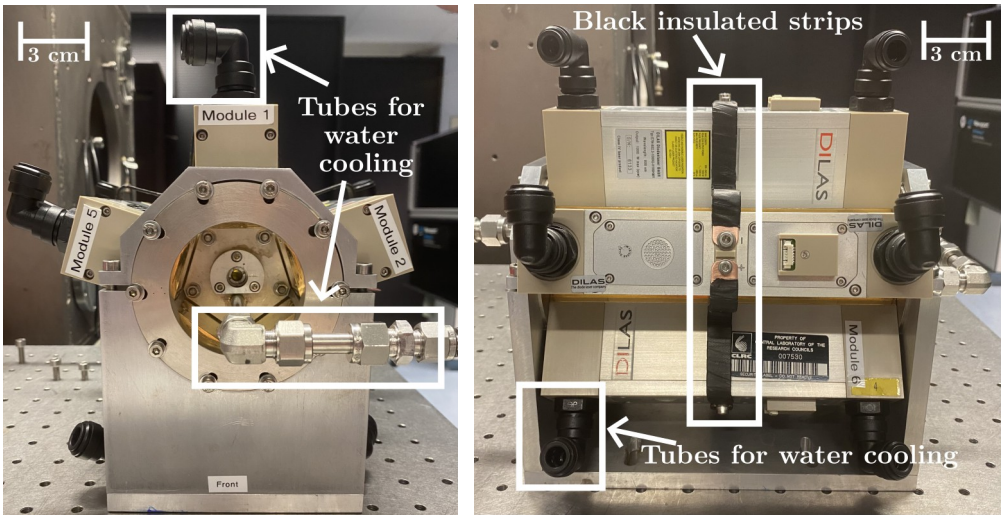
Figure 5.12: Beam profiles, taken using a WinCamD-LCM – 1" CMOS Beam Profiler, showing the mode before and after passing through the linear regenerative amplifier (regen).
(a) Regen input mode.
(b) ASE from the (unseeded) regen.
(c) Regen output mode.
(d) Regen output mode (ring input mode) immediately after passing through a 3.5 mm serrated aperture. The high-frequency triangular profile is produced by the serrated edge of the aperture – these are subsequently smoothed out by a VSF within the multi-pass ring booster loop.

the fine ($\sim$ 100 µm 'spike' feature size) serrated edges of the aperture are filtered out by the VSF. This additionally helps keep the total B-integral down to a minimum, as the VSF suppresses the onset of self-focusing wave-fronts within the beam. A Cr:YAG saturable absorber within the ring booster loop is used to suppress amplification of the unwanted regen ASE pedestal.

5.4.2 The Ring Booster

Shown in Figure 5.13, the 6-pass ring amplifier is a custom-built system based around a geometric 4-pass large aperture gain module developed by the Rutherford Appleton Labs [**100**, **205**]. Also run at 100 Hz, it is made up of 5 Dilas GmbH 1kW, 802 nm QCW E7N-802.3-1000Q-H189 pump diode modules connected electrically in series with a five-fold symmetry, together positioned around a deep-etched Nd:YLF rod (100 mm long and 6 mm in diameter). This symmetry allows the rod to be pumped as uniformly and homogeneously as possible. The process of deep-etching here refers to the chemical

removal of the outermost surface layer of the crystal to aid direct water cooling at high pump fluences. The process also works to reduce scattering and improve coating adhesion by eliminating polishing damage and micro-cracks.



(a) (b)

Figure 5.13: Diode pumped Nd:YLF ring regenerative amplifier.
(a) Front view.
(b) Side view. The black insulated copper strips show how the modules are electrically connected in series.

Each diode module (190 × 400 mm) contains a stack of nine sub-arrays, which themselves are made up of thousands of individual emitters, all cooled by integrated microchannel water cooling. The Nd:YLF rod is immersed directly in the water flow, both to provide efficient cooling and to minimise surface reflection losses, such that at least 95% of the diode output may be coupled into the crystal. The flow rate of the system (15 l/min in the flow tube, 3.3 l/min for individual diode modules) is set to maintain a temperature of 18°C within the optical components [100]. This is specifically set so that the wavelength emitted by the diodes (dependent on their temperature) is optimally matched to the absorption wavelength of the laser rod.

YLF has an anisotropic negative thermo-optic coefficient ($\partial n/\partial T$) and, as such, can act as a defocusing thermal lens. The consequence of this is that the beam will naturally be deformed upon each of the 6 amplifying passes, and evolve from a circle into an ellipse if this process is not well controlled. To attempt to average out these effects, the beam is spatially rotated 90° using a mirror-based periscope (with the polarisation maintained in a constant direction using a subsequent half-waveplate), and inverted using a VSF on each of the 6 passes. Despite this cumulative beam shaping, the output beam profile retained some non-uniformity, appearing slightly S-shaped. This was not particularly problematic as we still possessed sufficient energy to continue through the system; however, since the output of this amplifier pumps several OPA stages, a sub-optimal spatial profile results in inefficient coupling, meaning that a significant portion of energy in the beam wings may be wasted.

To combat this, we imposed a small ($\sim 6°$) twist onto one of the intra-cavity lenses in the VSF assembly. To correct further, we introduced two fluffy apertures (6 mm inner diameter) at opposing diagonals. This allows us to effectively shave down the beam upon each pass and results in a much more uniform output mode. The optical beam path of the ring booster cavity is shown in Figure 5.14. Not only is the profile significantly improved, but we also see a pronounced increase in the ring's output energy, despite the relatively aggressive spatial filtering it undergoes on each pass. Previously delivering 35 mJ of output energy with 1% RMS stability, the output of the ring amplifier now reaches 48.4 mJ with 0.6% RMS. Both energy measurements were recorded with a GENTEC QE12SP energy meter, taken over 10,000 shots. This allows us additional energy control within the future beam path, either to provide more allowance to underperforming stages when required or, should the damage threshold allow, provide more routine OPA pump energy. Beam profiles before and after these additions, measured after being compressed, approximately 2 m after the ring output, are shown in Figure 5.15.

An effect caused by heat deposition in the gain medium, thermal lensing, can change dramatically with intensity. This means that the output beam can evolve differently depending on the energy of the input seed coming from the regen and the drive current delivered to the laser head (normally kept constant). The diode modules surrounding the Nd:YLF laser rod are electrically pumped (in series) at a current of 80 A, with an electrical pulse duration of 400 µs, while a typical input energy into the ring is 2.80(1) mJ. As with the regen, for maximum output energy stability, it is important to drive the ring as much as possible into a saturated state, with the cavity well optimised to combat any non-uniformities in the spatial mode. As the diodes themselves are well beyond their lifetime specification (and now impossible to replace), the diodes should not be pumped harder than this – should saturation not be reached, this is a compromise that must be made for routine operation, so as not to risk over-driving (and damaging) the diodes.

After amplification in the ring booster, 30% of the pulse is compressed to ~ 8 ps using a pair of 1741 l/mm gold diffraction gratings in a Z configuration, becoming the pump for both the 800 nm OPCPA sub-system as well as for the first MIR OPA stage. The remaining 70% is stretched to ~ 600 ps before further amplification within a flashlamp pumped chain of Nd:YLF amplifiers.

5.4.3 Flashlamp Pumped Chain

The 1053 nm light, which will go on to pump the final mid-infrared OPA stages, is amplified in a chain of water-cooled flashlamp-pumped Nd:YLF amplifiers. There are two amplifier systems currently in use: two Continuum heads (salvaged from an Nd:YAG system and retro-fitted with Nd:YLF rods), both used in a double-pass configuration to amplify the 1053 nm pump to approximately 100 mJ for OPA2, and 4 Litron heads (3 single pass, 1 double pass), which form the final amplification stage and deliver up to 1 J of 1053 nm light to pump OPA3a+b. The rods used in this chain vary between 7 mm, 9.5 mm and 12 mm in diameter.

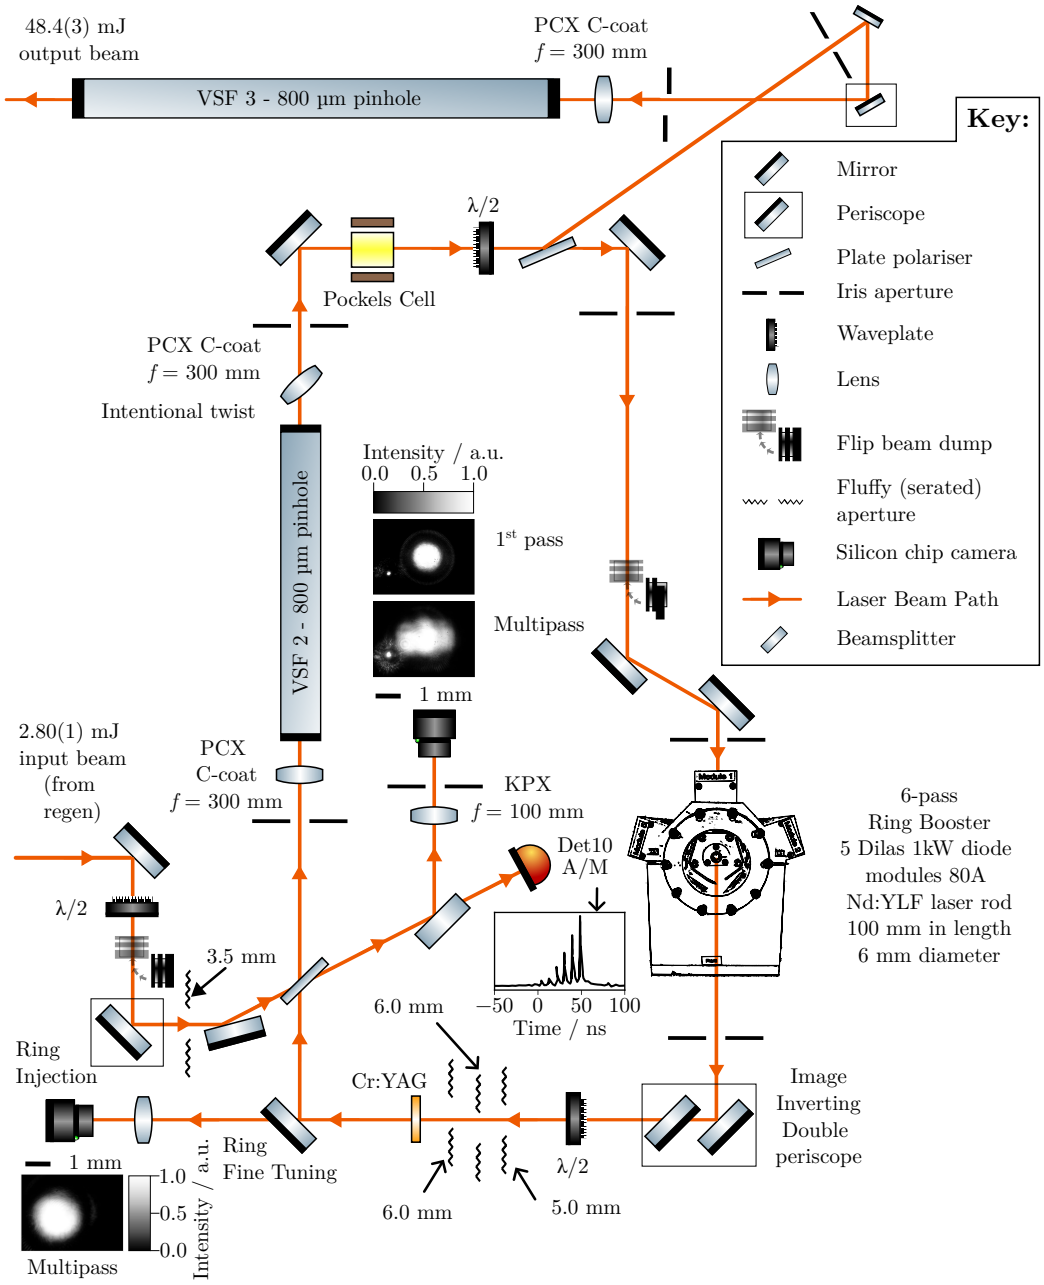


Figure 5.14: The Ring Booster Amplification Loop. The first-pass alignment through the ring loop is controlled day-to-day using the ring injection mirror. The multi-pass alignment is adjusted using the ring fine-tuning. Energy through the loop is controlled using the pre-cavity half-waveplate, with the amplifier head typically running at a fixed drive current. To combat thermal lensing effects within the amplifier, the beam is spatially rotated and inverted upon each amplifying pass using the indicated double periscope and vacuum spatial filter (VSF 2). An intentional $\sim 6°$ horizontal twist is introduced on one of the intracavity lenses to combat astigmatism introduced in the amplifier. Fluffy apertures are used to further control the shape of the beam: The 5.0 mm fluffy aperture is used to impose a consistent beam size through the system, and the following two 6.0 mm apertures are inserted at opposing diagonals to manipulate the output beam shape. Beam profiles used for day-to-day alignment (taken on standard silicon chip CCD cameras) are included, as is the photodiode trace showing the multi-pass gain, recorded on a fast oscilloscope. The y (voltage) axis has been removed for clarity. Diagram not to scale.

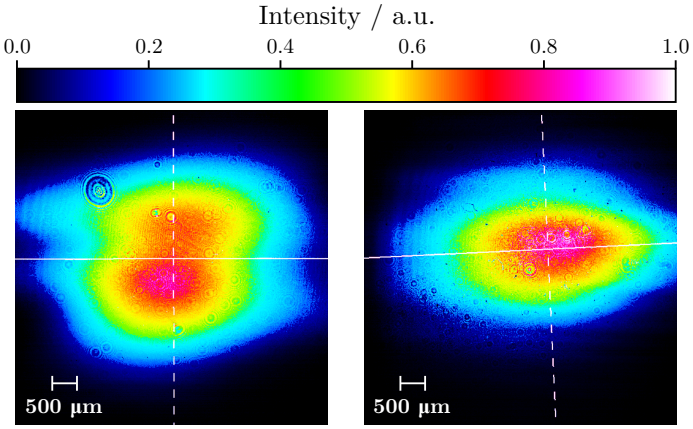


Figure 5.15: Beam profiles, taken using a WinCamD-LCM - 1" CMOS Beam Profiler, showing the output mode of the ring regenerative amplifier before and after the addition of significant beam shaping. These were measured after compression, roughly 2 m after the output of the ring. Small circular diffraction patterns in the images are due to dust particles on the surface of the camera chip. There are some faint Airy diffraction rings visible on the left-hand side of both profiles. These appear due to a small clip through an aperture in the beam path.
(a) Slightly S-shaped output beam prior to additional beam shaping. The image shows what appears to be two separate lobes in the beam due to the imperfect inversion correction on the ring multi-pass geometry. Measured using a GENTEC QE12SP energy meter over 10,000 shots, the energy contained in the pulse is 35.0(4) mJ.
(b) Beam profile taken after the additional beam shaping. The beam becomes much more uniform, with a single central lobe that propagates more elegantly. The energy contained in the pulse, measured using a GENTEC QE12SP energy meter over 10,000 shots, is improved to 48.4(3) mJ.

Before entering the flashlamp chain, the 1053 nm light from the ring is stretched to $\sim$ 600 ps using a folded Martinez-style grating stretcher (1741 l/mm). The Fourier plane intrinsic to this stretcher enables discrete control of spectral phase and amplitude for arbitrary pulse shaping, although that has not yet been implemented. In order to most efficiently use the pump energy in the OPAs, it is important to retain as much spatial and spectral quality as possible in each pump beamline. For this purpose, vacuum spatial filters (VSFs) are placed between major amplification stages to remove high-spatial-frequency noise. The beam is focused into the Fourier plane of the filter and passed through a pinhole, typically implemented as a short length of hollow-core glass tube operated under vacuum, often in combination with serrated apertures at the VSF input. This arrangement selectively removes high-spatial-frequency components, such as diffraction spikes and fine-scale intensity modulations, while preserving the low-frequency beam structure.

Due to the birefringence and polarisation-dependent wavelength of Nd:YLF, double-pass operation is achieved with Faraday isolator setups, combined with half-waveplates and polarisers, which maximise energy extraction while preventing back-reflections. This arrangement is applied only to the lower-energy modules; at higher energies, it is avoided to minimise pre-pulse generation from co-aligned surface reflections. A schematic diagram

of the full chain, together with representative beam profiles at key stages, is shown in Figure 5.16.

5.5 Optical Parametric Amplification and Mid-Infrared Generation Stages

Optical Parametric Amplification (OPA) is a key technology in the Chimera laser system, its primary advantage being the ability to amplify broad bandwidths of unusual wavelengths of light, such as within the mid-infrared. OPAs can also be relatively efficient, typically operating with very high gain and minimal heat deposition, which, unlike storage gain amplifiers that suffer from thermal loading, makes them well-suited for high repetition rate regimes. There are four OPA stages within the Chimera laser system. One of these is an 800 nm OPA sub-system before the mid-infrared generation stage. The remaining three stages (the second of which is double-passed) are dedicated to amplifying the generated mid-infrared light.

5.5.1 800nm OPA sub-system

To produce the 527 nm wavelength required to pump the 800 nm OPA stage, 30% of the 1053 nm ring amplifier output is compressed within the region of 8 - 10 ps, with a further 30% directed to a BBO nonlinear crystal for wavelength conversion by means of second harmonic generation (SHG). The compression stage uses a folded Treacy geometry. A trace, taken using an in-house TOAD (Third Order Autocorrelator – see Section 4.1.1) and fitted with a Gaussian temporal envelope, is shown in Figure 5.17. This results in a deconvolved pulse duration of 8.0(2) ps, with a Gaussian deconvolution factor of $1/\sqrt{3}$ (due to the autocorrelation signal scaling as the cube of the electric field intensity). The quoted uncertainty is the standard error from the Gaussian fit.

The crystal currently in use is a $6 \times 6 \times 3.45$ mm^3 Type I β-barium borate (BBO) crystal with a 23° phase matching angle. Typically, we aim to direct 2.2 mJ of 1053 nm light into the SHG stage, focused and collimated to a beam size of 1.6 mm^2 FHWM, resulting in a peak intensity of 1.37×10^7 W/cm^2. This produces roughly 1.1 mJ of 527 nm output, with a conversion efficiency of 50%. The 527 nm beam has a FWHM beam size 1.1×1.1 mm^2 and serves as the pump for the 800 nm OPA sub-system. Spatial profiles before and after the SHG stage are shown in Figure 5.18.

The 800 nm seed for this sub-system originates from a pick-off from the front-end oscillator pulse after it has been pre-chirped by a pair of chirped mirrors (Thorlabs DCMP175) and subsequently stretched in a two-pass fused silica block. For a transform-limited oscillator pulse between 7 - 8 fs, we estimate a net chirp of approximately $+2616$ fs^2 after propagation through these dispersive elements.

The chirped mirrors contribute six reflections at -175 fs^2 per bounce, while the fused silica block (thickness $L = 50$ mm) is traversed in a double-pass at an incidence angle of $\theta \simeq 14°$, contributing $+3666$ fs^2 – the block angle can be adjusted for fine tuning of the GDD. Producing $\sim$ 20 µJ at 100 Hz, the pulse is split immediately after the

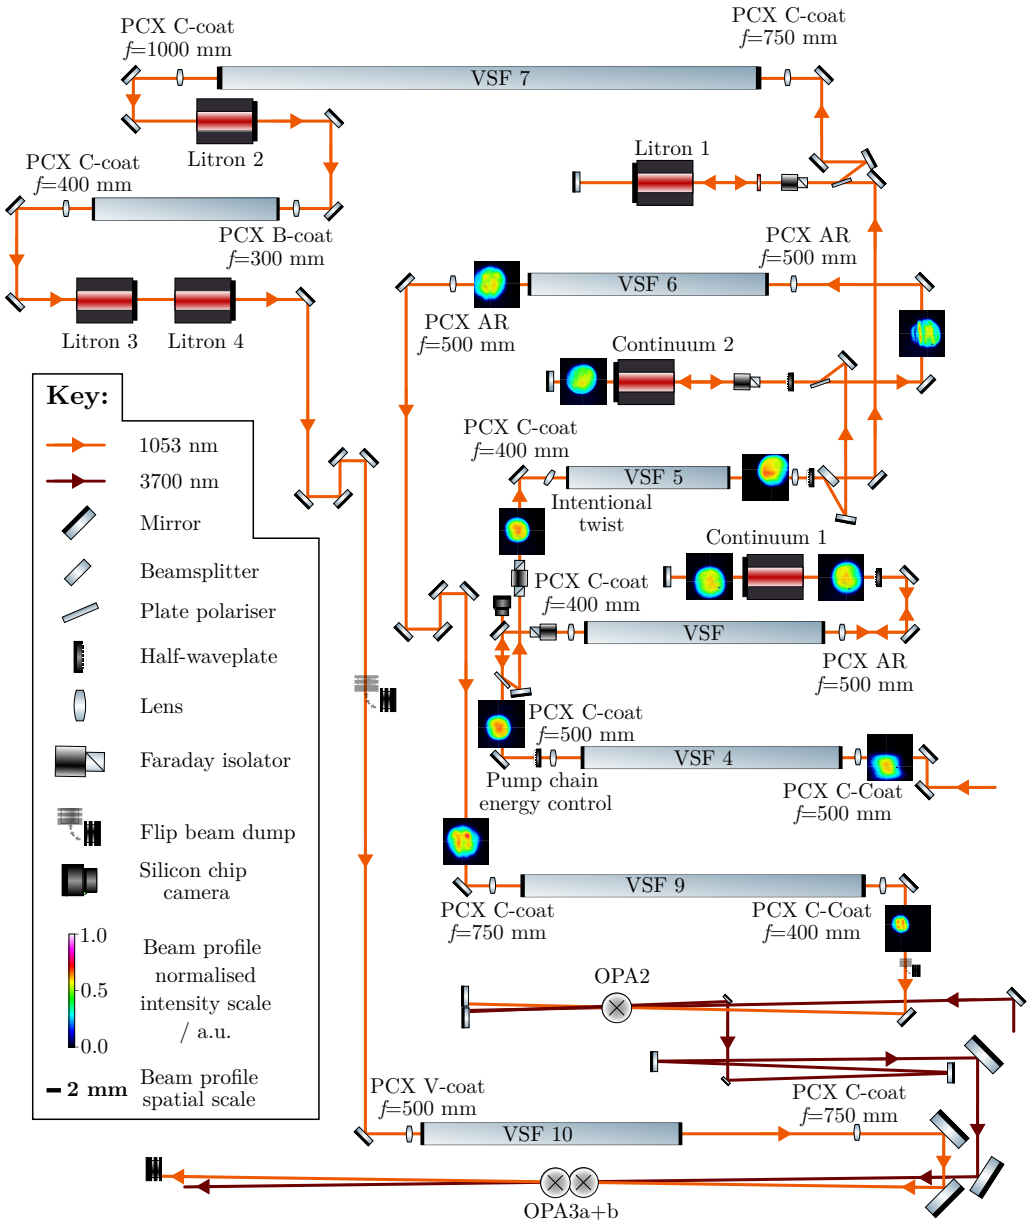


Figure 5.16: A schematic diagram of the Nd:YLF flashlamp pumped chain, showing the beam-path through successive amplifiers and OPAs. The chain consists of two double-passed Continuum heads, which amplify the 1053 nm beam to $\sim$ 100 mJ for pumping OPA2, and four Litron heads (one double-pass, three single-pass), which amplify the beam to $\sim$ 1 J for pumping OPA3a+b. Representative beam profiles, measured with a taken using a WinCamD-LCM – 1" CMOS Beam Profiler, are shown at key stages. Historic damage to Continuum 2 is visible in its output beam profile – the beam quality is marginally improved through passage of VSFs 6 and 9. Diagram not to scale.

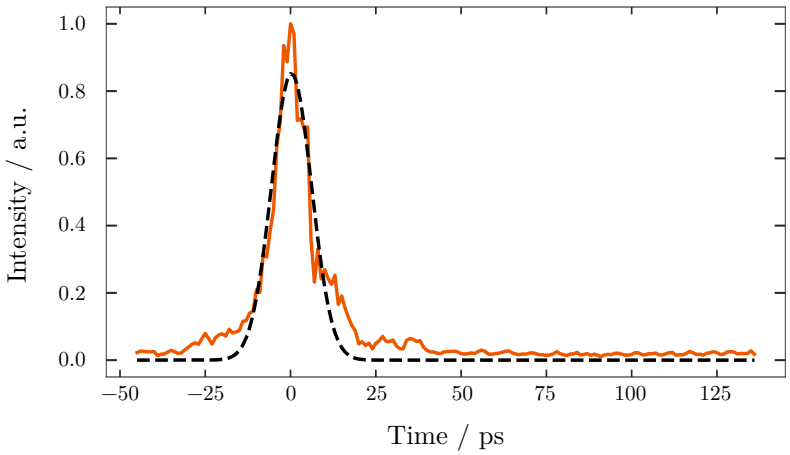


Figure 5.17: A TOAD trace (orange) showing the ring amplifier output (at 1053 nm) after compression. Using a Gaussian fit (black dash) and deconvolution factor, we retrieve a FWHM pulse duration of 8.0(2) ps.

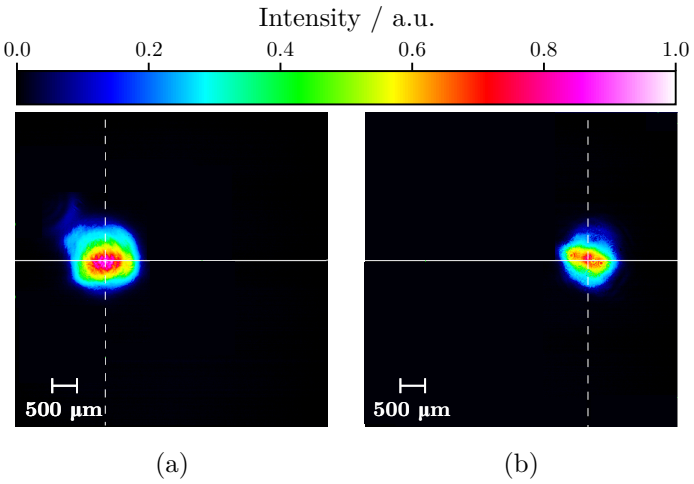


Figure 5.18: Beam profiles, taken using a WinCamD-LCM – 1” CMOS Beam Profiler, showing the input and output modes of the SHG stage.
(a) Shows 2.2 mJ of 1053 nm light at the direct input of the BBO crystal.
(b) Shows 1.1 mJ of 527 nm light at the output of the BBO crystal. The profile is taken ~ 65 cm from the crystal to prevent interference from residual 1053 nm light.

OPA: one branch is directed to a diagnostic line, while the other initiates generation of the mid-infrared seed via difference frequency generation (DFG). The spectrum of the 800 nm line, before and after amplification in the 800 nm OPA subsystem, is shown in Figure 5.19.

Before being sent to the DFG crystal, the amplified 800 nm beam is stretched once again, through 22 bounces on a second pair of Thorlabs DCMP175 chirped mirrors. At -175 fs^2 per bounce, this contributes a negative GDD of -3850 fs^2, resulting in an overall GDD incident on the DFG stage of -1234 fs^2.

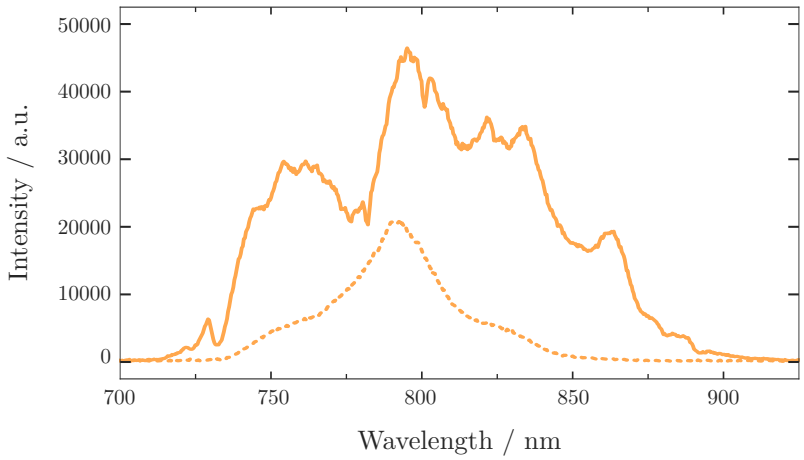


Figure 5.19: The 800 nm seed spectrum before (dash) and after (solid) amplification in a BBO OPA stage. There is $\sim$ 150 nm of amplified bandwidth.

5.5.2 Mid-Infrared Generation

To generate the mid-infrared (MIR) seed, we use a difference frequency generation (DFG) process. This is a particularly unusual feature of the Chimera laser system, as generating the MIR in this way produces a very low energy pulse which, in order to be used in high intensity experiments, must be further amplified using additional nonlinear processes. Most MIR systems worldwide generate mid-infrared light by extracting the idler from an OPA stage. As a result, MIR generation often occurs at the final stage of the amplification chain, inherently limiting the extractable energy due to constraints on pump availability.

An additional challenge, although less frequently reported, is the spectral quality of the MIR pulse. While a big advantage of the OPCPA process is that it allows the signal to preserve the spectral phase characteristics of the seed (independent of the pump), the idler is much more sensitive to distortions [206]. We see this in the phase equations derived from solving the imaginary parts of the coupled wave equations for three-wave mixing (see Section 2.2.2). The case in which we are concerned is one where there is initially zero idler intensity, and thus zero idler phase. Following [206], this condition allows us to write phase equations for the signal, idler and pump ($\phi_{\mathrm{s,i,p}}$) as:

$$\phi_{\mathrm{s}} = \phi_{\mathrm{s}}(0) - \frac{\Delta k z}{2} + \frac{\Delta k \gamma_{\mathrm{s}}^2}{2} \int \frac{\mathrm{d}z}{f + \gamma_{\mathrm{s}}^2} \tag{5.1}$$

$$\phi_{\mathrm{i}} = \phi_{\mathrm{p}}(0) - \phi_{\mathrm{s}}(0) - \frac{\pi}{2} - \frac{\Delta k z}{2} \tag{5.2}$$

$$\phi_{\mathrm{p}} = \phi_{\mathrm{p}}(0) - \frac{\Delta k}{2} \int \frac{f \mathrm{d}z}{1 - f} \tag{5.3}$$

Where Δk is the wavevector mismatch per unit length, $f = 1 - I_\mathrm{p}/I_\mathrm{p}(0)$ is the fractional depletion of the pump beam and $\gamma_\mathrm{s,i} = \sqrt{\frac{\omega_\mathrm{p}}{\omega_\mathrm{s,i}} \frac{I_\mathrm{s,i}(0)}{I_\mathrm{p}(0)}}$ are the input photon intensity ratios.

The consequence of these phase relationships is that any high-order spectral phase or spatial inhomogeneities in the pump can be transferred to the idler, degrading both its temporal and spatial coherence. This significantly complicates pulse compression and focusability, both of which are critical for any high-intensity interaction. As a result, OPA idler-based MIR pulses are often longer and less temporally clean. There are several reports in the literature of how to combat this problem – for example, grating and prism-based systems used to correct angularly dispersed idler beams – but this comes at the expense of a more complex experimental setup and day-to-day stability and usability [207–210].

Although using an OPA idler as a MIR source can appear more convenient, particularly given the alignment challenges posed by the mid-infrared, it comes with considerable limitations in compressibility, energy scalability and beam quality. In contrast, generating the MIR via DFG provides greater control over the pulse characteristics and allows for staged amplification, ultimately enabling higher energy delivery and better compatibility with strong-field physics experiments. By instead producing the MIR using DFG, we are afforded much more precise control, as well as the potential ability for longer-term energy scalability.

It is worth noting that whilst we argue against using an OPA idler for MIR generation, the MIR pulse produced via DFG in our system is, strictly speaking, also an idler, being the longer-wavelength output of the three-wave mixing process. As such, it is still susceptible to many of the same issues discussed above, particularly regarding phase distortions in the pump. To mitigate this, the DFG stage is positioned very early in the system, where the 800 nm pump pulse can be carefully managed. Here, the pulse remains low in energy and passes through very few optics, greatly reducing the likelihood of accumulating higher-order phase distortions. This level of control is far more difficult to achieve with a high-energy pump pulse feeding a final-stage OPA, where optical complexity and thermal effects can introduce significant phase and spatial aberrations.

The DFG stage on the Chimera laser system uses two quasi-phasematched magnesium oxide doped periodically poled lithium niobate (MgO:$LiNbO_3$ or MgO:PPLN) crystals in series (5% MgO dopant – Covesion UK). Each crystal has dimensions of $0.5 \times 3 \times 0.4$ mm^3, with the 0.5 mm dimension aligned along the propagation axis. While birefringent phase matching can achieve 'perfect' phase matching, it is constrained by material, wavelength and phase matching compatibility. To allow for a larger range of wavelengths, quasi-phasematching using periodic poling is a useful option, as the poling period can be custom-matched to the wavelength of light you are hoping to generate. PPLN is a highly efficient medium for this, particularly for applications requiring broadband wavelength conversion, due to its high nonlinear coefficient as well as a large available interaction

length. PPLN also exhibits a temperature-dependent phase-matching condition, meaning that changing the crystal temperature allows some degree of wavelength tuning in your interaction. Despite these attractive criteria, however, undoped PPLN requires high operating temperatures (100 - 200°C) and experiences considerable problems with power instabilities and photorefractive damage due to thermal lensing effects [211]. To combat this, it is common practice to use MgO as a dopant. This maintains both a high nonlinear coefficient and large interaction length while providing a higher resistance against photorefractive damage, improving its wavelength tunability and capacity to run well at room temperature [212].

The MgO:PPLN crystal used in Chimera is custom-manufactured with a linear chirped grating of 19.4 - 22.0 µm. Upon first installation in 2018, the crystal was temperature-tuned using a PPLN oven (Covesion UK), the results of which can be seen in Figure 5.20. The output spectrum was recorded for each temperature step (± 0.01°C stability) in order to maximise the usable bandwidth. The results from this scan provided assurance regarding the bandwidth available to us, with $\Delta\lambda \sim 965$ nm at a central wavelength of $\lambda_0 = 3700$ nm at room temperature. The consistency in the data in the 20 - 50 °C range informs a further design choice in the system: the PPLN temperature oven is not required in our case for fine temperature control, and thus a layer of complexity may be removed.

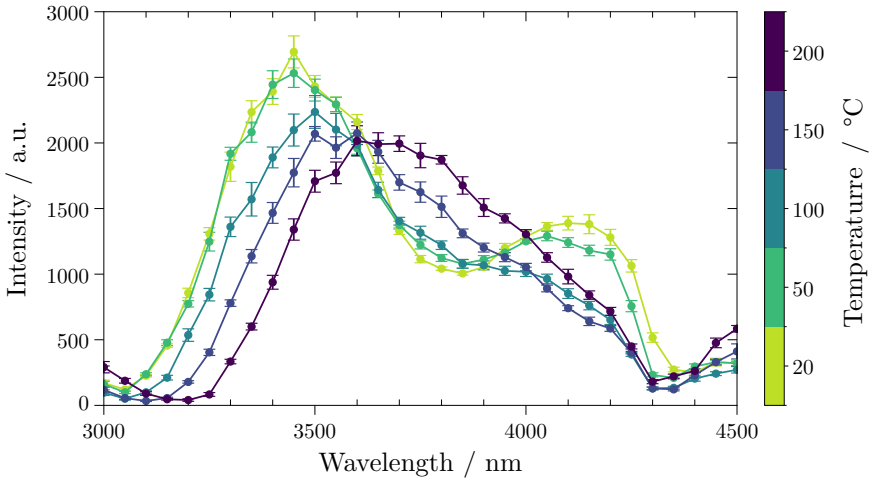
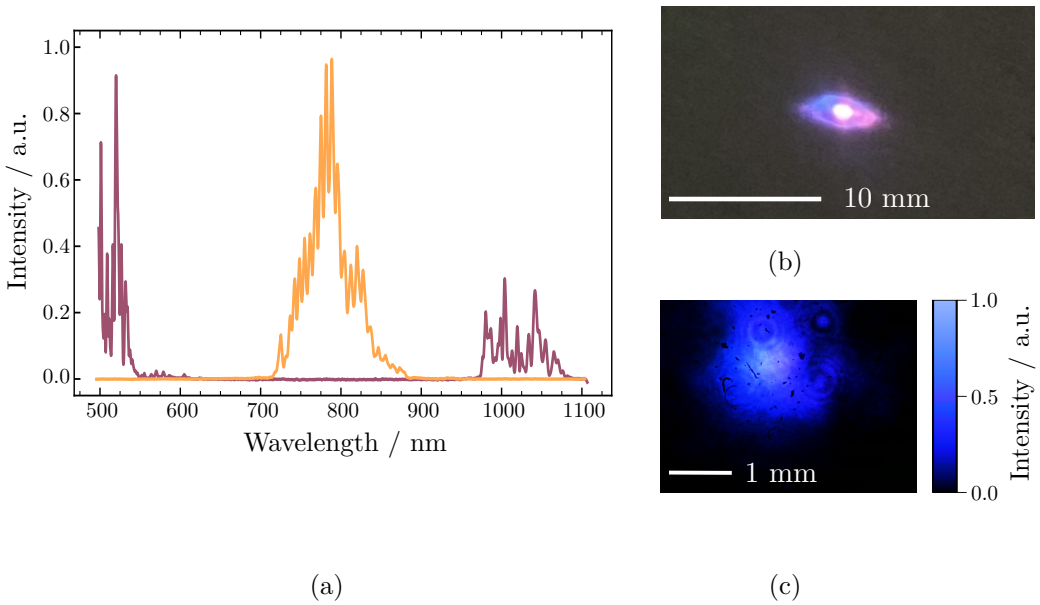


Figure 5.20: The output MIR spectrum of the DFG stage as a function of PPLN temperature. At room temperature, there is ∼ 965 nm of avaliable bandwidth at a central wavelength of $\lambda_0 = 3700$ nm. Data recorded by Dr Nicholas Stuart on 23/04/2018. The error bars in y represent the standard deviation of the photodiode signal as recorded on the oscilloscope.

As the input pulse is chirped to ∼ 190 fs, with an estimated group delay dispersion (GDD) of -1234 fs^{-2}, each wavelength component must be phase matched to a different position in the crystal. This means that the poling period varies linearly along the optical axis to enable more broadband phase matching. This is particularly important for us as we are looking to generate and amplify a very large amount of MIR bandwidth of order 1 µm, to support a few-cycle pulse after compression. Several nominally identical PPLN sections (5 - 10) are available transversely within the crystal. In our current configuration,

two crystals are used in series – a modification first tested during installation, where the addition of the second crystal increased the MIR seed energy by a factor of 5 - 10. At present, these are mounted in a relatively simple, non-optimised setup.

At this stage of the system, we expect the DFG stage to produce $\sim$ 10 nJ of MIR light, with a central wavelength of 3700 nm. The pulse energy is too low to be detected on an infrared card, energy meter, or mid-infrared beam profiler, but it can be detected on a MIR photodiode. A more routine experimental signature arises from the broad bandwidths involved: successful MIR generation is accompanied by the appearance of collinear white and blue light from the PPLN. As shown in Figure 5.21, this is visible to the naked eye and can be readily detected with a standard silicon chip camera or a spectrometer.



(a) (b) (c)

Figure 5.21: Experimental output of the MgO:PPLN DFG stage.
(a) Measured spectra of the output (smoothed using a 1D Gaussian filter), taken both with no filter (orange) and with a blue filter to suppress the 800 nm signal (dark pink). The filtered trace highlights the additional short-wavelength (visible blue) light generated during the DFG process. The feature in the 950 - 1100 nm spectral range is attributed to the DFG signal, generated alongside the MIR to conserve energy and momentum. This may also have its own utility – for example, in seeding a high contrast $\sim$ 1000 nm OPCPA system. Fine structure in both spectra likely arises from the temporary diagnostic setup (non-optimised sampling optics). Spectra were recorded using a Thorlabs SPLICCO spectrometer on 15/04/24. Intensities are not normalised between traces.
(b) A photograph, taken with an iPhone 11, of the PPLN output, visible on a piece of white card. The red wing of the beam is infrared light, not visible to the naked eye.
(c) Filtered beam profile showing the blue DFG signature (in real colour), recorded on a standard silicon chip CCD camera.

The setup used for this measurement is not yet permanently integrated into the system, but, particularly with the addition of a broader-wavelength spectrometer, could become a valuable tool for characterising and routinely monitoring the DFG output.

Specifically, it would enable a full characterisation of both the visible blue signature and the DFG signal in the 950 - 1100 nm spectral range, providing a direct means of optimising MIR generation. Moreover, as these additional wavelengths are collinear with the MIR, they can be exploited for alignment purposes using simple silicon detectors, not sensitive to the MIR.

5.5.3 Mid-Infrared OPA Stages

There are four mid-infrared OPA stages currently in place on the Chimera laser. The first is pumped with a portion of the compressed output from the ring, while the remaining stages (the first of which is double passed) are driven by the Nd:YLF flashlamp pumped chain. Since the beginning of this PhD in October 2021, only the first stage has been operational, and, due to a number of technical challenges with the front end of the system, at the time of writing, the author has not observed gain in any of the subsequent stages.

The first stage, referred to as OPA-B1, is a $5 \times 5 \times 10$ mm^3 Type I potassium titanyle arsenate (KTA) crystal with a 48° phase matching angle. The crystal is pumped with 4.1 mJ of 1053 nm light from the ring booster. Once amplified, the MIR pulse energy is high enough to be detected directly using a Pyrocam (the MIR beam profiler in use, described in Section 4.1.2). The MIR beam profile, taken using the Pyrocam ~ 75 cm after the output of OPA-B1, is shown in Figure 5.22.

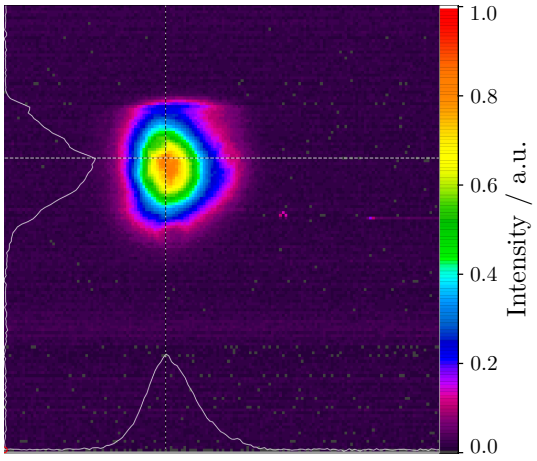
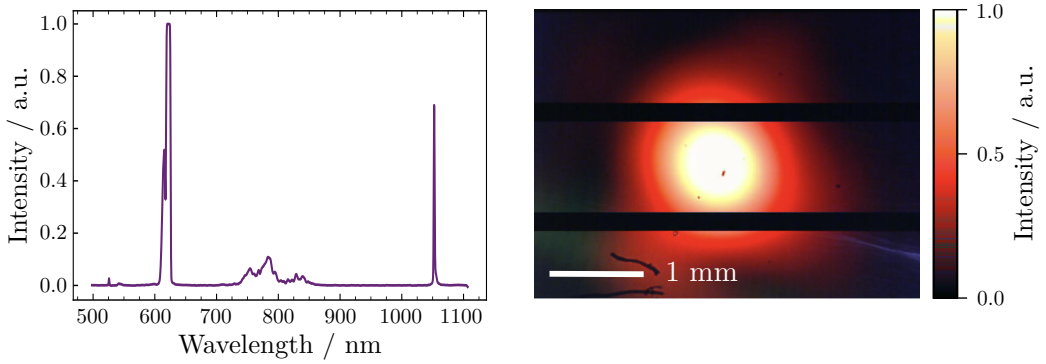


Figure 5.22: A beam profile of the amplified MIR beam (centred at 3700 nm) after the first MIR OPA stage (OPA-B1), measured with a Spiricon Pyrocam IIIHR beam profiler. Axis lineouts were recorded using Spiricon BeamGage software, giving D4σY and D4σY of 9.7 mm and 9.7 mm, respectively. Here, D4σ denotes the ISO standard 'second-moment' beam width, calculated as four times the standard deviation of the intensity distribution. The software also measured a pulse energy of 174.9 µJ and a peak fluence of 2.65×10^3 µJ/cm^2.

At this stage of the system, the average MIR pulse energy previously measured was 70 - 100 µJ. With improved oscillator and ring performance, we now routinely achieve 90 - 140 µJ, with the highest recorded energy reaching 174.9 µJ (measured using Spiricon BeamGage software). The beam profile also appears smooth and largely free from aberrations – the radial clip visible at the top of the profile arises from the beam

impinging on the edge of the camera sensor. This is unavoidable when measuring at this location, given the size of the camera and the fixed position of the beam.

Highly specialised, bulky and expensive, although the Pyrocam is invaluable for MIR alignment and characterisation of the spatial profile, it is not convenient for routine operations. For everyday diagnostics, we instead make use of additional light produced 'for free' in the nonlinear interaction. In the case of MIR generation, this manifests as visible blue light and the 950 - 1100 nm signal; in the OPA process, a visible red beam is generated, which can again be detected using a standard silicon chip CCD camera. This, alongside the measured spectra of the OPA-B1 output, is shown in Figure 5.23.



(a) (b)

Figure 5.23: Experimental output of the OPA-B1 stage.
(a) Measured spectra of the output (smoothed using a 1D Gaussian filter). The trace highlights the additional short-wavelength (visible red) light generated during the OPA process, as well as the residual amplified 800 nm light left over from the DFG process. The feature at $\sim$ 1050 nm shows the depleted 1053 nm OPA pump. There is also evidence of a small amount of visible green light being produced at $\sim$ 550 nm. This can often be observed around the KTA crystal when being pumped and is attributed to the over-pumping of the process, causing an additional second harmonic generation process to occur from the 1053 nm beam. Spectrum recorded using a Thorlabs SPLICCO spectrometer on 15/04/24.
(b) Beam profile showing the red OPA signature (in real colour), recorded on a standard silicon CCD camera. The black horizontal lines are due to excessive input power on the sensor.

For routine monitoring, optimising the brightness and beam quality of the visible red beam provides a useful proxy for MIR performance when observed side by side. For accurate characterisation beyond day-to-day checks, however, the Pyrocam remains the preferred diagnostic.

As in the case of the DFG stage, this OPA stage would benefit from the addition of a broader-wavelength spectrometer – particularly reaching wavelengths up to 2 µm, such that the idler beam at 1.5 µm, generated in the OPA process, may be fully characterised and monitored as an additional day-to-day diagnostic to reflect the quality of the amplified MIR pulse.

We have also trialled the use of direct MIR detection with silicon cameras as a lower-cost routine diagnostic for the MIR beam. Although they cannot directly detect MIR light under normal conditions, when the beam is tightly focused onto the sensor, a

detectable signal appears due to a three-photon absorption process in the chip. Using the metrics given by the Spiricon BeamGage software (using D4σ metrics: the ISO standard 'second-moment' beam width, calculated as four times the standard deviation of the intensity distribution), which measures a beam area of 9.7 × 9.7 mm^2, a MIR CaF_2 lens with $f = 50$ mm will produce a focal spot which has a diameter of ~ 23.6 µm.

Mounted on a translation stage, the lens produced a focus with intensities high enough to quickly damage the chip. At positions slightly away from best focus, however, the intensity was sufficient to generate a bright, non-destructive signal representative of the beam, as shown in Figure 5.24. Despite the image shown being saturated, we observed surprisingly good resolution – even to the point of resolving Airy rings when closing an aperture in the beam path. This highlights the potential value of the direct three-photon method: if placed at strategic points in the system, such cameras could serve as useful low-cost, day-to-day alignment diagnostics for the MIR.

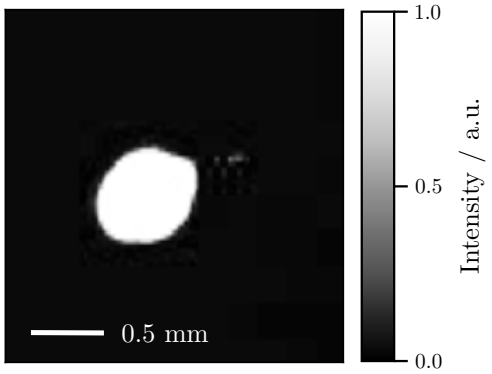


Figure 5.24: Saturated, bright MIR detection at a peak intensity of $\mathcal{O}(10^{11}$ W/cm$^2)$. The visible signature of the MIR beam on a silicon ship CCD camera is generated by three-photon absorption in the detector when the 3.7 µm pulse is tightly focused.

For the profile shown in Figure 5.24, we can estimate a beam size of 0.66×0.85 mm^2. Assuming a pulse duration in the range 100 - 200 fs, this results in peak intensities of order 10^{11} W/cm^2.

The remaining two MIR OPA stages have never been operated by the author; however, we include them here for completeness. Before further amplification, the MIR pulse is stretched to an estimated pulse duration of ~ 370 ps using an Öffner stretcher (see Section 5.6).

The second MIR OPA, referred to as OPA2, is a $10 \times 10 \times 30$ mm^3 Type I KTA crystal with a 41° phase matching angle. Pumped by ~ 100 mJ of 1053 nm light amplified in the Nd:YLF flash lamp pumped chain, it is run in a double pass configuration. When working well, the OPA2 is found to produce $\mathcal{O}(1$ mJ) of pulse energy with ~ 500 nm of available bandwidth. This is equivalent to a ~ 50 fs, 3-cycle pulse at its transform limit. An example of the MIR spectrum of the output of OPA2 is shown in Figure 5.25.

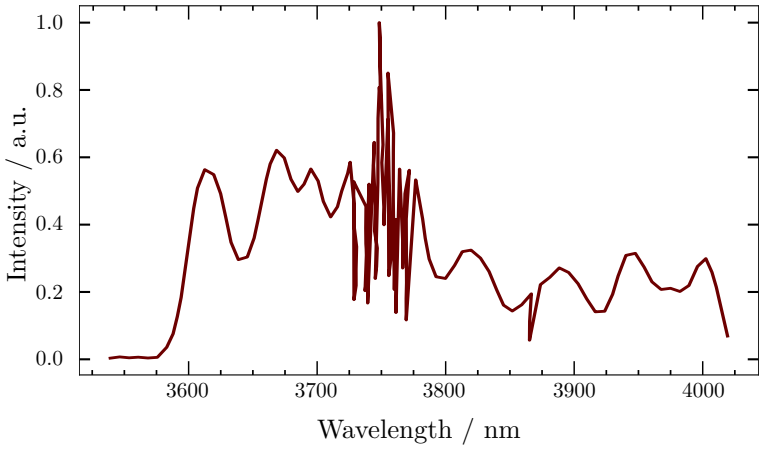


Figure 5.25: A spectrum taken at the output of the OPA2 by Dr Nicholas Stuart. The central feature is due to damage to the detector. The wavelength calibration is estimated, not measured.

The third MIR OPA, referred to as OPA3, contains two separate KTA crystals (both $20 \times 20 \times 15$ mm^3, phase matching angle 41°) and is pumped with $\sim$ 1.2 J of 1053 nm light from the flashlamp chain. Contrary to modelling predictions carried out by Dr Nicholas Stuart, which initially suggested a single longer crystal, higher gain was found experimentally by using two shorter crystals in sequence with independent angle tuning.

Although not verified during this work, the author was present when a new energy record was set in September 2021 (prior to the start of this PhD). At that time, a 31 mJ, 100 fs (uncompressed) 3.7 µm signal was recorded at the output of OPA3. Due to subsequent failures of several key system components, a new and well-characterised energy record attempt remains a work in progress, with recent work instead being focused on improvements to the stability and performance of the front end of the system.

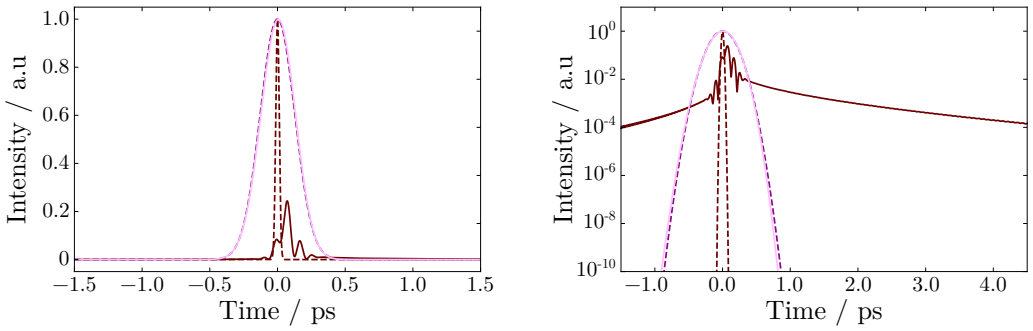
5.6 Stretching and Compressing the Mid-Infrared Pulse

Similar to the challenge of mid-infrared generation, there is much variation and discussion over the best way to stretch and compress a mid-infrared pulse, such that the highest intensity may be extracted. For systems in which the MIR is generated from the idler of a final-stage OPA, a dedicated stretcher is typically unnecessary as the pulse is not subject to further amplification. If, however, the aim is energy scalability, it is crucial to design an efficient stretcher system that preserves the pulse bandwidth while enabling amplification.

A big problem faced in any ultrafast system is the accumulation of high-order spectral phase during stretching and compression. In the mid-infrared, this issue becomes particularly important due to the broader bandwidths required and the therefore increased sensitivity of the pulse shape to dispersion. For this reason, stretcher designs in the MIR must aim to minimise both third-order (TOD) and fourth-order (FOD) dispersion

as much as possible to preserve the compressibility of the pulse. The vast majority of systems presented in Table 5.1 use bulk material to compress (and stretch where relevant) their MIR pulse. While material dispersion can be tailored to provide both positive and negative group delay dispersion (GDD), it inevitably introduces significant third-order phase, often limiting the achievable pulse duration after compression.

To understand the impact of higher order, particularly third- and fourth-order phase, we can use the Sellmeier equation (Equation (2.45)) to estimate the dispersion accumulated by a 3700 nm, 30 fs pulse propagating through 1 cm of silicon (a potential bulk stretcher or compressor element), and compare it to the case of a longer 300 fs pulse. This example would result in phase terms of $\phi_2 = 4.186 \times 10^4$ fs^2, $\phi_3 = 4.204 \times 10^5$ fs^3, $\phi_4 = 1.823 \times 10^7$ fs^4. With the impact of these TOD and FOD terms shown in Figure 5.26, whilst this has a negligible effect on longer pulses, few-cycle pulses are extremely sensitive to higher-order dispersion. An obvious (but high-cost solution) is to introduce a MIR Dazzler (acousto-optic pulse modulator used for active high-order phase control) [213] into the system. For the 2000 - 3700 nm spectral range, Fastlite specifies a maximum programmable delay of 5.5 ps at 2800 nm. This corresponds to roughly 2×10^3 fs^3 TOD or 1×10^9 fs^4 FOD. To preserve pulse quality in ultrashort regimes, our system design therefore aims to limit third-order phase to below $\sim 10^4$ fs^3.

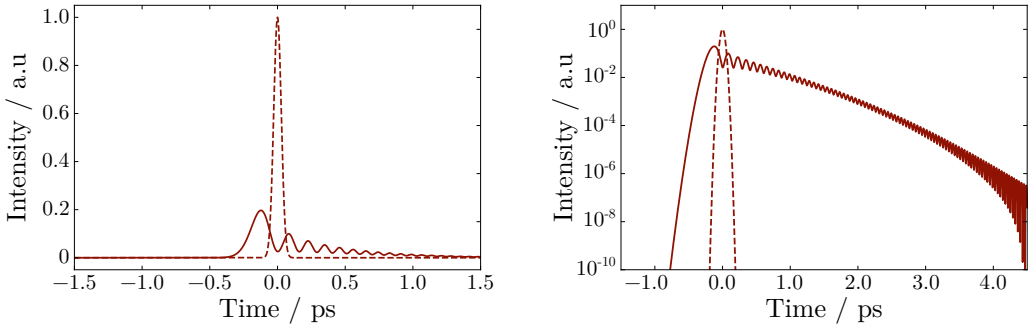


(a) (b)

Figure 5.26: The effect of third- and fourth-order spectral phase on the compressibility of an ultrashort pulse. For Chimera parameters of $\lambda = 3.7$ µm and a transform-limited pulse duration of 30 fs (dark red), we compare propagation through 1 cm of Si with that of a longer, 300 fs pulse (pink). This corresponds to spectral phase terms of $\phi_3 = 4.204 \times 10^5$ fs^3 and $\phi_4 = 1.823 \times 10^7$ fs^4. Dashed curves represent the transform-limited pulses, while solid curves show the effect of third- and fourth-order phase.
(a) Shows this on a linear intensity scale, while (b) uses a logarithmic scale.
With the addition of higher-order dispersion, the compressed pulse is limited to a minimum FWHM duration of 59.3 fs. We note, however, that the FWHM metric does not capture smaller high-intensity features below the half-maximum level, which still can contribute significantly to the pulse envelope.

Comparing this now to a system which uses bulk material to perform their stretch, for example, the system at RIKEN, Japan [187], we show the effect of their high-order phase upon compression in Figure 5.27. They report a transform-limited pulse duration of 66 fs, propagating through 14 cm of bulk Si. This will result in phase terms of: $\phi_2 = 6.610 \times 10^4$ fs^2, $\phi_3 = 5.951 \times 10^5$ fs^3, $\phi_4 = 2.312 \times 10^7$ fs^4. Propagating through this

bulk material adds significant dispersion to the pulse, which will have to be compensated for if they wish to perform an adequate compression. They later chose to compress the pulse using 60 cm of bulk CaF_2. While this material does provide the opposite-sign GDD necessary to compensate for the second-order chirp, it also introduces additional higher-order phase. With this compression producing phase terms of : $\phi_2 = -8.950 \times 10^4$ fs^2, $\phi_3 = 5.790 \times 10^5$ fs^3, $\phi_4 = -3.222 \times 10^6$ fs^4, we therefore find residual phase terms of: $\phi_2 = -2.340 \times 10^4$ fs^2, $\phi_3 = 1.1741 \times 10^6$ fs^3, $\phi_4 = 1.9898 \times 10^7$ fs^4. These residuals result in a FWHM pulse duration of $\sim$ 300 fs, although we note that when the pulse shape no longer resembles a Gaussian, this metric fails to capture its true shape.



(a) (b)

Figure 5.27: The effect of third- and fourth-order spectral phase on the compressibility of an ultrashort pulse. For RIKEN parameters of $\lambda = 3.3$ µm, and a transform-limited pulse duration of 66 fs [187], we show the spectral phase terms acquired through propagation through 14 cm of Si, followed by 60 cm of CaF_2. This corresponds to $\phi_2 = -2.340 \times 10^4$ fs^2, $\phi_3 = 1.1741 \times 10^6$ fs^3 and $\phi_4 = 1.9898 \times 10^7$ fs^4. Dashed curves represent the transform-limited pulses, while solid curves show the effects of high-order phase.
(a) Shows this on a linear intensity scale, while (b) uses a logarithmic scale.
With the addition of higher-order dispersion, the compressed pulse is limited to a minimum FWHM duration of 299.8 fs. If the second-order order-phase was perfectly compensated for, this would result in a minimum FWHM duration of 143.5 fs. This simplified representation does not account for additional dispersion contributions within the RIKEN system, and furthermore, the FWHM metric alone cannot capture smaller substructures in the pulse envelope, such as high-intensity features below the half-maximum level.

All materials commonly used in the mid-infrared spectral range exhibit positive TOD, meaning that bulk materials alone cannot compensate for higher-order dispersion. As a result, proper compression requires the use of additional dispersive elements such as prism gratings (grisms) or an Öffner triplet to adequately correct for the accumulated TOD.

The bulk material implementation used at RIKEN was, in fact, trialled in an earlier iteration of the Chimera MIR stretcher/compressor design. Using 14 cm of bulk Si to stretch, and 20 cm of bulk CaF_2 to compress, (with additional dispersion caused by propagating through the later OPA stages also accounted for), the residual dispersion was found to be: $\phi_2 = -1683$ fs^2, $\phi_3 = 1.01 \times 10^6$ fs^3, $\phi_4 = 2.23 \times 10^7$ fs^4. If the remaining chirp were perfectly compensated, the pulse would compress to 150 fs, compared to a transform-limited duration of 60 fs.

Previous OPCPA modelling performed by Dr Nicholas Stuart suggested that we might be able to amplify up to 700 nm of bandwidth for 5 - 10 mm of KTA. For a Gaussian pulse, this would provide a transform-limited pulse duration of $\sim$ 30 fs (2.3 optical cycles). The goal, therefore, is to design a system that can support 700 nm of FWHM bandwidth through both a MIR stretcher and compressor in order to produce a short, few-cycle pulse, close to the transform-limit, whilst simultaneously suppressing higher-order phase. Bulk materials offer a clear advantage for this purpose due to their straightforward, robust alignment. Indeed, we use a YAG-based stretcher prior to the first MIR OPA stage. However, the large optical path lengths required to achieve the necessary GDD for later stages mean that the accumulated TOD – as illustrated by the RIKEN system – becomes challenging to compensate for. We turn instead to alternative dispersion management schemes, which are better suited to scalable, ultra-broadband MIR pulse shaping.

For the pulse durations we are hoping to reach, it is preferable to use reflective rather than refractive optics as much as possible in order to minimise chromatic aberrations. Potential candidates for the MIR stretcher included reflective geometries such as the Martinez or Öffner stretcher. These can then, in principle, be perfectly matched to a Martinez or Tracey-style grating compressor following the final amplification stage. Prism compressors were also considered in the design; however, due to multi-photon absorption when using materials transparent to the mid-infrared at our expected intensities, they are not well suited to compensate for third-order phase [214–218].

Using the Optometrika raytracing code [219], modified to include gratings and pulse dispersion analysis by Dr Nicholas Stuart, different stretcher-compressor configurations can be modelled. Although in theory a Martinez stretcher-compressor pair provides perfect dispersion compensation and the flexibility to reach arbitrary stretched pulse durations while maintaining compressibility, lab implementation and raytracing revealed significant shortcomings. Specifically, when aligned, the Martinez stretcher system exhibited strong aberrations and a vertically chirped output.

These issues arose from the use of off-axis parabolas: when the injected beam does not precisely intersect the focal points of the parabolas, spherical aberrations are introduced. Since angled beam injection is necessary to retrieve output beams in this geometry, this configuration cannot be used to stretch our MIR pulse.

In contrast, if the diffraction grating is placed at the radius of curvature from the concave and convex mirrors, an Öffner all-reflective triplet design can perfectly correct any such spherical aberrations introduced by parabolas. Designed for aberration-free imaging across a wide field of view, the Öffner triplet uses three mirrors to provide positive GDD whilst attempting to minimise any spatial or temporal aberrations. It is a fully reflective design, meaning there is no dispersion to consider, and is suitable for broadband operation.

A standard Öffner stretcher typically requires large and expensive optics. The concave mirror must be sufficiently oversized to accommodate the convex mirror between it and the grating, with the beam propagating both above and below it. Reflective optics must additionally exhibit surface waveform errors below $\sim \lambda/30$ to avoid introducing random aberrations in the spectral phase. This is one of the few cases in which working in the MIR is advantageous: typical research-grade optics specified at $\lambda/10$ for 600 nm correspond to a much improved $\lambda/60$ tolerance when used at 3700 nm, making these components more forgiving in MIR systems. The Öffner stretcher used on Chimera also uses a 4° off-axis configuration, where the beam propagates orthogonally to the dispersion plane. This orientation preserves the spectral phase characteristics of the standard on-axis geometry, whilst allowing for a more compact layout and avoiding the need for excessively oversized optics, helping to further minimise cost and optical footprint. The convex and concave mirrors are 2-inch, $f = 7.5$ cm, and 4-inch, $f = -15$ cm, respectively. The optical layout of the Chimera MIR off-axis Öffner stretcher is shown in Figure 5.28.

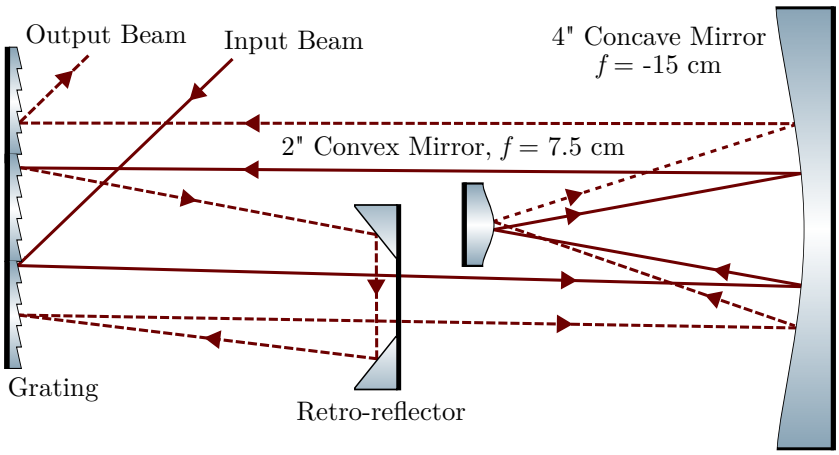


Figure 5.28: Optical layout of the Chimera mid-infrared off-axis Öffner stretcher. The input and output beams are vertically displaced from one another. There is an off-axis angle of 4°. Diagram not to scale.

Using the Optometrika raytracing code, with the Öffner design written by Dr Nicholas Stuart, we can model our setup to retrieve both the spectral phase and 2D spatial profile of a MIR pulse interacting with the stretcher configuration, allowing us to identify any spectral clipping caused by finite optic apertures at the design stage. This is shown in Figure 5.29.

This model was run to stretch the MIR pulse to ~ 320 ps, with GDD of 4.42×10^6 fs^2. No aberrations are observed when using this configuration – the output beam profile remains round and free from spatial chirp, making it ideal for seeding subsequent amplifier stages. For a 48.2° angle of incidence into the stretcher, the Taylor expansion of spectral phase coefficients (around the central wavelength) are solved to be: $\phi_2 = 4.42 \times 10^6$ fs^2, $\phi_3 = -5.08 \times 10^7$ fs^3, $\phi_4 = 9.23 \times 10^8$ fs^4. These coefficients can then be used to guide compressor design, enabling us to compensate for residual phase terms and deliver the shortest possible pulse to the target.

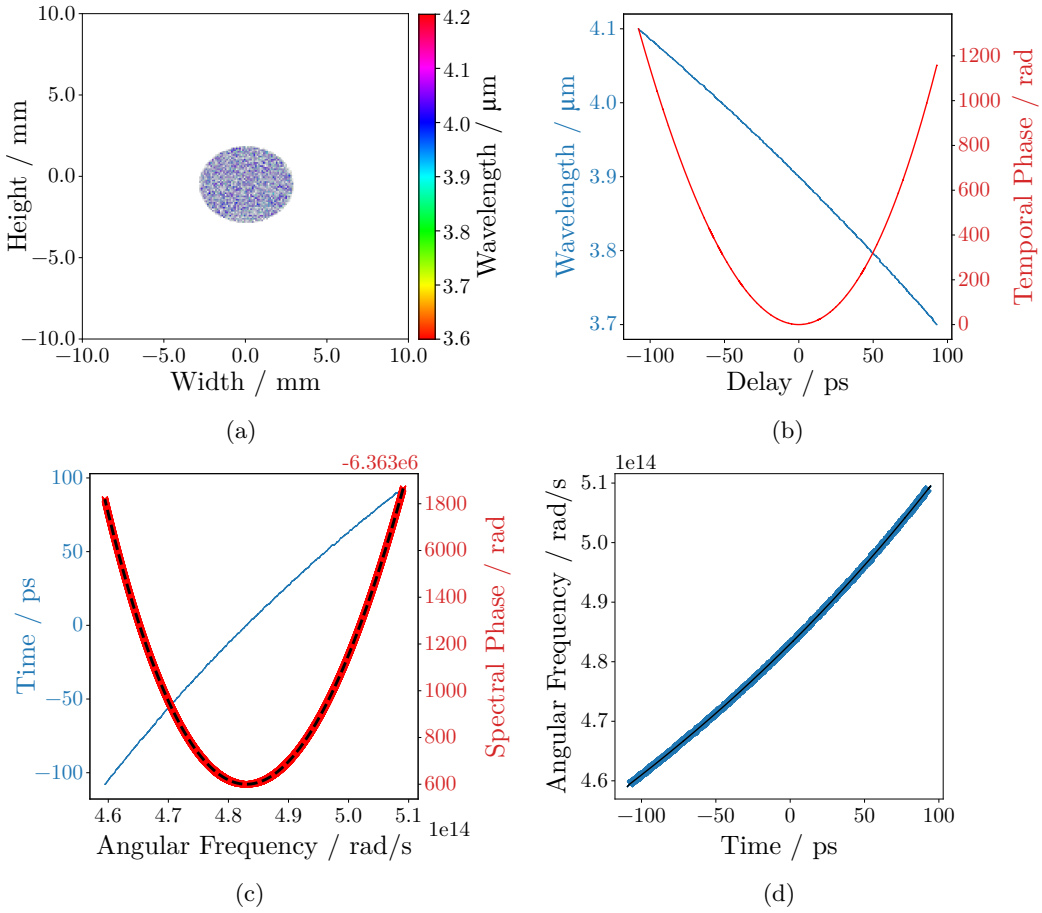


Figure 5.29: Image-on-screen and phase analysis output of the Optometrika raytracing code simulating an Öffner style stretcher system for the Chimera MIR laser. The code for this Öffner stretcher design was written by Dr Nicholas Stuart.

(a) 2D spatial profile of the beam on a screen. Each pixel encodes wavelength via the colour map – the spatial distribution shows how the spectrum varies across the beam profile. For the Öffner stretcher design, the beam is circular and symmetric, with no visible spatial chirp or angular dispersion. The lack of horizontal or vertical asymmetry indicates there is no spatial chirp or tilt – i.e., the wavelength doesn't vary across the beam significantly.

(b) Temporal Phase vs Delay. Blue curve: Shows how the central wavelength of the pulse varies with delay. The linear nature of this curve, which decreases with wavelength, indicates a linear chirp in time – shorter frequencies (higher wavelengths) arrive earlier. In terms of frequency, it has a positive chirp. Red curve: Temporal phase as a function of time delay (using a polynomial fit) – how the phase of the electric field evolves in time. The quadratic nature indicates a pure linear (positive) chirp in the time domain.

(c) Spectral Phase vs Frequency. Blue curve: Ray-traced time delay vs frequency (i.e. GDD). This reinforces the fact that we see a consistent chirp, meaning that the pulse is temporally stretched in a predictable way. Red crosses: Spectral phase fitted from a polynomial. Black dashed curve: Spectral phase fitted from a Taylor expansion. These show the frequency-domain equivalent of the temporal chirp, with quadratic spectral phase indicating that the pulse is temporally broadened. Since this matches the time-domain phase (b), it confirms that the chirp is linear and well-behaved.

(d) Instantaneous Frequency vs Time. Blue crosses: Instantaneous frequency extracted from the electric field phase derivative. Black curve: Fitted polynomial to the instantaneous frequency (used to model chirp). The curve is smooth, monotonic, and the frequency increases with time, confirming a positive chirp in frequency.

By using an Öffner stretcher rather than bulk material, not only do we significantly aid the problem of large amounts of TOD accumulated through material dispersion, but we also allow ourselves flexibility to scale up to higher energies. To our knowledge, no other MIR system worldwide currently offers the capability for such a large stretch. As discussed in Section 5.5.2, where the challenges of MIR pulse generation are considered, energy scalability is a fundamental requirement for any system intended for high-intensity experiments. Pulse durations (and dispersion compensation) are naturally a huge factor in this goal, but if there is no capacity to provide more energy to a pulse, we are quickly limited in the intensities we can reach. Systems that rely on long lengths of bulk material for stretching or compression not only suffer from severe higher-order phase accumulation, but must also contend with the damage thresholds of those materials – a serious obstacle when scaling to high pulse energies. Although all systems, including ours, must contend with material damage thresholds, the capacity for a large stretch allows us to prioritise energy extraction and efficiency in the OPA process, rather than being constrained by the stretching or compression architecture. In the case of Chimera, the main energy-scaling limitation arises from how much energy can be extracted from the KTA crystals. Further energy scaling is, in principle, feasible using larger-aperture crystals; however, in practice, the maximum crystal size is constrained by current growth capabilities.

An alternative approach, demonstrated both in simulation and experiment by Flemens *et al.* [220], uses a parasitic nonlinear process in $CdSiP_2$ to quench idler amplification and thereby improve conversion efficiency in OPA systems. Reporting to boost efficiency at 3390 nm from 18% to 44% by coherently suppressing the idler through an additional parametric channel, this technique may be transferable to KTA by doping the crystal with a material that selectively suppresses the idler while remaining transparent to the pump and signal. Although such an approach could raise concerns regarding thermal loading, the authors argue that the parametric nature of the process should largely avoid dissipative heating. In our case, any residual thermal effects may be further mitigated by the relatively low repetition rate of the system (10 Hz in the final MIR OPA stages), potentially allowing thermal management via oven stabilisation or active cooling. When combined with the inherent energy scalability offered by our MIR generation and Öffner stretching architecture, this could be an interesting future path to further increase pulse energy for use on high-intensity experiments.

Once stretched, the MIR pulse is further amplified by several KTA OPA stages (including one double-pass), which inevitably introduces additional dispersion. Any compressor design must therefore also account for the cumulative TOD and FOD introduced by these amplification stages. A practical compressor counterpart to the Öffner stretcher is the Tracey compressor, which uses a pair of gratings in a simple, compact (double-passed) geometry. This design is attractive for its ability to provide substantial GDD while minimising complexity. The primary challenge lies in sourcing gratings that are both broadband and highly efficient in the mid-infrared – a region where options are severely limited.

When considering grating line (or groove) density, longer wavelengths favour smaller groove densities (i.e., larger groove spacings). This follows from the grating equation (Equation (2.46)). If λ is too large, and d is too small (i.e. the groove density is too high), the sine terms in the grating equation can exceed 1, making diffraction unphysical. At normal incidence, the maximum wavelength that can be diffracted in the first order is approximately $\lambda_{\max} = d$. In practice, however, gratings are rarely used at normal incidence. They are often designed for operation near the Littrow angle, θ_{L}, where the incident and diffracted beams are symmetric, such that:

$$\sin\theta_{\mathrm{L}} = \frac{m\lambda}{2d} \tag{5.4}$$

Operating near Littrow maximises diffraction efficiency, but also tightens the wavelength range over which only the desired order exists. Moving away from Littrow relaxes this somewhat, but additional diffraction orders can then enter at shorter wavelengths. This makes the choice of line density a balance: too low, and extra orders appear across the MIR bandwidth; too high, and the usable wavelength range is truncated. For example, for our MIR pulse, the shortest wavelengths are $\sim$ 3.4 µm. To ensure diffraction is possible at these wavelengths, the groove spacing must be at least 3.4 µm, corresponding to a groove density of approximately 294 l/mm. At Littrow, the allowable range is reduced further, meaning somewhat lower line densities are desirable in principle.

Consultation with Plymouth Grating Laboratory (PGL) led to the eventual groove density being manufactured as 320 l/mm. The reason for this design choice was a compromise between manufacturability and the predicted spectral efficiency. Although a lower groove density grating (e.g. 150 - 200 l/mm) would, in principle, make the optical design of the stretcher easier, such gratings are extremely difficult to fabricate for MIR wavelengths. Additionally, at lower groove densities, the grating can diffract some wavelengths into higher diffraction orders, particularly when working with large bandwidths, leading to energy leakage, reduced efficiency in the desired (first) order, and potential distortion of the compressed pulse. For the wavelength range of 3200 - 4100 nm, the chosen groove density of 320 l/mm was quoted to produce an overall diffraction efficiency of 90%, with a uniformity of 5% over the 1000 nm of specified bandwidth.

A numerical raytrace was once again constructed to model a grating compressor. The results of this, using a 15 cm grating separation, are shown in Figure 5.30. Regarding the grating aperture and spacing, this was chosen to balance cost with stretched pulse duration. The chosen dimensions are $50 \times 50 \times 15$ mm^3 and $70 \times 50 \times 15$ mm^3. The beam size on the first grating is set by the estimated laser-induced damage threshold, a metric which is not well established in the mid-infrared. A conservative operating limit of 50 - 100 mJ/cm^2 was therefore assumed. To remain below this fluence, the beam must be expanded sufficiently after the OPA amplifier chain, fixing the minimum grating aperture. The separation between the gratings then determines the amount of temporal

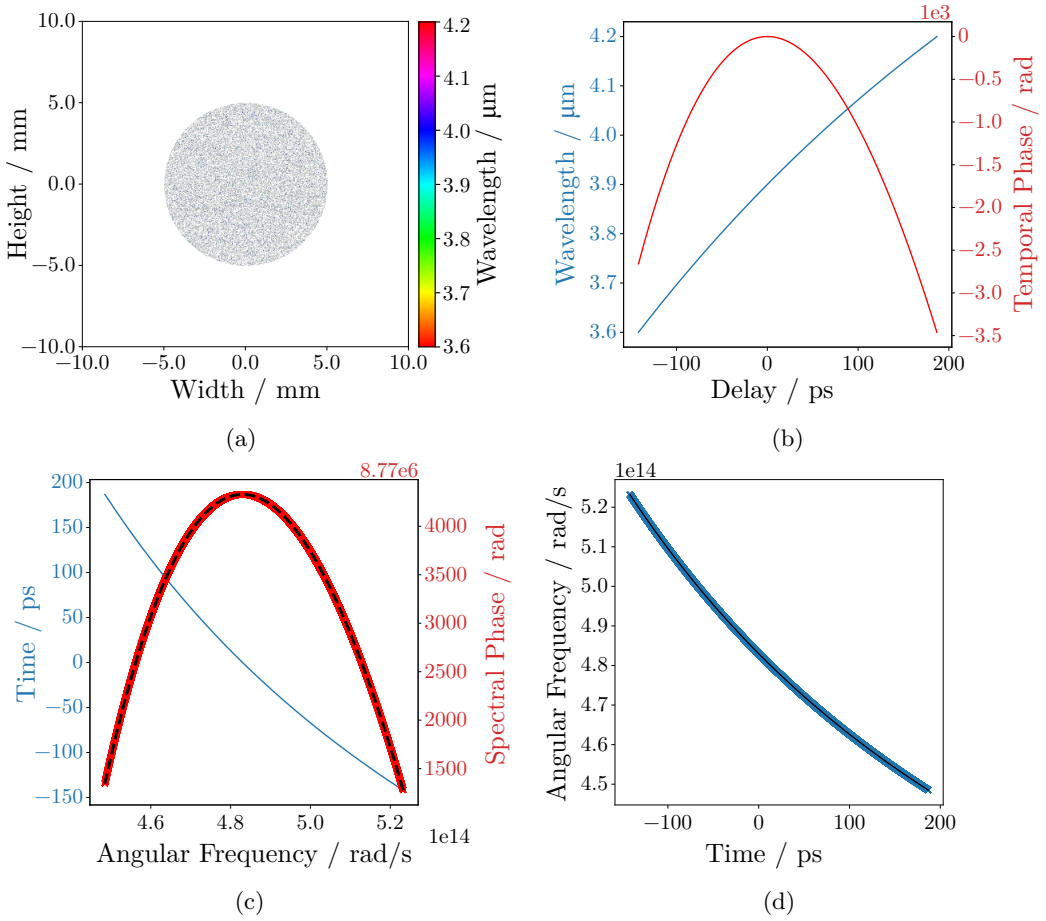
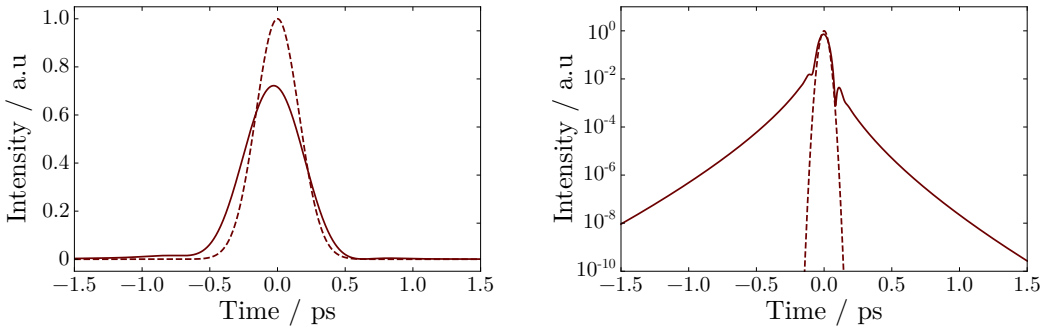


Figure 5.30: Image-on-screen and phase analysis output of the Optometrika raytracing code simulating a Tracey style compressor system for the Chimera MIR laser. The code for this Tracey compressor design was written by Dr Nicholas Stuart.
(a) 2D spatial profile of the beam on a screen. Each pixel encodes wavelength via the colour map – the spatial distribution shows how the spectrum varies across the beam profile. The beam is circular and symmetric, with no visible spatial chirp or angular dispersion.
(b) Temporal Phase vs Delay. Blue curve: Shows how the central wavelength of the pulse varies with delay. The linear nature of this curve, which increases with wavelength, indicates a linear chirp in time – higher frequencies (shorter wavelengths) arrive earlier. In terms of frequency, it has a negative temporal chirp. Red curve: Temporal phase as a function of time delay (using a polynomial fit) – how the phase of the electric field evolves in time. The clean downward-opening parabola indicates a pure linear (positive) chirp in the time domain.
(c) Spectral Phase vs Frequency. Blue curve: Ray-traced time delay vs frequency (i.e. GDD). Red crosses: Spectral phase fitted from a polynomial. Black dashed curve: Spectral phase fitted from a Taylor expansion. It is a smooth, downward-opening quadratic curve, indicating dominantly second-order spectral phase (GDD).
(d) Instantaneous Frequency vs Time. Blue crosses: Instantaneous frequency extracted from the electric field phase derivative. Black curve: Fitted polynomial to the instantaneous frequency (used to model chirp). The pulse has a clean, negative linear chirp, with the instantaneous frequency decreasing over time.

stretching and simultaneously constrains the maximum allowable footprint on the second grating to avoid clipping.

When using a Tracey-style compressor with the quoted grating apertures and 15 cm separation, the model produces a projected output beam that remains circular and free from spatial chirp. Spectral and temporal phase analyses show a clean quadratic dependence, consistent with a well-behaved linear chirp. The Taylor expansion of spectral phase coefficients (around the central wavelength) are solved to be: $\phi_2 = -4.33\times10^6$ fs^2, $\phi_3 = 4.97\times10^7$ fs^3, $\phi_4 = -9.74\times10^8$ fs^4. These well compensate the coefficients produced by propagation through the Öffner stretcher, confirming that the compressor design can compensate the stretcher phase and deliver near-transform-limited pulses. The residual spectral phase, after propagation through the the 10 cm YAG stretcher, 25 cm of KTA crystal, the Öffner stretcher and finally the Treacy compressor, is given by the coefficients: $\phi_2 = 8$ fs^2, $\phi_3 = -2.61\times10^4$ fs^3, $\phi_4 = -3.63\times10^6$ fs^4. This is illustrated in Figure 5.31. These phase terms broaden a 50 fs, 3.9 µm pulse to 68.6 fs, and are something that will require a full experimental characterisation, through the use of either a MIR FROG or SPIDER.



(a) (b)

Figure 5.31: The effect of high order spectral phase on the compressibility of an ultrashort pulse. For Chimera parameters of $\lambda = 3.7$ µm, and a pulse duration of 50 fs, we show the spectral phase terms acquired through propagation through a 10 cm YAG stretcher, 25 cm of KTA crystal and the Öffner stretcher, before being recompressed in a Treacy-style compressor. This corresponds to $\phi_2 = 8$ fs^2, $\phi_3 = -2.61\times10^4$ fs^3, and $\phi_4 = -3.63\times10^6$ fs^4. Dashed curves represent the transform-limited pulses, while solid curves show the effects of high-order phase. (a) Shows this on a linear intensity scale, while (b) uses a logarithmic scale. With the addition of higher-order dispersion, the compressed pulse is limited to a minimum FWHM duration of 68.6 fs.

Although we succeed in the retention of a relatively Gaussian temporal pulse envelope, Figure 5.31b) shows a significant intensity pedestal due to the residual high-order phase, which may limit applications where high temporal contrast is critical, such as when using thin or nanostructured solid targets.

To test the compression setup, custom grating mounts were designed and 3D printed. These mounts were built to fit within standard 4-inch precision kinematic mirror mounts, allowing for quick integration onto the optical table. This means that in addition to the horizontal and vertical plane adjusters, we can also set the rotational position of

the gratings. The mounts are designed so that the grating slides in from the side and is secured using three nylon-tipped grub screws. These are threaded into hand-tapped M6-to-M4 adapters, which are glued into the plastic body of the mount. A ~ 2 mm solid silicone rubber sheet is inserted between the grating and the plastic surface to protect the grating from excessive point clamping loads when in direct contact with the plastic. A thin metal load-bearing plate is installed on the top surface to further reduce mechanical stress from being concentrated by the grub screws. These mounts will eventually be replaced with custom stainless steel versions to reduce mechanical drift and thermal expansion. The CAD grating mount design is shown in Figure 5.32.

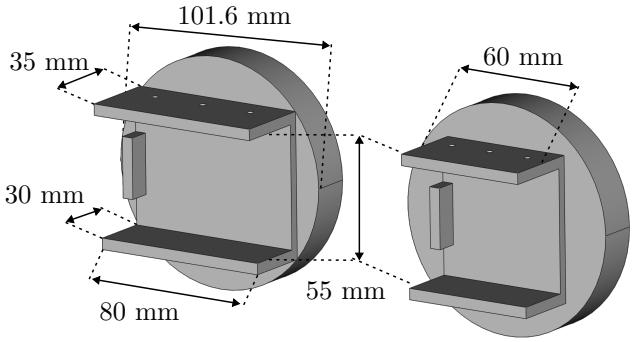


Figure 5.32: The compressor grating 3D printed mount designs. The gratings have dimensions $70 \times 50 \times 15$ mm^3 (left) and $50 \times 50 \times 15$ mm^3 (right). The three holes in the upper support (deliberately protruded to protect against dust) are hand-tapped to M6 after printing.

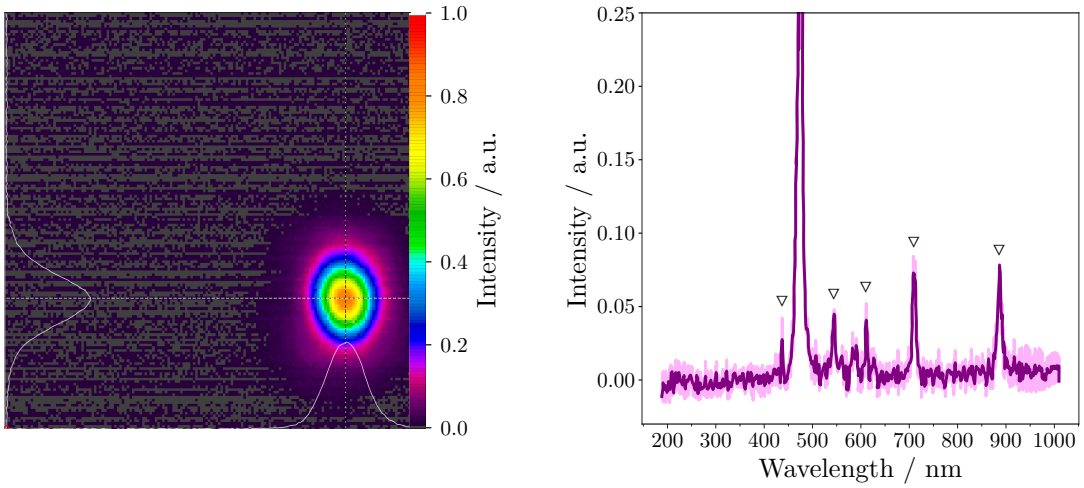
Although the system is not yet ready to receive an incident MIR beam, once the OPAs return to regular operation, we will be in a strong position to begin tests towards MIR pulse compression, enabling the commissioning of high-intensity experiments.

5.7 Experimental Output of the Chimera Laser System

Despite currently lacking the intensity capability required for high-intensity laser–plasma experiments (e.g. laser wakefield acceleration, requiring intensities of $\mathcal{O}(10^{18}$ W/cm$^2)$), the mid-infrared beam from the Chimera laser can be commissioned for experiments that demand lower intensities. Although we do not yet have the capacity to directly measure the MIR pulse duration (meaning that the laser intensity on target remains an estimate), experimental commissioning has begun through solid high-harmonic generation (HHG), marking the first experiment performed on the Chimera laser system since its conception.

With a pulse energy of 90.8 µJ measured by the Pyrocam, and a D4σ beam size of 5.20×4.54 mm^2, the beam was picked off from the first MIR OPA stage and focused, using a $f = 50$ mm PCX CaF_2 lens, onto a $10 \times 10 \times 0.3$ mm^3 sample of ZnSe. From this, we estimate, using the 100 Hz repetition rate of the system, an average intensity of 4.09 W/cm^2, assuming a spot diameter at focus of 47.1 µm. For a pulse duration in the range 100 - 200 fs, this corresponds to a peak intensity on target of 2.05 - 4.09 $\times 10^{11}$ W/cm^2. The MIR beam profile, along with the HHG spectrum collected

from the sample, is shown in Figure 5.33. Although a typical HHG spectrum contains odd harmonics only, due to ZnSe being a non-centrosymmetric crystal, it presents even harmonics as well.



(a) (b)

Figure 5.33: (a) A beam profile of the amplified MIR beam (centred at 3700 nm) after the first MIR OPA stage (OPA-B1), before being sent to a solid sample of ZnSe, measured with a Spiricon Pyrocam IIIHR beam profiler. Axis lineouts were recorded using Spiricon BeamGage software, giving D4σY and D4σY of 5.20 mm and 4.54 mm, respectively. Here, D4σ denotes the ISO standard 'second-moment' beam width, calculated as four times the standard deviation of the intensity distribution. The software also measured a pulse energy of 90.8 µJ and a peak fluence of 2.65×10^3 µJ/cm^2.
(b) High harmonic spectral data obtained when directing a 1000 nm bandwidth 3700 nm pulse, with duration 100 - 200 fs, at a sample of crystalline ZnSe. The pink trace shows the raw data, recorded on an Ocean Optics USB2000+ spectrometer and reconstructed from a digital image of the spectrum. The purple shows a smoothed version of the raw data, using a finite impulse response (FIR) filter. The prominent central feature at 470 nm is attributed to photoluminescence from the sample. Strong 4$^{\text{th}}$ and 5$^{\text{th}}$ harmonic signals at $\sim$ 890 nm and $\sim$ 710 nm are seen. Weaker signals are seen at $\sim$ 610 nm, $\sim$ 545 nm and $\sim$ 435 nm, possibly representing 6$^{\text{th}}$, 7$^{\text{th}}$, and 8$^{\text{th}}$ harmonics. Harmonic signals are indicated on the plot with triangles.

A prominent signal is observed at 470 nm in the spectral data, which we attribute to photoluminescence in the sample. This was accompanied by a visible blue glow around the sample *in situ*, consistent with the 2.70 eV band gap of ZnSe [221]. Although the MIR beam is centred around 3700 nm, at this stage of the system, it exhibits a broad bandwidth of nearly 1000 nm. Strong 4$^{\text{th}}$- and 5$^{\text{th}}$-order harmonic signals, observed at $\sim$ 890 nm and $\sim$ 710 nm respectively, indicate that the spectral region around 3500 nm contains the highest intensity, and is the central wavelength generating the harmonic spectrum. We also observe weaker features at $\sim$ 610 nm, $\sim$ 545 nm, and $\sim$ 435 nm, which could correspond to HHG reaching the 8$^{\text{th}}$ harmonic.

Using a peak fitting procedure to detect significant local maxima in the HHG spectrum, we can match these to integer multiples of a chosen fundamental energy (within a 0.07 eV tolerance, corresponding to a 0.2 tolerance in harmonic order), and thus estimate the

implied photon energies, as well as the approximate bandwidth of each harmonic. For this, we choose a fundamental energy of 3500 nm (equivalent to 0.35 eV photon energy), which better reflects the highest-intensity region of the MIR spectrum (see Figure 5.20), rather than the central wavelength of 3700 nm. The resulting harmonic structure as a function of photon energy is shown in Figure 5.34.

With each MIR pulse containing approximately 10^{15} photons at 0.35 eV, the ZnSe sample thickness of 0.3 mm defines a finite interaction length over which the nonlinear polarisation responsible for HHG can accumulate. The observation of multiple harmonic orders within this thickness indicates that the harmonic emission per unit length is sufficiently strong to overcome both absorption and out-coupling losses, yet remains weak compared to competing channels such as photoluminescence. This constrains the HHG process to a low-efficiency but finite regime, consistent with an order-of-magnitude conversion efficiency estimate in the $10^{-9} - 10^{-6}$ range. To obtain the true conversion efficiency, a measurement of pulse energy would have to be obtained, both before and after interacting with the sample. For context, literature reports on ZnSe indicate high second-order (SHG) conversion efficiencies of up to $\sim 35\%$ under phase-matched conditions [**222**, **223**]. These values demonstrate the strong intrinsic nonlinearities of ZnSe; however, they cannot be directly compared to the results reported here. The spectrometer on which these results were recorded does not cover the wavelength range required to measure the second-order harmonic, and conversion to the fourth and higher orders is not a simple extension of SHG. The efficiency of higher-order harmonics is governed by strongly nonlinear processes, influenced by absorption, phase-mismatch, and non-perturbative electron dynamics, preventing any straightforward extrapolation from low-order values.

At this stage, the laser stability is not sufficient to measure intensity dependence or spectral broadening of the MIR harmonics, such as those reported by Chin *et al.* in 2001 [**222**]. Nevertheless, this would be an interesting avenue to pursue in future Chimera experiments when more stable operation is restored. In both [**222**] and the more recent work of Werner *et al.* in 2019 [**223**], experiments were performed with thicker ZnSe samples than used here, and in both cases an increase in harmonic signal and efficiency with thickness was reported. This, as well as investigating additional solid-state semiconductors, should form part of future iterations, particularly given that the energy available to us (even from the first front-end MIR OPA of Chimera) far exceeds the ~ 35 µJ per pulse reported in the literature. When the final-stage MIR OPAs return to regular operation, this experiment could also be repeated both before and after compression, offering a wide parameter space to investigate.

Looking ahead, we aim to extend high harmonic generation experiments from the Chimera laser beyond purely spectral measurements. Once pulse intensities reach the threshold for intense laser-plasma interactions, a full characterisation of the MIR pulse will become essential. As the MIR spectral range is notoriously difficult to detect directly, this is something that remains challenging and raises the question: can the visible spectra

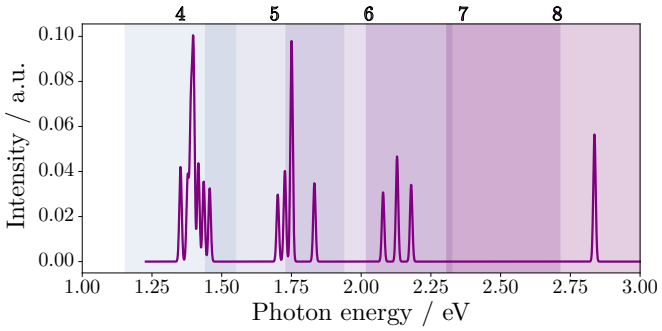


Figure 5.34: Measured MIR harmonics as a function of photon energies. Significant local maxima in the HHG spectrum were identified by a peak-detection routine and matched to integer multiples of a fundamental energy, assuming a central wavelength of 3500 nm (0.35 eV) with a tolerance of 0.2 in harmonic order (equivalent to 0.07 eV). Harmonic bands are indicated using the colour map, where the harmonic order is shown for the wavelength range of 3200 - 4300 nm.

from solid HHG provide indirect experimental information about the MIR pulse, serving as a non-destructive diagnostic at the point of interaction? A recent publication from Awad *et al.* [224], demonstrating the use of HHG in ZnO as a pulse-duration diagnostic at 3.2 µm, suggests that extending this approach into our spectral regime may indeed be possible.

Given our initial demonstration of high harmonics on the Chimera laser, further exploration of this method need not wait for the return of the entire MIR OPA chain or for the successful compression of our MIR pulse. Rather, it depends on the construction of one or more pulse-duration diagnostics, such as a MIR FROG, SPIDER, or autocorrelator, with which the results can be corroborated. If this approach proves viable, then, combined with spectral beam-profile characterisation via three-photon absorption in silicon chip cameras, high-intensity MIR-driven experiments could become far more diagnosable than previously anticipated.

5.8 Multi-beam Capabilities of the Chimera laser system

In this chapter, we have focused primarily on the main mid-infrared (3.7 µm centred) pulse of the Chimera laser system; however, this is not the full extent of Chimera's eventual capabilities. Designed as a multi-beam platform, although only the MIR arm is currently in full operation, Chimera contains two additional wavelength arms at 0.8 µm and 1.5 µm.

The 0.8 µm arm is the most straightforward candidate for commissioning, owing to its more conventional wavelength. Indeed, the Chimera project began with the re-use of much of the hardware from a previous 0.8 µm OPCPA laser, with the initial intention being to 'simply' amplify the idler at ~ 2 µm; however, pushing further into the MIR was found to hold more interesting scientific opportunities. The 0.8 µm seed originates from the split following the 0.8 µm OPCPA subsystem: of the ~70 µJ available after amplification, 50% is directed to generate the MIR seed while the remainder is sent to

a diagnostics line. While this diagnostics branch is of great use for day-to-day system operation, it is also an obvious point for an upgrade via further amplification, either using standard Ti:Sa gain media or, for higher contrast, in an additional series of OPA stages pumped at 527 nm. At present, this line also incorporates a NIR Dazzler despite spectral phase control not being required for diagnostics. In an upgraded configuration, however, the Dazzler would be much better placed in the 0.8 µm arm, where it could provide dispersion management and enable compression to few-cycle durations.

The 1.5 µm arm presents another promising upgrade pathway. This wavelength appears naturally as the idler in the MIR OPA process – effectively provided 'for free' by energy and momentum conservation. In the earlier OPA stages (OPA-B1 and OPA2), this 1.5 µm light will be used for diagnostics, as it is easier to detect and characterise than the MIR. In the final stage (OPA3), however, there is the opportunity to exploit this higher-energy idler as a secondary beamline. With the addition of a dedicated 1.5 µm compressor, this pulse could be commissioned as an electronically and temporally synchronised probe for experiments driven at 3.7 µm. An outstanding challenge, however, would be to design and build a suitable compressor as its chirp would inherit the opposite sign of the seed.

If fully realised, Chimera's multi-beam architecture would establish it as a truly unique system, enabling high-intensity, multi-beam experiments at the university-laboratory scale.

5.9 Summary

In this chapter, we have described the design, development and characterisation of the Chimera multi-beam mid-infrared OPCPA laser system at Imperial College London. We presented the Chimera front end, including the Femtolasers Rainbow oscillator, the wavelength-shifting photonic crystal fiber, and the regenerative amplifiers, followed by the nonlinear amplification stages: an 800 nm OPCPA subsystem, four MIR OPA stages (one operated in double-pass), and the flashlamp-pumped amplification chain. Particular attention was given to improvements made to the oscillator and ring booster, as well as unresolved peculiarities observed in the linear regenerative amplifier. We also discussed in detail the design considerations for generating the MIR pulse, together with its stretching and compression architecture, placing this in the context of other state-of-the-art high-intensity MIR systems reported in the literature. Finally, we described Chimera's first commissioning experiment: in a high-harmonic generation experiment using a crystalline ZnSe sample, we observed signals consistent with harmonics up to the 8^{th} order. In addition to the main MIR arm, Chimera was designed with two further beamlines at 0.8 µm and 1.5 µm. Although not yet commissioned, these offer clear upgrade pathways: the 0.8 µm arm could be amplified and compressed for few-cycle operation, and the 1.5 µm idler from the MIR OPAs could be exploited as a synchronised probe. Realising this multi-beam capability would establish Chimera as a uniquely versatile platform for high-intensity experiments.

6 Mid-Infrared Laser Wakefield Acceleration Simulations : An exploration of high electron energies

The challenge of creating a high-intensity mid-infrared short-pulse laser has led to very few experimental or computational results in the literature to which a comparable baseline may be established. This allows us a certain amount of freedom in our investigations. Before we commence any kind of experimental planning, we may instead turn to a simulation campaign purely to satiate a curiosity. The query to resolve is: will a mid-infrared (MIR) driven laser wakefield acceleration (LWFA) experiment produce comparable results to a more typical (and much more extensively documented) 800 nm experiment?

In asking this question, we must ask ourselves what exactly we are doing by changing the wavelength of the driving laser. When expressed using normalised quantities, as is standard in particle-in-cell (PIC) simulations, Maxwell's equations are dimensionless and therefore parametrically scalable. This suggests that the wavelength itself has little direct effect; instead, its influence arises through how it reshapes the relationships between other laser and plasma parameters. In principle, a change in wavelength would be inconsequential if all other quantities were scaled accordingly; in practice, however, such perfect similarity is rarely achievable, and it is precisely these departures that make longer-wavelength drivers distinctive.

Considering first what an increase in laser wavelength could potentially mean, we can look at how this will affect other experimental parameters. Beginning with the normalised vector potential, a_0, and the ponderomotive force, $\mathbf{F}_\mathrm{p}$ (Equations (3.4) and (3.17)), there is a strong wavelength scaling, such that:

$$|\mathbf{F}_\mathrm{p}| \propto a_0^2 \propto \lambda^2 \tag{6.1}$$

Although this implies a very positive wavelength scaling towards how strongly we can drive a nonlinear wakefield, the expression is naïve: a_0 also depends on laser energy, which, in practice, is far lower in millijoule-class mid-infrared systems compared to joule-class 800 nm systems.

Using the dispersion relation (Equation (3.25)), we also define a critical plasma density n_c (the density at which the plasma frequency ω_p is equal to the frequency of the propagating laser within the plasma ω_L) as:

$$n_c = \frac{\epsilon_0 m_e \omega_L^2}{e^2} \tag{6.2}$$

A change in wavelength will then necessarily impact the critical plasma density, such that:

$$n_c \propto \frac{1}{\lambda^2} \tag{6.3}$$

In a scenario where we change the wavelength of an experiment, we take this to imply that dimensionless parameters, such as the ratio of plasma density to a critical plasma density n_e/n_c and the normalised laser vector potential a_0, should stay the same. For n_e/n_c and a_0 both constant, this implies the relationship $\lambda_p \propto \lambda$, with:

$$\frac{n_e}{n_c} \propto \mathrm{C} \propto \frac{\omega_p}{\omega_L} \propto \frac{\lambda_L}{\lambda_p} \tag{6.4}$$

Such that longer-wavelength lasers drive proportionally larger plasma bubbles. Using Equations (3.48) and (3.49), it also implies longer dephasing and depletion lengths for longer wavelength drivers.

What this thought experiment lacks is the fact that this whole scenario is simplistic – a longer wavelength laser system is simply not a scaled-up version of a shorter one. To provide some perspective, let us consider a 'standard' 800 nm LWFA experiment, using perhaps the Gemini laser at the Rutherford Appleton Laboratories, as well as a less conventional one, for example, the Salle Noire laser at the Laboratoire d'Optique Appliquée (LOA). Using [225], the Gemini laser has a wavelength and pulse duration of 800 nm and 43 fs, respectively, with 11 J of energy on target. This would produce a peak power P_0 of 260 TW. By contrast, the LOA system operates at the same wavelength but with a much shorter pulse of 3.4 fs and 2.1 mJ on target, yielding a peak power of 0.6 TW [66]. The Chimera laser has a longer wavelength of 3700 nm, a target pulse duration of 30 fs, and 25 mJ of energy on target, producing a peak power of 0.8 TW.

These examples highlight the complexity of the parameter space: comparing experiments by wavelength alone is misleading, since factors such as pulse duration and energy strongly influence the resulting dynamics. Instead, we consider only dimensionless quantities – the ratio of plasma density to a critical plasma density n_e/n_c, the ratio of laser power to a critical laser power P/P_{crit}, the normalised laser vector potential a_0, and the number of optical cycles contained within the laser pulse – to provide some amount of reasonable comparison to other LWFA experiments previously reported in the literature. Here P_{crit} is the critical power – the threshold power for a laser to relativistically self-focus in a plasma.

This is defined as:

$$P_{crit} = \frac{8\pi\epsilon_0 m_e^2 c^5 \omega_L^2}{e^2 \omega_p^2} \propto \frac{\omega_L^2}{\omega_p^2} \propto \frac{\lambda_p^2}{\lambda_L^2} \propto \frac{n_c}{n_e} \tag{6.5}$$

Where the subscripts 'L' and 'p' refer to laser and plasma, respectively.

Using the relation for plasma frequency $\omega_p = \sqrt{n_e e^2/\epsilon_0 m_e}$, (and taking an experimental plasma density of 2.8 $\times 10^{18}$ cm^{-3}), we can calculate and compare relevant parameters for the Gemini, Salle Noire and Chimera lasers. These are displayed in Table 6.1.

	Gemini	Salle Noire	Chimera
Wavelength λ	800 nm	800 nm	3700 nm
FWHM Pulse Duration t_{FWHM}	43 fs	3.4 fs	30 fs
Laser Energy E_L	11 J	2.1 mJ	25 mJ
Number of Cycles	16.1	1.3	2.4
Power P	260 TW	0.6 TW	0.8 TW
Plasma density n_e	2.9 $\times 10^{18}$ cm^{-3}	1.2 $\times 10^{20}$ cm^{-3}	2.8 $\times 10^{18}$ cm^{-3}
Critical Power P_{crit}	10.8 TW	0.249 TW	0.501 TW
P/P_{crit}	23.7	2.41	1.58
Critical Density n_c	1.7 $\times 10^{21}$ cm^{-3}	1.7 $\times 10^{21}$ cm^{-3}	8.0 $\times 10^{19}$ cm^{-3}
n_e/n_c	0.002	0.07	0.03
FWHM spot size d_{FWHM}	42 µm	3.5 µm	13.6 µm
a_0	2.4	1.4	1.9

Table 6.1: A comparison of parameters for a LWFA experiment driven by either the Gemini, Salle Noire or Chimera laser systems. The numbers used for the Gemini and Salle Noire parameters are taken from [225] and [66]. The Chimera parameters are those used in this simulation campaign.

This exercise highlights that the problem spans a truly multidimensional parameter space. For a direct comparison, the Chimera parameters would need to be drastically altered.

To match the Gemini laser, we would need to include changes such as a 138 fs pulse duration (to provide the same number of optical cycles). To reach the same P/P_{crit} would require a peak power of 12 TW. With the longer pulse duration, this would mean

having a laser energy of 1.7 J – a significantly higher energy than is currently achievable in the mid-infrared, due both to the lack of brute force gain storage amplifiers in that wavelength range and the challenge of growing large enough crystals for a high-energy OPA process. To reach the same a_0 (with the above parameters), this would then mean an extremely large FWHM spot size of 4.2 cm. Essentially, then, for a 3.7 µm driven LWFA experiment to take place under the same conditions as an 800 nm one, driven by Gemini, we would have to commit drastic changes, most of which laser technology cannot yet meet.

The LOA laser, on the other hand, may provide a more sensible comparison. Like Chimera, it operates in the few-cycle regime, which makes its parameter space (aside from wavelength) appear closer to our own. Providing the same number of optical cycles would mean a pulse duration of 16 fs, a peak power of 1.2 TW and a laser energy of 19 mJ. Disregarding the technical challenge of further reducing the pulse duration within the MIR, these parameters are not dissimilar to Chimera's design regime. The main difference lies in the normalised vector potential a_0. Due to its strong wavelength dependence, even modest changes in other parameters translate into dramatic variations in a_0. For these Salle Noire-adjusted laser parameters, our original simulated spot size of 13.6 µm would result in an $a_0 \approx 2300$ – to be comparable to LOA, our spot size would have to increase to 22 mm, which would preclude strong wake excitation, self-guiding, and operation in the blowout regime. Rather than representing a straightforward geometric upscaling, this illustrates that the strong wavelength dependence of a_0 fundamentally constrains which combinations of intensity, spot size, and plasma density remain physically meaningful at longer wavelengths. The MIR regime therefore accesses a qualitatively different region of parameter space, rather than a scaled version of the 800 nm case.

In short, we are now delving into a completely different regime than that which we have come to understand and expect from our regular LWFA experiments using 800 nm driving lasers – and thus should compare the two with some caution.

6.1 Plasma Density Dependence of Accelerated Electron Energy

In 2007, a phenomenological framework of laser wakefield acceleration in the 3D nonlinear regime was set out by Lu *et al.* [74]. Discussed in more detail in Section 3.3.5, this pioneering paper established a set of scaling laws to which the vast majority of experiments still conform. Amongst these scaling laws is one to determine the maximum energy $W_{\max}$ reached by accelerated electrons in a 3D nonlinear LWFA regime. This is given by:

$$W_{\max}/m_{\mathrm{e}}c^2 = \frac{2}{3}a_0\frac{n_{\mathrm{c}}}{n_{\mathrm{e}}} \tag{6.6}$$

We use this relationship as a starting point in this simulation campaign: by scanning through the starting plasma density, we can track how it affects $W_{\max}$. To begin our investigation, we chose variables which we (optimistically) predict the Chimera laser

system to reach. For a FWHM spot size of $d_{\mathrm{FWHM}} = 13.6$ µm ($w_0 = 11.6$ µm), pulse energy 25 mJ and FWHM pulse duration of ~ 30 fs, this corresponds to an a_0 of 1.9. We note that this is not the matched spot size – the optimal spot size suggested by simulations for which laser-self guiding may occur is density dependent, and is given by $w_{\mathrm{m}} = 2\sqrt{a_0}\ c/\omega_{\mathrm{p}}$, where ω_{p} is the plasma frequency [226]. The full set of parameters used in this density scan is given in Table 6.2.

Wavelength	3.7 µm
FWHM (intensity) beam waist	13.6 µm
FWHM (intensity) pulse duration	30 fs
Pulse energy	25 mJ
Polarisation	Linear, out of the simulation plane (z)
Temporal Pulse Shape	Gaussian
a_0	1.9
Critical density	8.0×10^{19} cm^{-3}

(a) Laser Parameters

Particle density	1.8 - 5.0 $\times 10^{18}$ cm^{-3}
Profile	Trapezoidal density ramp
Plateau length	800 µm
Ramp length	100 µm at beginning and end

(b) Plasma Parameters

Geometry	Azimuthal Mode Cylindrical
Number of azimuthal modes	2
Particles per cell	10
Box size	94 $\times$ 57 µm
Resolution	$\lambda/100 \times \lambda/20$
Timestep	0.060446 = 0.999977 x CFL
Initial Electron Temperature	0 K

(c) Simulation Parameters

Table 6.2: The full set of parameters used to simulate the Chimera laser system using the PIC code Smilei.

Here we scan through densities $n_{\mathrm{e}} = 1.8 - 5.0 \times 10^{18}$ cm^{-3}, keeping all other parameters constant. To extract energies of the accelerated electrons, we use a particle tracking diagnostic to record the positions and momenta of macroparticles over the course of the simulation. Limiting this to track only macroparticles reaching energies greater than 5 MeV (to exclude noise from background plasma electrons), we can extract particle energies over all simulated timesteps. Figure 6.1 shows an example electron waterfall plot (energy-time distribution) and energy spectrum, with Figure 6.1(a) highlighting the evolution of the maximum and 90$^{\mathrm{th}}$ percentile macroparticle energies, $E_{x\mathrm{max}}$ and $E_{x90\%}$, along the longitudinal direction, at each timestep of the simulation.

Looking at this plot, we can see that there is a wide spread in electron momenta, and that identifying the very highest accelerated electron energies is perhaps not very

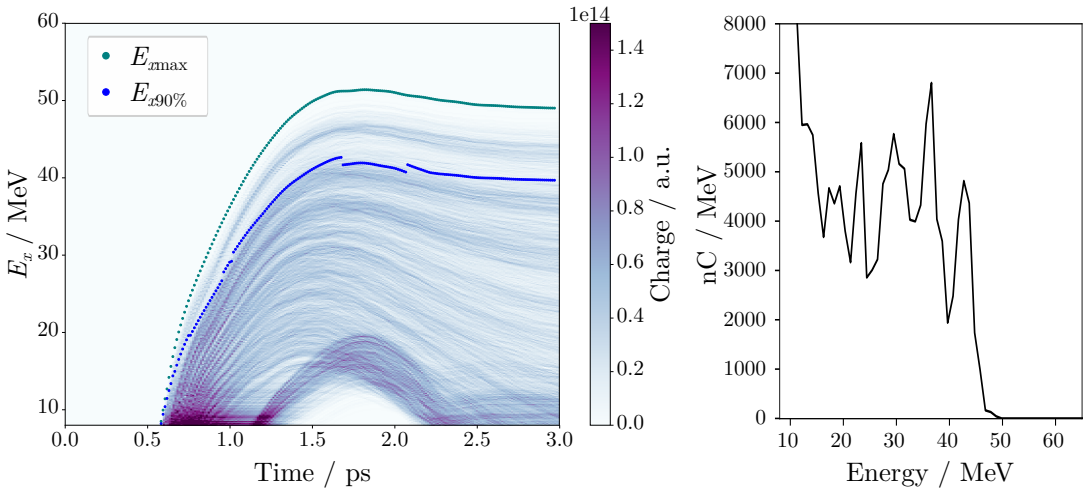


Figure 6.1: (a) Electron waterfall plot (energy-time distribution) for a LWFA simulation at a starting plasma density of 2.8×10^{18} cm^{-3} and using simulation parameters given in Table 6.2. The maximum ($E_{x\max}$) and 90$^{\text{th}}$ percentile ($E_{x90\%}$) electron energies E_x of distributed charge are extracted and overlaid onto the waterfall plot.
(b) Corresponding electron energy spectrum, expressed as the differential charge distribution $\mathrm{d}Q/\mathrm{d}E$.

physical due to the small percentage of electrons (or numerical noise) reaching this point. We will choose to work instead with energies at the 90$^{\text{th}}$ percentile of distributed charge ($E_{x90\%}$). We can also extract additional information from this plot regarding the injection events of the accelerator. The first injection event occurs very early on, at roughly 0.6 ps, when the laser self-focuses to a minimum spot size. Due to the disorder and random nature of the self-injection mechanism, this first injection event is quickly followed by many additional injections, which reach decreasingly lower energies as the simulation progresses. There is one further injection event of note, at about 1.2 ps into the simulation.

We can corroborate these injection events in the electron energy spectrum with the transverse laser profile seen through the course of the simulation. Figure 6.2 is a 3D waterfall plot (with an intensity scale relative to its maximum intensity) showing the evolution of the transverse FWHM beam waist of the laser, simulated with a starting density of 2.8×10^{18} cm^{-3}. With the colour scale representing the FWHM spot size d_{FWHM}, we see evidence of the laser self-focussing into the plasma, reaching a minimum spot size of 10.13(1) µm at 0.55 ps. This matches well with the first injection event. The laser then appears to self-guide for roughly 1 ps before starting to self-focus again. At the very end of the simulation, the laser rapidly approaches total energy depletion and diffraction effects take over. The cause of the second injection event is unclear and can not be unpicked by the transverse laser evolution alone.

Extracting the highest energy along the $E_{x90\%}$ line, and equating this to the $W_{\max}$ defined in Equation (6.6), we find a $W_{\max}$ of 44.3 MeV, 1.8 ps into the simulation. We

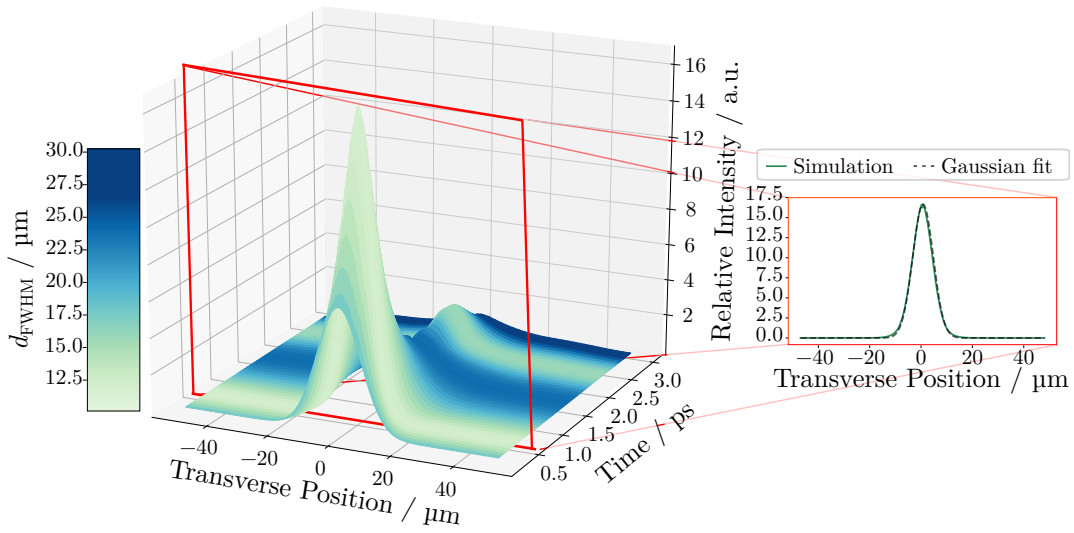


Figure 6.2: A 3D waterfall plot showing the transverse laser evolution for a simulation with starting density 2.8×10^{18} cm^{-3}. The intensity of the pulses is normalised relative to the highest intensity. The red rectangle highlights the minimum FWHM spot size position at 0.55 ps, where, using a Gaussian fit, we find the laser to self-focus down to 10.13(1) µm.

repeat this analysis for each starting plasma density in the parameter scan, obtaining the maximum energy reached by the 90$^{\text{th}}$ percentile of charge over the course of each simulation. These values are plotted as pink diamonds for densities in the range $1.8 - 5.0 \times 10^{18}$ cm^{-3} in Figure 6.3. As a cross-check, we apply a second extraction method. The first method tracks a random selection of individual electrons, taking the highest energy reached by the 90$^{\text{th}}$ percentile distribution, shown in blue in Figure 6.1. The second method (orange circles in Figure 6.3) identifies the 90$^{\text{th}}$ percentile of charge directly from the binned energy spectrum (as in Figure 6.1(b)) for each simulated density. Whilst this second method includes all simulated macroparticles (unlike the first case), it comes at the expense of energy resolution.

Plotted alongside $E_{x90\%}$, we include the aforementioned Lu energy scaling, given in Equation (6.6), and see that by using an initial a_0 of 1.9, this fares extremely poorly compared to the extracted energies of our dataset. Indeed, we see that the Lu scaling is only followed if we impose an a_0 of 4.7.

This result is extremely unexpected. The Lu framework is one which is very well established within the LWFA community for work at 800 nm – a very large proportion of simulation and experimental results agree with these scalings remarkably well, and the framework is one used by many, both to plan and interpret their experiments.

It is, of course, both possible and expected that there will be some evolution of the propagating laser corresponding to a larger a_0 than that which we have inputted. This evolution occurs due to the self-focussing and temporal compression of the pulse as it propagates (see Sections 3.2.2 and 3.2.3). If we assume that the pulse will evolve to its

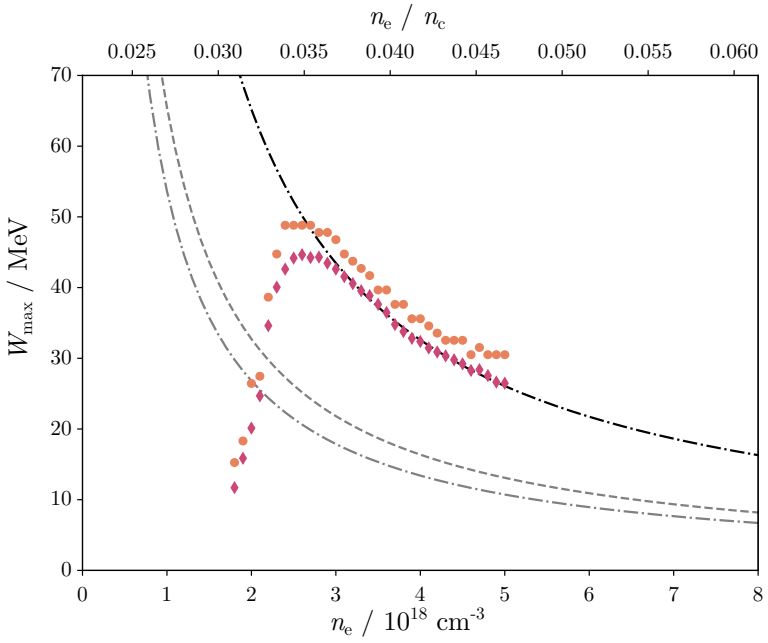


Figure 6.3: The 90th percentile of accelerated electron energies for each starting density $1.8 - 5.0 \times 10^{18}$ cm^{-3}, obtained using two methods. Pink diamonds: tracking a random selection of electrons and taking the maximum energy reached by the 90th percentile distribution during the simulation. Orange circles: identifying the 90th percentile of charge from the binned electron energy spectrum at each density. The grey and black dot-dash curves show the scaling predicted by Lu *et al.* [74] for $a_0 = 1.9$ and $a_0 = 4.7$, respectively, while the grey dash curve shows the Lu energy scaling predicted when using an evolved $a_{0\mathrm{max}} = 2.36$, calculated using Equation (6.7).

matched spot size (corresponding to a maximum $a_{0\mathrm{max}}$), we can use another of Lu's scalings, $a_0 \simeq 2(P/P_c)^{1/3}$ [74], to extract what we anticipate $a_{0\mathrm{max}}$ to be. From this (and taking an example starting density of 2.8×10^{18} cm^{-3}) we calculate an $a_{0\mathrm{max}}$ of 2.34 – only a fraction higher than the inputted a_0 at the start of the simulation and still below the a_0 which might explain our high electron energies. Thinking further about what happens to the laser as it evolves, Bloom *et al.* [150] proposed an extension to Lu's initial scaling. They do this by taking the laser power P to be some final power P_f, depending both on a pulse temporally compressed down to some final pulse duration τ_f and the fraction α of laser energy E_L which is trapped within the bubble, resulting with the final expression:

$$a_{0\mathrm{max}} \simeq 2(P_f/P_c)^{1/3} \simeq 2(\alpha E_L/(\tau_f P_c))^{1/3} \tag{6.7}$$

Taking an α of 0.5 (and noting that this is the maximum for a Gaussian pulse and thus would be considered relatively high in an experiment), we calculate an evolved $a_{0\mathrm{max}}$ of 2.36 (included in Figure 6.3 in grey dash) and are no further into our investigations than before.

To highlight the unusual nature of the results that we present, we take the same data plotted in Figure 6.3 and normalise it to a maximum predicted energy W'_{max}, using

Equation (6.6) with a_0 updated to the maximum $a_{0\max}$ in Equation (6.7). This allows us to directly compare our results to a wide range of published data, including both experiments and simulations. This is shown in Figure 6.4. The publication depository from which all the presented results are taken can be found at [227].

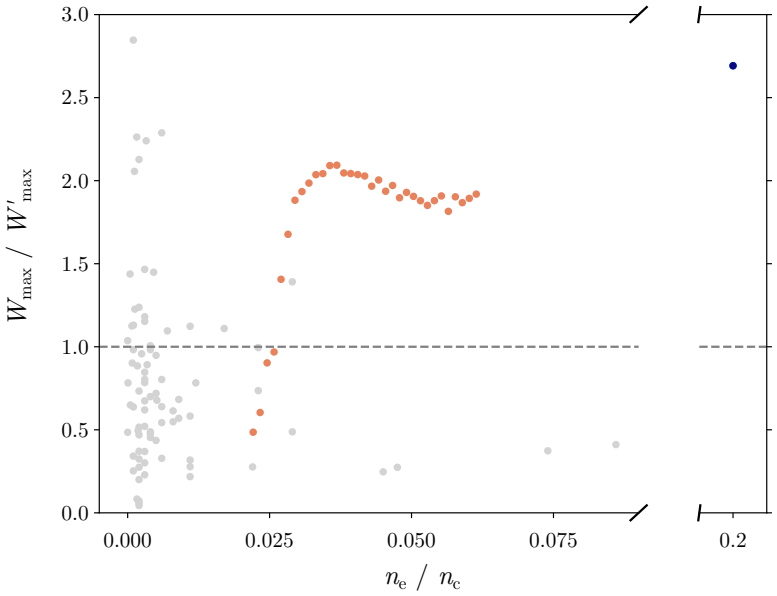


Figure 6.4: The highest electron energy with plasma density for a selection of laser-driven wakefield experiments. The grey data points represent publications reporting both simulation and experimental results. These data are from references which have been placed in a depository found at [227]. The orange data points show the simulation results presented in this chapter: the highest energy reached by the 90^{th} percentile distribution of binned charge in the electron energy spectrum for simulated densities between $1.8 - 5.0 \times 10^{18}$ cm^{-3}. Plotted in dark blue, we also identify the only comparable LWFA experiment reported thus far to the simulations we present – driven by a 3.9 µm laser at the University of Maryland [185]. The y-axis shows the maximum reported energy reached by accelerated electrons, normalised using the maximum predicted energy $W_{\max}$ from Equation (6.6), with $a_{0\max}$ calculated from Equation (6.7). The reference line at $y = 1$ indicates what would be a perfect agreement between the results and the updated model.

Here we see that the simulation data presented in this chapter (orange circles) appear to stand out markedly from the rest. While differences in n_e/n_c can in part be attributed to the change in laser wavelength, most published results (with the exception of this work and a few others) lie below the $y = 1$ line, indicating some 'underperformance' regime. This is most likely due to effects such as laser depletion or dephasing, preventing electrons from reaching the theoretical maximum energy. Conversely, publications lying above $y = 1$, and especially those above $y = 2$, may be considered to 'over-perform' relative to Lu's scaling.

In addition to our own results, we highlight the dark blue point in Figure 6.4, which represents, to our knowledge, the only experimental LWFA result obtained with a driving laser of comparable characteristics to those modelled here. This result, published by Woodbury *et al.* in 2018 [185], with laser parameters of $\lambda = 3.9$ µm, $\tau_{\text{FWHM}} = 118$ fs,

$E_L = 23$ mJ, is also a significant high-performing outlier, providing a valuable sanity check towards the fidelity of our simulations.

To characterise further, we distinguish between different experimental configurations: self-injected/self-guided, guided, staged, optical injection, ionisation injection and downramp injection. Using a logarithmic x-axis and including only a single representative data point from this work (at $n_e = 2.8 \times 10^{18}$ cm^{-3}), this is plotted in Figure 6.5.

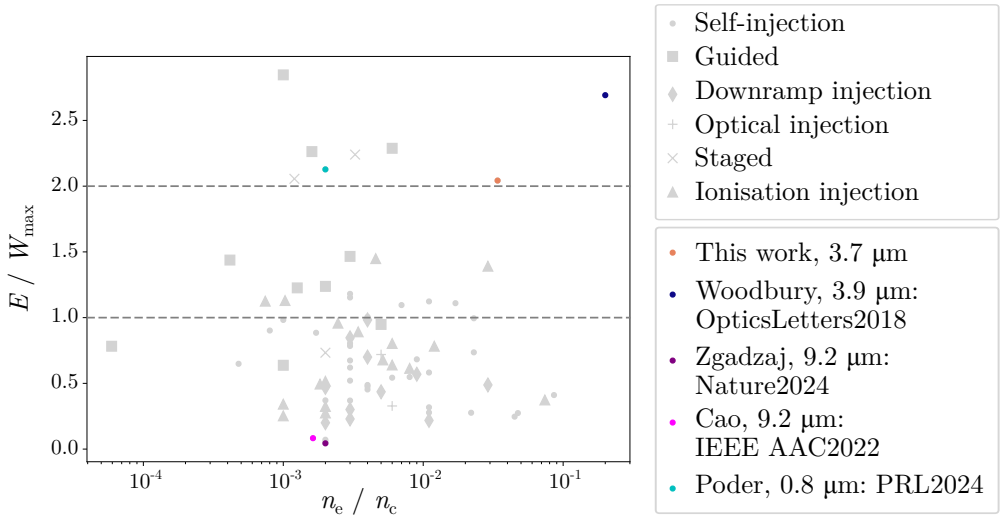


Figure 6.5: Using a logarithmic scale on the x-axis, we replot the data presented in Figure 6.4, with individual results classified into experiments which are self-injecting/self-guided (circles), externally guided (squares), staged (crosses), or using downramp injection (diamonds), optical injection (plus signs) and ionisation injection (triangles). The individually highlighted results are: orange – a single representative point from this work; dark blue – Woodbury *et al.* [185]; Purple – Zgadzaj *et al.* [228]; Cao *et al.* [229]; and Põder *et al.* [225]. The y-axis shows the maximum reported energy reached by accelerated electrons, normalised using the maximum predicted energy W_{max} from Equation (6.6), with a_{0max} calculated from Equation (6.7). Lines at $y = 1$ and $y = 2$ are plotted for reference.

From this plot, it is immediately clear that results lying above the $y = 2$ line originate predominantly from guided or staged experiments. Such cases are more difficult to interpret within the Lu framework: staged setups are not considered, and guided configurations typically operate in a quasi-linear regime, where direct comparison to the Lu scaling (defined in a fundamentally nonlinear regime) becomes less meaningful.

There are three self-injection/self-guided experiments above the $y = 2$ reference in Figure 6.5. The orange circle again marks the representative simulation result from this chapter, while the dark blue circle corresponds to the MIR LWFA experiment of Woodbury *et al.* [185]. We also highlight the pink and purple circles, representing simulated and experimental LWFA results using a 2 J, 4 ps CO_2 driving laser at 9.2 µm [228, 229]. These results are especially striking, as they challenge the intuition that the apparent outliers in our work could simply be explained as a wavelength-driven effect.

The only near-infrared experiment within this 'over-performing' region, which is neither guided nor staged, is a recently reported one from Põder *et al.* [225] in 2024. Here, the authors explored the use of various focusing geometries and showed that initiating the laser with an intentionally over-matched spot size slows the onset of self-focusing. This, in turn, stabilises the bubble length at early times and reduces the beam loss of injected electrons, allowing those injected into the extreme fields at the back of the bubble to remain confined and reach high energies. In our simulations, the laser spot size is close to the matched condition and fixed, with plasma density being the only parameter scanned, suggesting that the mechanism identified by Põder *et al.* is unlikely to explain our results. In addition to this, for these simulations, the injection event containing the electrons which eventually reach these unexpectedly high energies also occurs very early on in the acceleration process. The mechanism described by Põder *et al.* concludes that initiating LWFA using an overlarge spot fundamentally delays injection. Across each of the tested starting densities, we see no such delay.

In this chapter, we will consider various mechanisms which might explain these anomalously high electron energies. These include tracking the evolution of a_0 within the simulation by means of laser redshift and electric field; monitoring the size of the blow out radius as the wakefield evolves; considering carrier envelope phase (CEP) effects; a discussion of whether the high energy electrons might have been accelerated directly by the laser itself, by means of a direct laser acceleration (DLA) scheme; and a discussion into the potentially positive effects of beam-loading.

6.1.1 Normalised Vector Potential

As discussed in Section 3.2.3, self-compression of the laser pulse as it evolves must correspond to some increase in bandwidth, as stated by the time-bandwidth product (Equation (2.3)). This increase in bandwidth materialises as the phenomenon of photon acceleration, resulting from a time-dependent refractive index within the co-propagating plasma. For a Gaussian pulse driving a wakefield through an underdense plasma, the front of the pulse can steepen, resulting in a redshift, whilst the back of the pulse can extend, causing a blueshift [135, 136]. Given the strong wavelength dependence of laser intensity and normalised vector potential, there is a high likelihood of a redshift, if present, contributing to an increase in the a_0.

Figure 6.6 shows a laser kymograph (laser spectrum streaked through time), showing very significant redshifting of the laser at our representative density step, $n_e = 2.8 \times 10^{18}$ cm^{-3}, reaching wavelengths of up to 10 µm. This effect seems to be most prominent at very late times in the simulation, with the biggest shift happening after 2 ps. We can use this to understand if and how this wavelength shift translates to an enhancement in a_0.

As PIC codes are dimensionless, the normalised vector potential a_0 throughout the simulation can be simply estimated by extracting the laser electric field E_z before

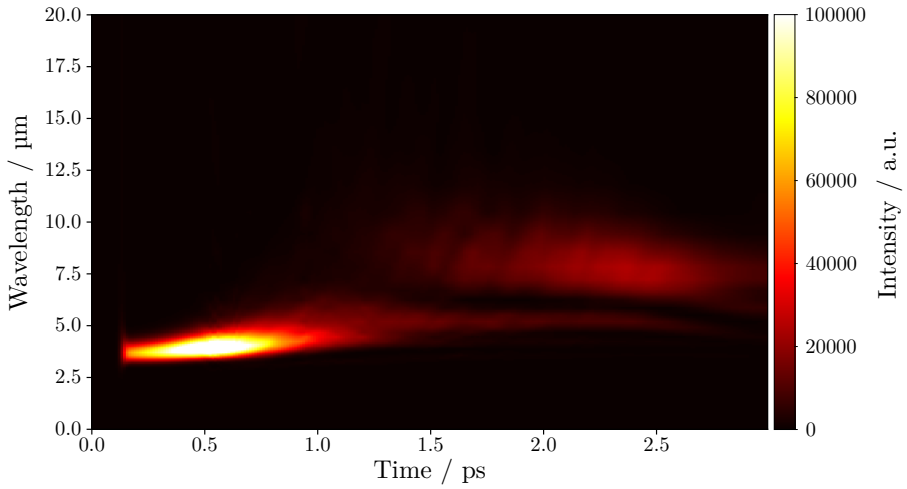


Figure 6.6: A laser kymograph for a simulation with starting density 2.8×10^{18} cm^{-3} and initiated wavelength 3.7 µm. We see significant redshifting of the laser pulse up to wavelengths of 10 µm.

converting the diagnostic into real units. As $a_0 = e|\vec{E_0}|/(m_e \omega_L c)$, we see that for a reference, normalised electric field $\hat{E}$, we can take this as a good approximation to the laser a_0.

We should also consider using alternative methods to calculate a_0. While using the normalised electric field can be effective, it relies on the assumption of a fixed reference frequency. As a result, any a_0 enhancement due to the redshifting of the laser spectrum is not captured in this approach. To compensate for this, the normalised field can be rescaled by the instantaneous wavelength shift observed in the laser spectrum, via the ratio $\lambda(t)/\lambda_0$, or equivalently $\omega_0/\omega(t)$. This correction enables a more accurate estimate of the effective a_0 in the presence of significant spectral evolution. Using this definition, we approximate a_0 as:

$$a_0 = \frac{\lambda(t)}{\lambda_0}\hat{E} \tag{6.8}$$

Estimating a_0 from the normalised electric field provides a useful first-order understanding of its evolution; however, a more rigorous definition can be obtained by returning to the fundamental relation between the electric field and the vector potential. In the Coulomb gauge ($\nabla \cdot \mathbf{A} = 0$) and assuming a negligible scalar potential (as is typical for a laser field in underdense plasma), Maxwell's equations give:

$$\mathbf{E}(x,t) = -\frac{\partial \mathbf{A}(x,t)}{\partial t} \tag{6.9}$$

From this, the vector potential can be recovered via:

$$\mathbf{A}(x,t) = -\int \mathrm{d}t\ \mathbf{E}(x,t) \tag{6.10}$$

The normalised vector potential is then defined as:

$$\mathbf{a} = \frac{\mathbf{A}e}{m_\mathrm{e}c} \tag{6.11}$$

$$a_0 = \frac{A_0 e}{m_\mathrm{e}c} \tag{6.12}$$

The challenge with this method from a computational standpoint is that the field diagnostics available in the Smilei PIC code do not provide a full time-resolved electric field at each spatial position, preventing a direct time integration. In an underdense plasma, the laser pulse propagates at the group velocity v_g. Using this assumption, we may perform a change of variables, converting the time integral into a spatial one and enabling a more physically complete estimate of a_0:

$$\mathbf{A}(x,t) = -\int \frac{\mathrm{d}x}{v_\mathrm{g}}\ \mathbf{E}(x,t) \tag{6.13}$$

For the representative density step, $n_\mathrm{e} = 2.8 \times 10^{18}$ cm^{-3}, Figure 6.7 shows the evolution of all the discussed methods of calculating a_0. Seen in the enhancement of a_0, we see the evidence of strong self-focusing effects at early times up to 0.55 ps – the point at which we see the laser pulse evolving to its tightest focal spot. We also see a strong a_0 enhancement at times corresponding to the redshift at later times, as seen in Figure 6.6.

Although both self-focusing and photon acceleration arise from nonlinear laser-plasma interactions, we highlight that they are different in origin and effect. Self-focusing is a transverse effect driven by the intensity dependence of the plasma refractive index and occurs immediately as the laser enters the plasma. This leads to rapid temporal compression of the pulse and an early increase in a_0. In contrast, photon acceleration is a longitudinal effect caused by the time-dependent refractive index introduced by the evolving plasma wave. In our simulation, significant redshifting (characteristic of photon deceleration) only becomes apparent at later times, particularly after 2 ps. While this redshifting does result in a secondary enhancement of a_0, it does not exceed the peak value established during the initial self-focusing phase. Consistent with this, the highest energy electrons in the simulation are injected early, coinciding with the period of strongest self-focusing. This suggests that the initial self-focusing phase (rather than the later spectral evolution due to photon acceleration) plays the dominant role in setting the energy gain of accelerated electrons. In both cases of a_0 enhancement, the maximum

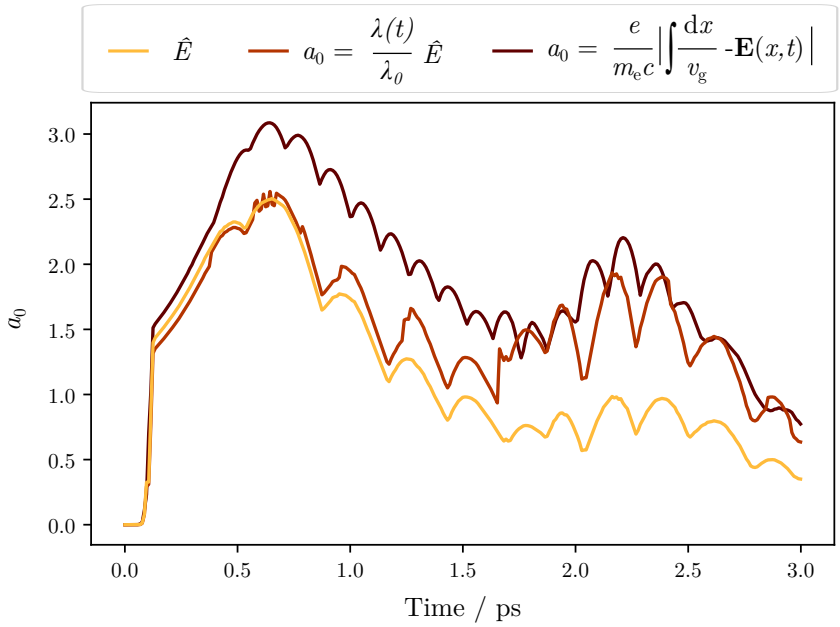


Figure 6.7: Evolution of the normalised vector potential a_0 through a simulation with starting parameters: $a_0 = 1.9$, $n_e = 2.8 \times 10^{18}$ cm^{-3}. The a_0 retrieval is done using three different methods: extracting the normalised laser field, the normalised laser field rescaled by the instantaneous laser wavelength shift, and the vector potential derived using Maxwell's equations.

a_0 achieved in the simulation remains below the threshold required to match the electron energies predicted by the scalings of Lu *et al.*

We can also refer back to the model proposed by Thomas (2010) [143], discussed in Section 3.3.2, which estimates the threshold a_0 required for electron injection in the bubble regime. According to this model, injection occurs when:

$$a_0 \gtrsim \ln(2\gamma_p^2) - 1 \tag{6.14}$$

Where γ_p is the Lorentz factor corresponding to the phase velocity of the plasma wave, or equivalently, the velocity of the back of the bubble in the etching regime. The model assumes $\gamma_p^2 = n_c/3n_e$. Using this expression, we find that the threshold for injection corresponds to $a_0 = 1.96$. Although we do not use this model to precisely determine the value of a_0 at the specific injection point in our simulation, given that the highest energy electrons are injected so early in time, these initial self-injected electrons likely experience a laser a_0 close to this threshold. If, as would be suggested by all methods of tracking a_0 presented here, this is the case, it remains unclear why electrons accelerated after being injected at this timestep reach energies up to twice those predicted by established models.

6.1.2 Carrier Envelope Phase (CEP)

When faced with discrepancies in irregular results, it is essential to re-examine the operating regime of the simulation. With the starting laser parameters used in these simulations, we are operating well within the few-cycle regime, at 2.3 optical cycles. As highlighted in Section 3.1.1, such short pulses can give rise to behaviour that deviates from standard LWFA expectations. This is because many of the theoretical models and diagnostic assumptions applied to both experimental and numerical studies of LWFA are built upon the slowly varying envelope approximation (SVEA). The SVEA underpins the derivation of the ponderomotive force, which is central to electron acceleration in the wakefield regime (see Section 3.1.1). When the pulse duration approaches only a few optical cycles, the SVEA can no longer be applied, and the associated models may fail to capture all aspects of the laser–plasma interaction.

One approach to investigating and mitigating potential few-cycle effects is to examine the influence of the carrier envelope phase (CEP) of the driving laser pulse. CEP effects in LWFA, and few-cycle behaviour more generally, remain active areas of research, with many groups worldwide working to identify their experimental signatures and explore ways to exploit them for performance gains [**66**, **127**, **230**]. Although no comprehensive theoretical model yet exists to predict electron dynamics in the presence of CEP variation, the effect can be straightforwardly explored in simulation by initiating pulses with different carrier envelope phases. In near-single-cycle pulses, changing the CEP modifies both the amplitude and phase of the laser field within the pulse envelope. If CEP plays a significant role in the dynamics, this should manifest as variations in the resulting electron spectra. In particular, assuming that trapped and accelerated electrons originate from the most intense part of the laser cycle, the CEP will control the exact initial conditions of their injection and, consequently, their subsequent acceleration [**231**].

Simulations were performed at a starting plasma density of $n_e = 2.8 \times 10^{18}$ cm^{-3}, with initial laser parameters $a_0 = 1.9$, $\lambda_0 = 3.7$ µm, $E_L = 25$ mJ, and $\tau_{FWHM} = 30$ fs. To investigate the influence of the CEP on the resulting electron characteristics, we can adjust the phase term ϕ within the simulated laser carrier envelope. Here we compare simulations run with: $\phi = 0$, $\pi/6$, $\pi/4$, $\pi/3$, $\pi/2$ and π.

Examining the resultant electron spectra in Figure 6.8, no significant differences are observed between the simulations. In all 6 cases, the spectra show initial injection occurring at 0.560(4) ps, with 90th percentile energies of 43.5(4) MeV/c, 1.7(10) ps into the simulations, where the uncertainties represent the standard deviation between the tested simulations. The consistency of the final electron energies across all CEP values suggests that varying the laser CEP has no measurable impact on the acceleration process in this regime.

While this may initially seem unlikely given that we are simulating in a regime of only 2.3 optical cycles, this becomes clearer when we examine the actual pulse field being simulated. Figure 6.9 shows the electric field for CEP values of $\phi = 0$ and $\phi = \pi/2$

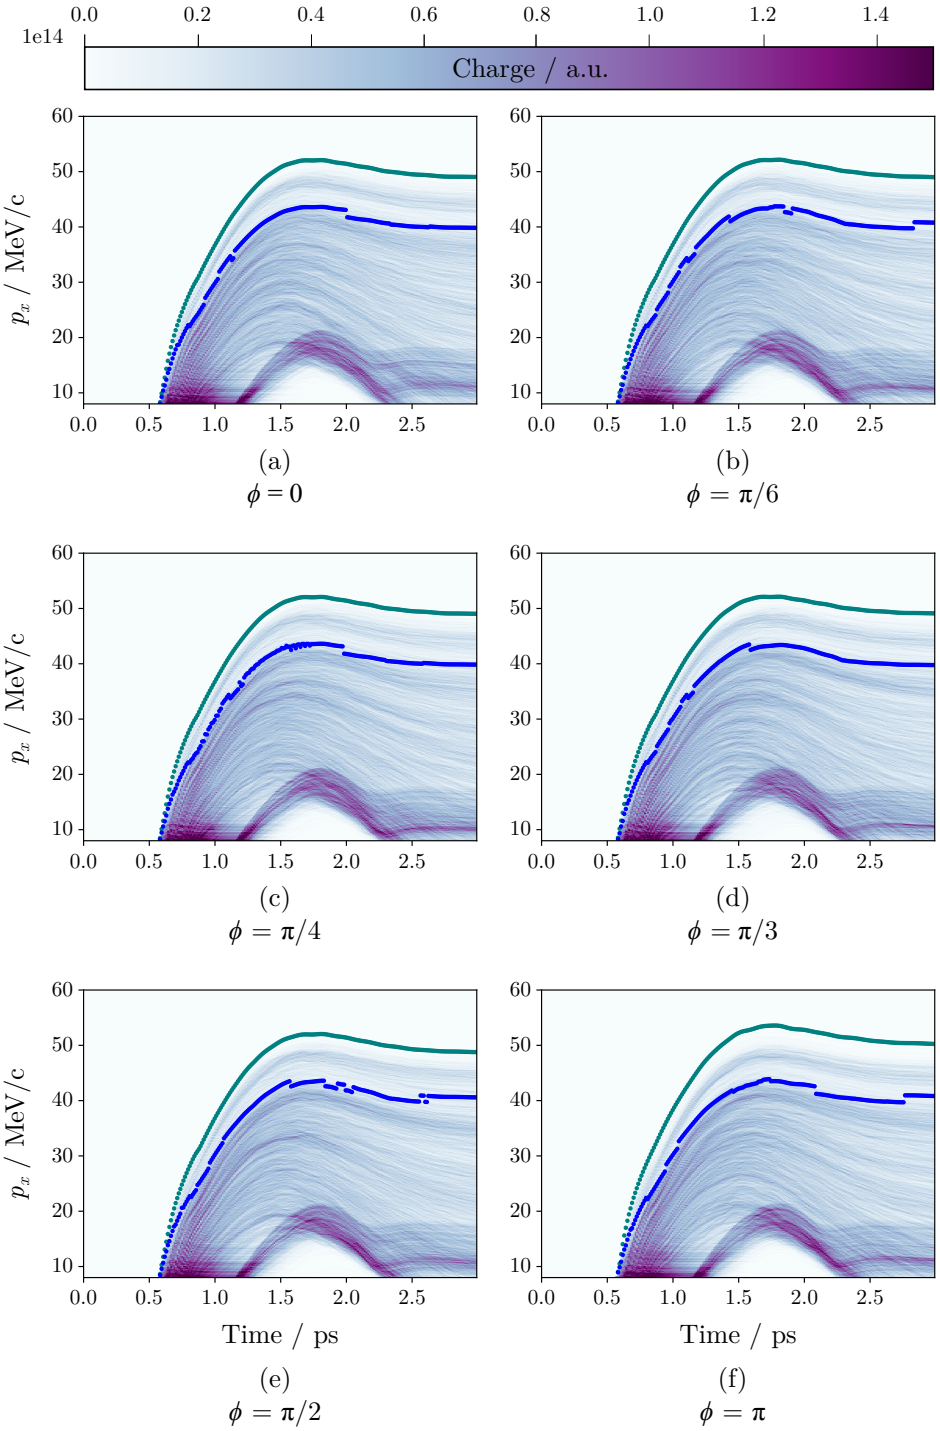


Figure 6.8: Electron waterfall energy spectra for LWFA simulations run while scanning the laser carrier envelope phase (CEP) term. Here we test: (a): $\phi = 0$, (b): $\pi/6$, (c): $\pi/4$, (d): $\pi/3$, (e): $\pi/2$ and (f): π. The maximum (green) and 90^{th} percentile (blue) electron momenta p_x of distributed charge are extracted and overlaid onto each plot.

(commonly referred to as cosine and sine CEP). Not only does the imposed phase offset appear relatively minor, but the overall pulse envelope is also largely unchanged. For CEP effects to be prominent, we would expect a substantial difference between the highest peak within the pulse envelope and its second-highest peak. Somewhat surprisingly, despite nominally containing just 2.3 optical cycles, the pulse does not appear to exhibit the strong asymmetry typical of a few-cycle waveform. This can be understood by recalling that the 'number of optical cycles' is defined relative to the pulse's full width at half maximum (FWHM) duration, rather than the total field duration.

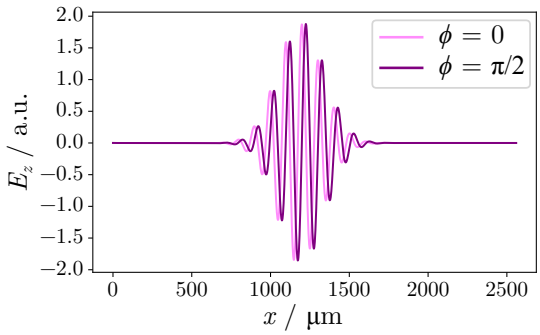


Figure 6.9: The laser electric field E_z for both cosine and sine carrier envelope phases ($\phi = 0$, $\phi = \pi/2$), 0.27 ps into the simulation. We see the offset driven by the change in carrier phase; however, there is no particular asymmetry caused by this. There is no substantial difference between the first and second highest peaks within the pulse envelope, implying that modifications in carrier envelope phase are not a significant concern for the simulations we present.

Given that changing the laser CEP produces only small changes to the waveform, it is, upon reflection, inevitable that we observe little variation in the electron spectra presented in Figure 6.8. Any genuine CEP-dependent variation is likely to be overshadowed by strong pulse evolution during propagation through the dispersive plasma, meaning that the residual differences will be predominantly statistical.

6.1.3 Direct Laser Acceleration (DLA)

If the electrons under consideration are not gaining energy via the ponderomotive force in the classical wakefield sense, they may instead be accelerated through a different mechanism. In contrast to LWFA, where energy gain arises primarily from the longitudinal electric fields of the plasma wake driven by the laser's ponderomotive potential, direct laser acceleration (DLA) involves a direct transfer of energy from the laser to the electrons themselves. As discussed in Section 3.4, DLA occurs when electrons undergo significant resonant transverse oscillations within the laser field, enabling a coherent transfer of energy from the laser to the particles. This could be an explanation for the significant mismatch of observed electron energies with theory, as the models used here do not consider any contribution from DLA.

DLA has a characteristic signature in electron momentum space: it exhibits a forked structure in its $p_\perp p_\parallel$ phase space. For classical LWFA, such a plot will typically show a single narrow streak along the longitudinal axis, indicating that most momentum is

in the direction of forward-propagation with only a small transverse spread. For DLA-accelerated electrons, however, two symmetric diagonal 'prongs' that split away from the main longitudinal spine will appear, arising from resonant betatron oscillations in the laser field. Betatron oscillations are a normal consequence of LWFA, but if the frequency of these oscillations (or a harmonic) matches the Doppler-shifted laser frequency in the electron's frame, the laser can do net work on the electron in the transverse direction. The resulting transverse momentum is then partially converted into additional forward momentum via the $\mathbf{v} \times \mathbf{B}$ force, producing the characteristic angled prongs.

If present, the fork's features can inform interpretation: the prongs' length and population relative to the LWFA spine provide a measure of the DLA contribution, while their angle directly indicates the ratio of transverse to longitudinal momentum – steeper angles corresponding to stronger transverse oscillations relative to forward acceleration. Shown in Figure 6.10, no such DLA signature is observed in our simulation.

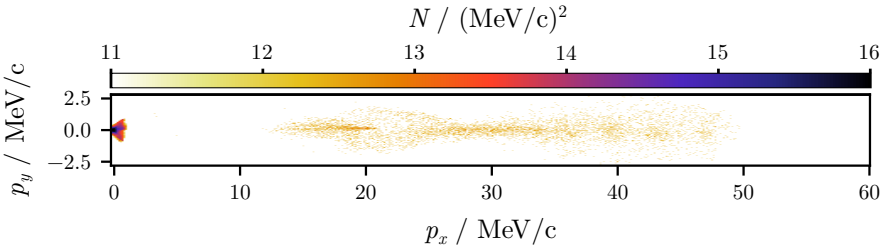


Figure 6.10: A map over the electron momentum phase space for a simulation with starting density 2.8×10^{18} cm^{-3}. The snapshot is taken at $t = 1.69$ ps – the time at which the electrons reach their highest energy. The spectrum shows only a small amount of transverse spread – it has a single narrow streak along the longitudinal axis, characteristic of a classical LWFA electron beam. There is no evidence of any 'forked' or 'pronged' structure which would be characteristic of DLA.

To confirm this result without relying solely on visual inspection, we can directly track the momenta of a representative set of electrons and quantify the mechanism by which they gain energy. For electrons accelerated via DLA, energy gain arises from work done by the transverse laser electric field [152]:

$$W_\perp = -e \int_0^t \mathbf{E}_\perp \cdot \mathbf{v}_\perp \mathrm{d}t' \tag{6.15}$$

In contrast, electrons accelerated in the LWFA regime gain energy from work done by the longitudinal electric field of the plasma wake:

$$W_\parallel = -e \int_0^t \mathbf{E}_\parallel \cdot \mathbf{v}_\parallel \mathrm{d}t' \tag{6.16}$$

Where $\mathbf{E}_\perp$, $\mathbf{E}_\parallel$ are the transverse and longitudinal electric field components, $\mathbf{v}_\perp$ and $\mathbf{v}_\parallel$ are the corresponding velocity components of the electrons, and e is the electron charge.

For all simulations in the density scan (starting plasma densities of $n_e = 1.8 - 5.0 \times 10^{18}$ cm^{-3}), all with initial laser parameters $a_0 = 1.9$, $\lambda_0 = 3.7$ µm, $E_L = 25$ mJ, and $\tau_{FWHM} = 30$ fs, we plot the DLA and LWFA contributions to the final energy of each of the randomly selected electrons versus their final energies. This is shown in Figure 6.11. We note that, due to the electron selection criteria (tracking a percentage of electrons that exceed 5 MeV at any point), the number of electrons tracked varies with density.

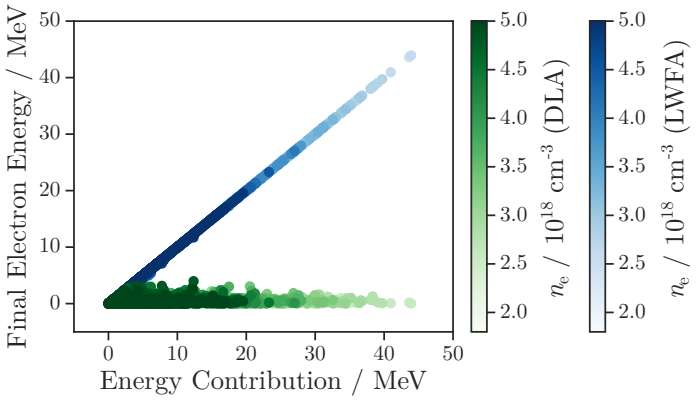


Figure 6.11: The DLA $W_\perp$ (green) and LWFA $W_\parallel$ (blue) contributions to the final energy of randomly selected tracked electrons versus their final energies for starting densities $n_e = 1.8 - 5.0 \times 10^{18}$ cm^{-3}.

This plot demonstrates that the maximum transverse energy contribution to the final electron energy (i.e., the DLA contribution) never exceeds 5 MeV. This is consistent with our interpretation of the momentum phase space in Figure 6.10: the DLA mechanism does not significantly contribute to the energies achieved by the accelerated electrons.

6.1.4 Bubble Size

Amongst the scaling laws prescribed by Lu *et al.* [74] is the prediction that, in the so-called bubble or blow-out regime of nonlinear laser wakefield acceleration, the radius of the plasma cavity is directly related to the normalised vector potential a_0 of the driving laser field. Building on earlier phenomenological work by Gordienko and Pukhov [232], and verified using particle-in-cell simulations, they find that the bubble radius R is given by:

$$k_p R = 2\sqrt{a_0} \tag{6.17}$$

Where k_p is the plasma wavevector. Although this relation is most often used to determine a matched spot size that favours self-guiding, it can also be inverted: by

measuring R in a simulation, we can estimate the corresponding a_0 at any point during the interaction. This provides an independent method for tracking a_0 evolution throughout the simulation. The procedure for obtaining a quantitative bubble radius is illustrated in Figure 6.12.

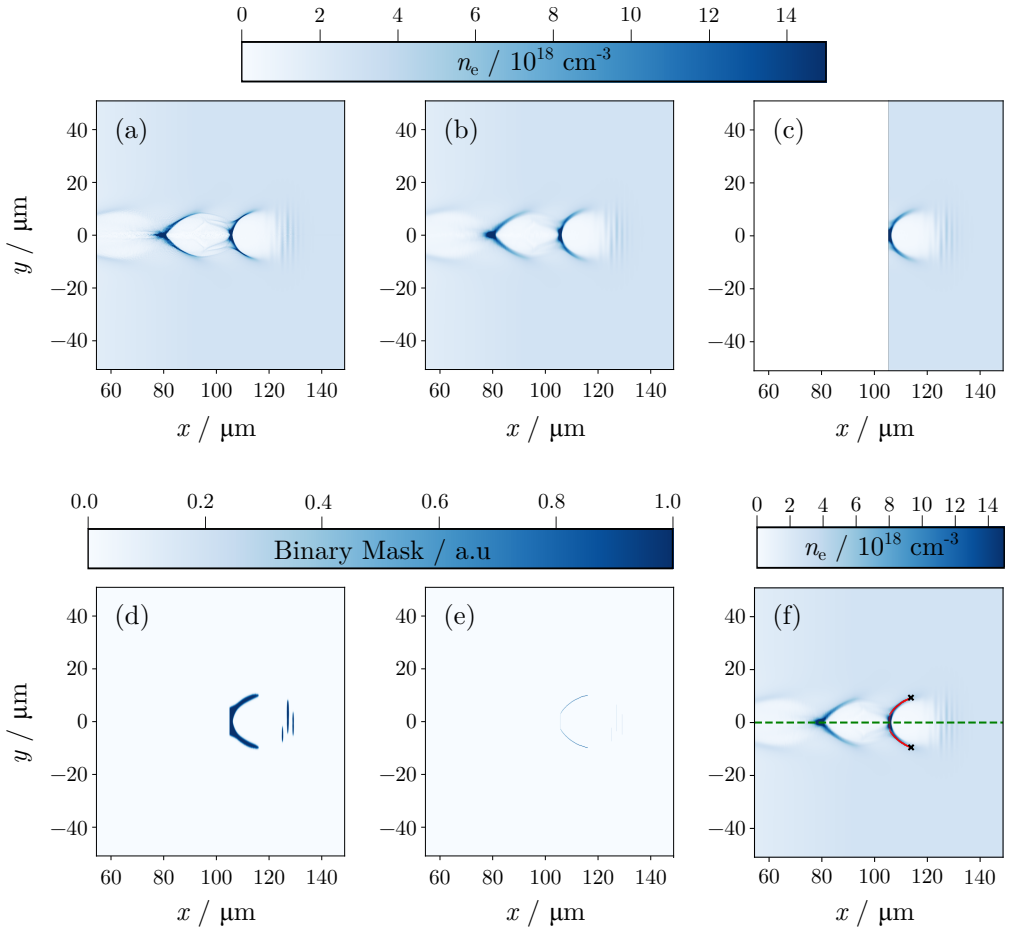


Figure 6.12: Procedure for quantitatively extracting the plasma bubble radius from a PIC simulation.
(a) Raw electron density map at the timestep of interest (here we are at 0.55 ps).
(b) Gaussian smoothing followed by a 10×10 flat average to suppress small-scale noise.
(c) Spatial threshold applied along the propagation axis to isolate the primary bubble.
(d) Density thresholding to produce a binary map of high-density regions, followed by Gaussian smoothing.
(e) Skeletonisation of the bubble boundary.
(f) Extraction of a smooth curve representing the bubble profile, from which the radius is measured. The black crosses identify the two endpoints of the bucket feature, relative to its central axis in green. This is overlaid onto the electron density map.

We apply this method at the timestep 0.55 ps – the point at which the electrons that reach the highest energies are injected. Starting with a simple electron density map taken at the timestep of interest (a), we firstly use a Gaussian filter to smooth the data by averaging nearby values before smoothing again using a flat average over a 10×10 neighbourhood (b). To isolate the bubble structure, we exclude all data before a specified x-position by applying a threshold along the propagation axis (c). We apply an

intensity threshold to the remaining data to create a binary map of regions above the chosen density, which is smoothed again using a Gaussian filter (d), before skeletonising the data (e), such that a single smooth curve can be extracted (f).

This method allows us to extract a transverse radius of the bubble, measured between the two endpoints of the bucket feature relative to its central axis. The problem that follows is that, rather than being a perfect sphere, the bubble structure in our case is actually an ellipse. As a result, we can also define a longitudinal radius: the distance from the central point along the axis, set by the transverse radius, to the back of the bubble, as determined from the extracted curve. For this simulation, the transverse and longitudinal radii are calculated to be 9.95 µm and 13.26 µm, respectively. The visual result of this fitting is shown in Figure 6.13.

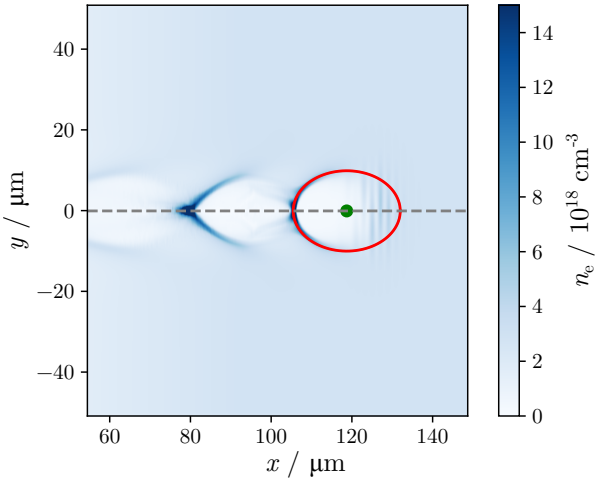


Figure 6.13: Electron density map, with starting plasma density $n_e = 2.8 \times 10^{18}$ cm^{-3} at 0.55 ps. The density map is overlaid with the calculated bubble centre (green dot), the central horizontal bubble axis (grey dashed) and the ellipse (red) inferred from tracking the bubble shape.

These two radii lead to markedly different values of a_0: using the transverse radius of $R = 9.95$ µm yields $a_0 = 2.46$, while the longitudinal radius of $R = 13.26$ µm gives $a_0 = 4.35$. The former of these values is consistent with our previous methods of tracking a_0, whereas the latter – if accurate – could potentially explain the high electron energies we observe.

There are several difficulties in pursuing this interpretation: firstly, the Lu model described by Equation (6.17) assumes a spherically symmetric bubble – something that this method clearly does not yield; secondly, defining the exact centre or edge of the bubble is non-trivial. In particular, the longitudinal measurement is highly sensitive to small changes in where the 'end' of the bubble is defined, whereas the transverse measurement is far less affected. For example, shifting the central position in Figure 6.13 (green circle) from 118.7 µm to 116.8 µm so that the far edge of the ellipse does not extend as far into the laser field reduces the longitudinally derived a_0 from 4.35 to 3.17, while the transverse a_0 changes only slightly, from 2.46 to 2.19. We further note that in

the matched regime, the laser spot size w_0 should be roughly equal to the bubble radius. For a FWHM spot size at focus of d_{FWHM} 10.13 µm, this would translate to a laser spot size w_0 of 8.58 µm and $a_0 \approx 1.82$.

We can look to perform this same exercise 1.69 ps into the simulation – the point at which the electrons are at their highest energies; however, we find that this quickly becomes nonsensical. Shown in Figure 6.14, there is no longer a well-defined 'bubble' to track, and the wake appears to be heavily beam-driven.

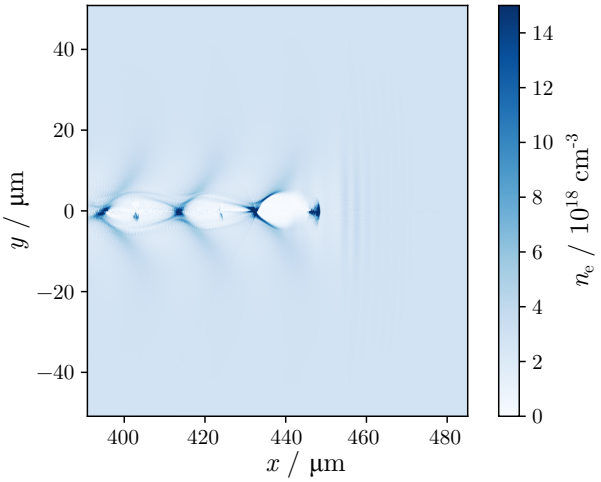


Figure 6.14: Electron density map, with starting plasma density $n_e = 2.8 \times 10^{18}$ cm^{-3} at 1.69 ps. The bubble no longer retains its approximately spherical shape – it has elongated in the longitudinal direction and contracted in the transverse direction. At this timestep, it appears to be heavily beam-driven – a high-density bunch is driving a blown-out electron cavity in its wake.

Such significant beam loading – to the point where the electron bunch is driving its own wake before the depletion of the laser pulse – gives rise to the suggestion that the high-energy electrons that we see in these simulations are a result of beam-driven effects.

6.1.5 Beam Loading

The concept of beam loading is not a new phenomenon [233]. Historically seen as an effect which will ultimately limit the charge that can be accelerated in a LWFA, as well as the quality of the accelerated beam, beam loading is often posed as a serious disadvantage to the prospects of laser-plasma accelerators. While pure beam driven plasma wakefield accelerators (PWFA), where the wakefield is driven by a very short particle beam rather than a laser [234–236], have the significant advantage of not being subject to the same dephasing limit as the laser-driven case, the combination of laser and plasma drivers (by way of beam loading), is not seen to be an experimental advantage. For the most part, investigations have been undertaken to control and characterise the effect, as, when uncontrolled, it can contribute to an extensive energy spread in the output electron beams [237–241]. On the other hand, there have been some proposals regarding the feasibility of such a hybrid scheme, to profit from the benefits of PWFA,

without the challenge of accessing a classical particle accelerator to provide the short electron pulse [242]. This so-called 'self-mode-transition' from LWFA to PWFA of a laser-driven plasma-based electron accelerator has been reported to contribute to electron energy gain in laser wakefield accelerators – specifically, it has previously been reported to enhance electron energies up to more than two times the Lu scaling predictions [243, 244].

We thus propose that it is the result of extreme beam loading, to the extent of a self-mode-transition to a hybrid LWFA/PWFA scheme, that causes the higher-than-predicted electron energies seen in these simulations. Here, we use the hybrid LWFA/PWFA term to mean that a particle bunch begins to drive its own wake, whilst still benefiting from the accelerating fields of the driving laser. This is different to the cases reported in [242–244], who describe a particle bunch driving the acceleration of a second bunch once the driving laser is fully depleted.

We first look to confirm that we really are within the regime of a beam-driven wakefield. Shown at the point of injection, as well as several timesteps following up to the point of maximum energy, Figure 6.15 shows the wakefield structure in density space. To quantify whether or not a snapshot can be classed as being beam driven, we overlay a 1D line-out of particle density: it is a consequence of a laser-driven wake that it can never be considered fully blown-out – that is to say, the particle density will never drop to zero to form a complete ion channel. This is not the case for a beam-driven wakefield, meaning that if, at any point, the particle density following an electron bunch is equal to zero, we may judge the wake that follows it to be beam-driven. In addition, we include a 1D line-out of the longitudinal electric field E_x. In a beam-loaded or beam-driven wake, the presence of a high-charge electron bunch modifies the wakefield, leading to a characteristic flattening of E_x across the bunch. The simultaneous observation of a fully evacuated ion channel and a flattened accelerating field therefore provides strong evidence that the wake is being driven, at least in part, by the electron beam rather than the laser alone.

This distinction arises from the different forces responsible for driving the wakefield, and is reflected both in the plasma density structure and in the longitudinal electric field. In LWFA, the ponderomotive force (Equation (3.17)) is proportional to the gradient of an oscillating, spatially inhomogeneous laser field, making it less effective at fully evacuating plasma electrons and sustaining a flat accelerating field than the peak field amplitude alone might suggest. In PWFA, by contrast, the dominant force arises from space-charge separation within the electron beam itself, which can much more efficiently expel plasma electrons and directly shape the accelerating field. As a result, even though the total energy carried by an electron bunch is often considerably smaller than that of a laser, it can excite a plasma wave more effectively and produce a more strongly beam-loaded wake [245].

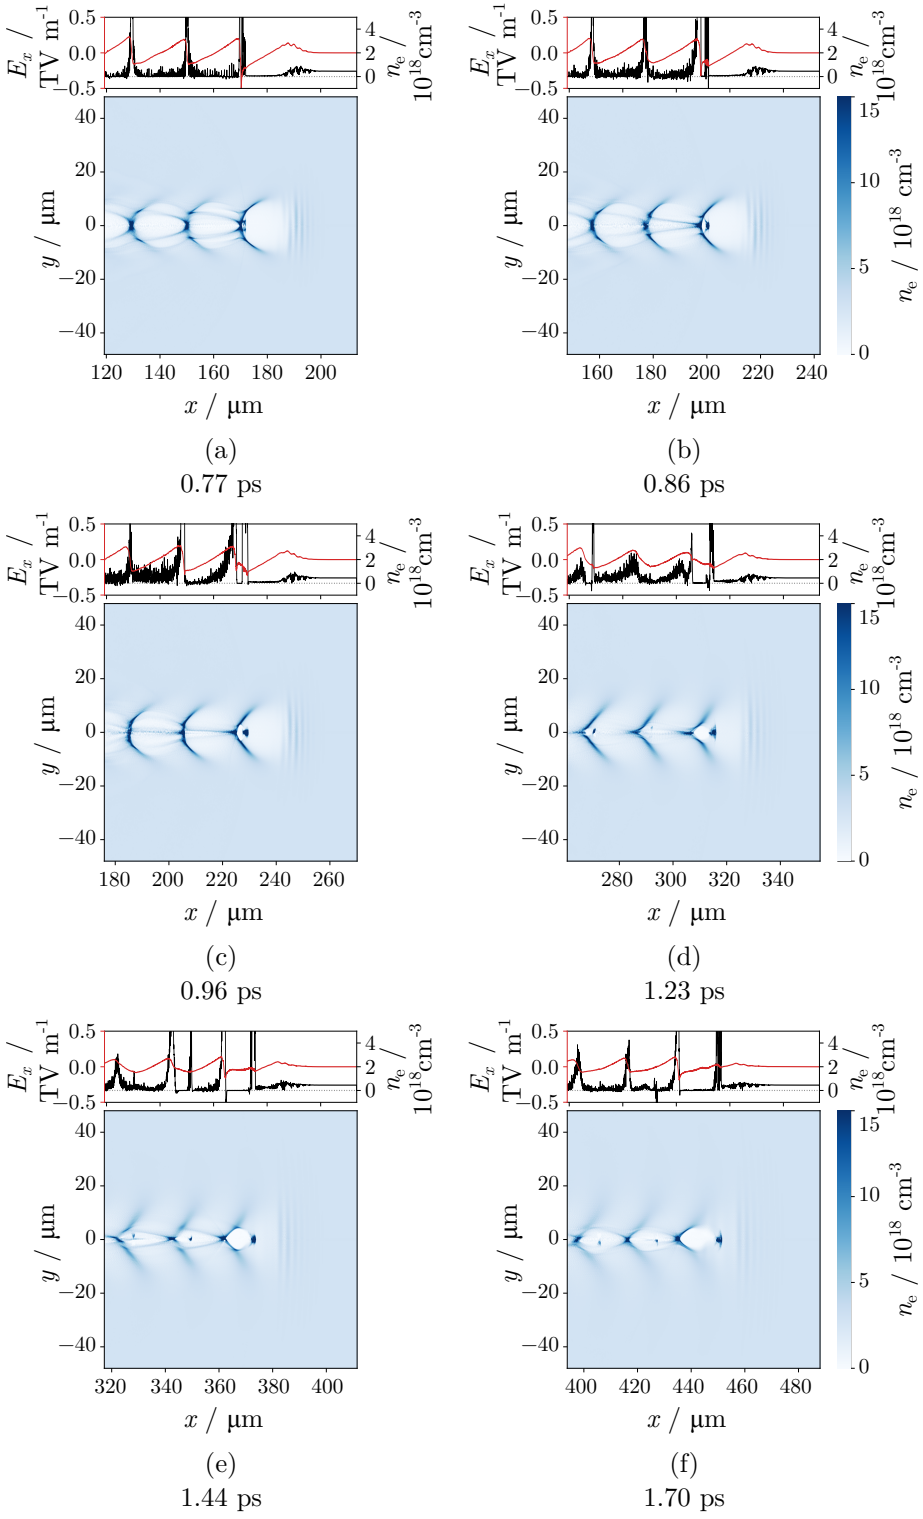


Figure 6.15: Electron density map with starting plasma density $n_e = 2.8 \times 10^{18}$ cm^{-3} with additional 1D density (black) and longitudinal field E_x (red) line-outs, showing beam-loading effects during the course of the simulation. Snapshots are shown at (a) 0.77 ps, (b) 0.86 ps, (c) 0.96 ps, (d) 1.23 ps, (e) 1.44 ps and (f) 1.70 ps. The intensity scale is shared for each row in the panel. With the exception of panel (a), the plasma density line-out drops to zero immediately behind the injected bunch (reference grey dotted line included at $n_e = 0$), indicating a fully blown-out, beam-driven wake. The corresponding flattening of E_x further supports the presence of beam loading.

The results in Figure 6.15 confirm that beam-driven effects are indeed present, and remarkably, they appear as early as 0.86 ps into the simulation (panel (b)). Previous studies have typically reported the onset of this hybrid mode only after the depletion of the laser in a long plasma target [243, 246], or when the laser diffracts due to ionisation-induced defocusing [247], leaving the electron bunch as the dominant driver [245].

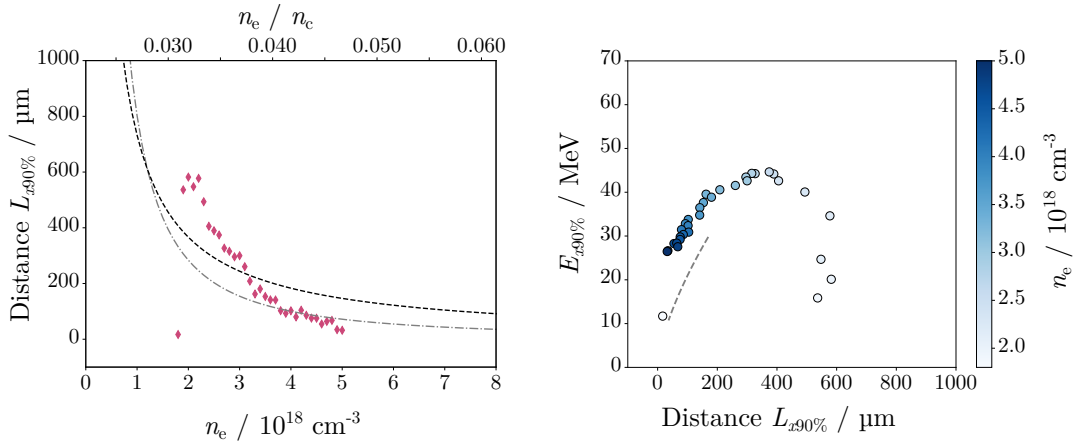
Although our simulations clearly show the presence of beam loading effects, the precise mechanism by which they enhance energy gain is less straightforward. The hypothesis that we propose is that this hybrid regime acts to slow down dephasing. Through particle tracking, we can determine that the electrons which reach these unexpectedly high energies are being injected very early on in the process, before the wake becomes beam-driven. We find that electrons are 'born' ~ 0.55 ps into the simulation. Since beam-driven effects appear almost immediately after injection, the resulting slowdown of dephasing would allow these early-injected electrons to remain in phase with the accelerating laser field for longer, and thus gain more energy.

If this is the case, we anticipate that we should exceed the dephasing scaling length described in Equation (3.48). Considering now all simulations in the density scan, with $n_\mathrm{e} = 1.8 - 5.0 \times 10^{18}$ cm^{-3}, we show in Figure 6.16 how this compares to the scalings. Figure 6.16a shows directly the distances reached by the 90$^{\mathrm{th}}$ percentile of accelerated electrons plotted alongside both the dephasing and pump depletion lengths (Equations (3.48) and (3.49)). Figure 6.16b presents the same data as a function of energy. To make this comparison meaningful, we do not use the absolute electron position in the lab frame. Instead, we define the effective acceleration length, $L_{x90\%}$, as the distance over which the electrons have been accelerated within the wake. This is obtained by subtracting the distance travelled by the laser prior to injection, thereby isolating the true acceleration length. To express the predicted energy gain as a function of acceleration length, we combine Equation (3.48) and Equation (3.51) as:

$$E = \frac{\pi\sqrt{a_0}}{\lambda_\mathrm{p}} L_\mathrm{d} \tag{6.18}$$

In Figure 6.16, we see that at lower plasma densities the accelerated electrons travel well beyond the predicted dephasing length. At higher densities, this is not the case, suggesting that while an extended dephasing length may be a contributing factor, it cannot be used as definitive evidence towards hybrid LWFA/PWFA playing a role in the observed overperformance of the accelerator. This is because even at high densities – where the dephasing length is not exceeded – the electrons still reach energies beyond the Lu predictions.

This highlights an important limitation of using dephasing alone as a diagnostic for the underlying acceleration mechanism. In practice, the maximum achievable energy is not controlled by dephasing length in isolation, but by a combination of factors including



(a) (b)

Figure 6.16: Comparison of simulated acceleration lengths with theoretical dephasing and pump-depletion scalings.
(a) Effective acceleration length $L_{x90\%}$ reached by the 90$^{\text{th}}$ percentile of electrons as a function of plasma density n_e, plotted alongside the predicted dephasing (grey) and pump-depletion (black) lengths.
(b) 90$^{\text{th}}$ percentile electron energy plotted as a function of acceleration length $L_{x90\%}$. The dephasing length, calculated using Equation (6.18), is included as the grey curve.
In both panels, $L_{x90\%}$ accounts only for the relative distance travelled after injection, subtracting the driver's motion to isolate the effective acceleration path.

the evolution of the wake phase velocity, pump depletion, beam loading, and the stability of the accelerating structure. In particular, strong beam loading can simultaneously flatten the accelerating field, reducing local energy gain while also modifying the wake dynamics in a way that delays dephasing for early-injected electrons.

There are also other reasons why an accelerator might appear to exceed its predicted dephasing length. One possibility is that the etching of the laser front, defined by Decker *et al.* [149], proceeds more slowly than expected. Since the etching velocity is given approximately by $\beta_{\text{etch}} = n_e/n_c$, a reduced etching rate would increase the effective phase velocity of the wake, thereby extending both the dephasing and pump depletion lengths beyond the simple scaling predictions.

More work is required to understand this result in its entirety; however, if it is the case that high-energy electrons can be accelerated through the proposed hybrid LWFA/PWFA mechanism, it would be of great interest to pursue this further, as it could provide a route to enabling higher energy gains without requiring larger, more powerful laser systems. In particular, it would be of interest to see if this effect can be further manipulated, especially when using alternative injection schemes such as downramp, ionisation or shock injection, which offer greater control over the injected beam's energy and energy spread. Successfully harnessing this hybrid regime could offer a means of producing higher-quality, higher-energy electron beams for both scientific and industrial applications.

6.2 Discussion

The possibility of a hybrid LWFA/PWFA mechanism provides a potential explanation for the higher-than-expected energies seen in our simulation campaign; however, to understand why these simulations may access this regime, it is necessary to place these results within a wider experimental and theoretical context.

Considering the normalised parameters – number of optical cycles, a_0, and P/P_c – which we previously suggested as directly comparable across the full LWFA parameter space, we investigate whether any systematic trends emerge that might help explain why both the MIR-driven simulations presented here perform unusually well compared to the rest of the literature. In addition to these parameters, we introduce a fourth: the ratio of the pulse length $c\tau$ to the plasma wavelength λ_p. Using the same dataset as Figure 6.5, we apply a colour map to visualise these normalised quantities across the full database. This is shown in Figure 6.17.

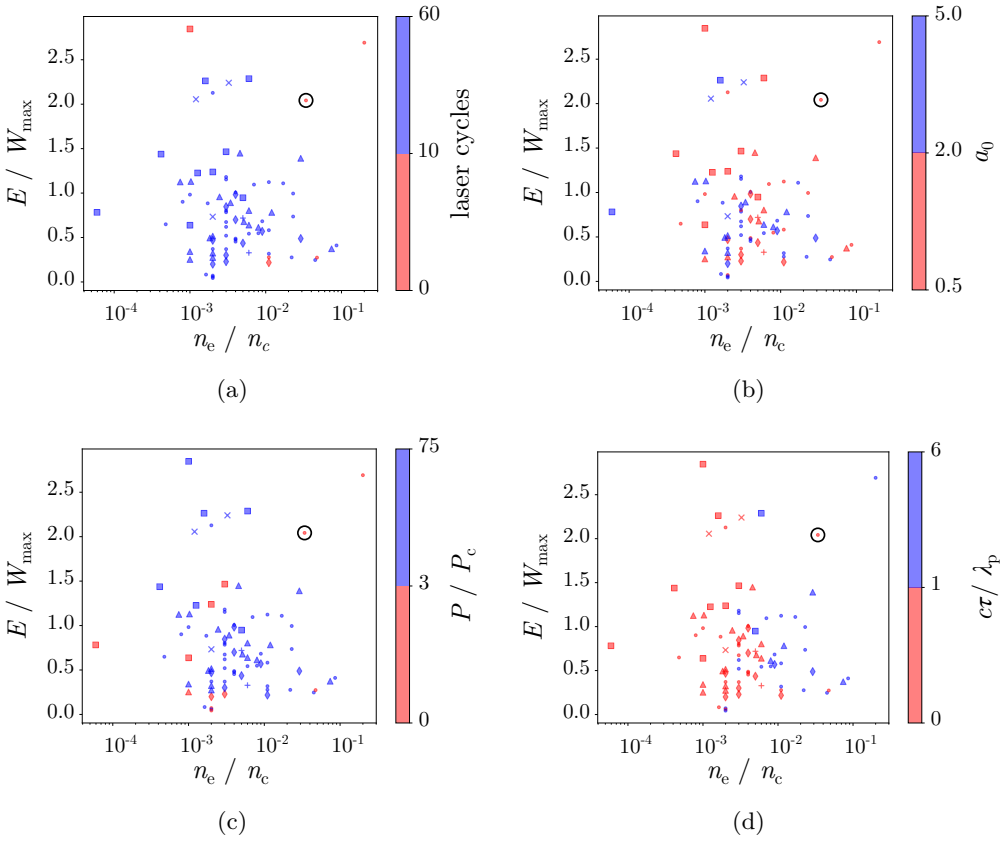


Figure 6.17: Using the same data as Figure 6.5, we apply a colour map to explore possible trends in normalised (and therefore directly comparable) quantities across the full database. The tested parameters are: (a) number of optical cycles, (b) a_0, (c) P/P_c and (d) $c\tau/\lambda_p$. All plots share the same normalised x- and y-axes. As before, circles denote self-injection experiments; squares, externally guided; diamonds, downramp injection; plus signs, optical injection; crosses, staged experiments; and triangles, ionisation injection. The representative simulation datum presented in this thesis is circled in black.

In the first three cases, the distribution of the tested normalised parameters appears essentially random: no clear trend or correlation emerges. Figure 6.17a examines whether

the number of laser cycles correlates with the maximum electron energy. As only a small number of experiments fall into the 'few-cycle' category (here taken to mean fewer than 10 optical cycles within the FWHM pulse envelope), the available data are insufficient to support a firm conclusion. Nevertheless, several of the lower-energy data points correspond to experiments explicitly designed to probe CEP effects (notably at Laboratoire d'Optique Appliquée, ENSTA/École Polytechnique) [66, 127, 230, 231], and the inclusion of these targeted results reinforces the absence of any discernible trend.

The a_0 and P/P_c ranges shown in Figure 6.17b and Figure 6.17c cover a much wider set of experiments but likewise show no correlation between either a_0 or P/P_c and the achieved electron energies. Given the unusually low P/P_c values arising from our simulation parameters, one might have expected a stronger signature to emerge, but no such pattern is observed.

Figure 6.17d, on the other hand, may allow some interpretation. For efficient wakefield excitation, the laser pulse length should be on the order of the plasma period, meaning that for shorter wavelength laser drivers, it is important to operate in regimes of lower plasma density. Assuming an essentially fixed pulse length $c\tau$ across experiments, the ratio $c\tau/\lambda_p$ naturally grows with increasing plasma density, meaning that the vast majority of experiments operating at high ratios of plasma to critical density n_e/n_c, must use driving lasers with a comparatively long pulse length to plasma wavelength ratio. The choice of laser wavelength thus affects this condition indirectly: because the critical density scales as $n_c \propto 1/\lambda$, longer wavelength drivers naturally operate at higher n_e/n_c. It is likely, therefore, that the parameter space we find ourselves in for these simulations is extremely unusual. Although the complexity of LWFA parameter space, where many quantities are strongly entangled and interdependent, make it difficult to isolate a single explanation for why our simulations appear to over-perform relative to the electron energies predicted by Lu *et al.* [74], it is very likely that it has something to do with the efficient high density region we are able to operate in.

We also note that much of the literature exploring and exploiting beam loading effects comes from the same few-cycle, Salle Noire system at LOA, to which we compared our parameters at the beginning of this chapter [239, 240, 247]. Although we initially argued that operating at longer wavelengths cannot be understood as a simple upscaling of 800 nm experiments – and in many respects this remains the case – when considering plasma parameters, a long-wavelength drive does offer several practical and logistical advantages that are consistent with such an upscaling perspective. In particular, when a longer wavelength driver corresponds to an increase in plasma wavelength, this naturally produces a larger blowout cavity, capable of supporting significantly more charge. This may help to explain why beam loading effects appear at much earlier stages in our simulations compared to at the dephasing length reported at LOA, and why we are able to reach unexpectedly high electron energies in regimes where their experiments do not – although this could also be influenced by the nearly single-cycle regime in which they operate, where CEP effects are significant. In practical terms, a larger wakefield

is advantageous both for diagnostics and for external injection mechanisms, permitting operation at lower plasma densities and simultaneously removing the requirement for extremely small targets, which are considerably more challenging to manufacture.

Given the fundamentally different parameter space that can be accessed with the Chimera laser system, it is perhaps unsurprising that our simulations suggest a regime of operation distinct from that described by the Lu framework of 2007. While Lu *et al.* presented their model as a "recipe for designing a LWFA for given laser and plasma parameters" [74], we argue that for parameter regimes far outside those accessible to most conventional laser facilities, this recipe cannot – and should not – be applied uncritically in the context of our experimental design. We note too that although the LWFA experiment performed on the 3.9 µm system at the University of Maryland [185] provides a degree of assurance towards the validity of this simulation campaign, we must nevertheless consider this work inconclusive until such time that we carry out this experiment ourselves.

6.3 Summary

In this chapter, we presented results from a density scan of laser wakefield acceleration (LWFA) simulations, driven with parameters representative of the Chimera laser system under development at Imperial College London. The scan revealed an unexpected outcome: accelerated electron energies exceeded the predictions of the widely adopted Lu *et al.* scaling law [74] to an extent that could not be readily explained.

To probe this discrepancy, we examined several possible mechanisms, including the evolution of the normalised vector potential a_0 through laser redshifting and self-focusing, the scaling of the blow-out radius, and potential contributions from carrier-envelope phase (CEP) effects, direct laser acceleration (DLA) and beam loading. While most of these mechanisms proved insufficient, beam loading effects may constitute an important factor in accounting for the high energies we observe. Our simulations suggest that an early transition to a hybrid LWFA/PWFA regime could extend the effective dephasing length, enabling electrons injected at early times to remain in phase with the accelerating field for longer.

Due to the longer wavelength of the mid-infrared Chimera driver that we simulate (3700 nm), compared to more typical experiments driven with an 800 nm beam, more extensively reported in the literature, the simulations we present operate at comparatively high ratios of plasma to critical density. Mid-infrared drivers naturally access regimes of higher n_e/n_c, affecting wakefield excitation and potentially modifying acceleration dynamics relative to the conditions assumed in conventional 800 nm experiments.

We therefore conclude that the parameter space accessed by these mid-infrared, few-cycle simulations differs fundamentally from the regimes for which the Lu framework was originally developed. While further investigation is required to fully understand these results, operating in this regime could allow access to highly efficient wakefield

excitation, even at lower laser energies and more achievable pulse durations compared to conventional near-infrared systems. It also opens a promising route for future research, combining the compactness of laser wakefield acceleration with the efficiency and stability of beam-driven plasma wakefield acceleration. In particular, it motivates exploration of tailored injection schemes, optimised plasma densities, and more controlled beam loading effects to harness the full potential of this hybrid acceleration regime in next-generation mid-infrared laser systems such as Chimera.

7 Design Considerations for a Laser Wakefield Acceleration experiment driven by the Chimera laser system

In laser wakefield acceleration (LWFA) experiments at major facilities, emphasis is typically placed on 'standard' laser and gas parameters such as pulse duration, spot size, normalised vector potential, gas constituents and gas density to produce the most stable, high charge, high energy, narrow spread electron beams. While these parameters are certainly of unquestionable importance, the most fundamental laser parameter, its wavelength, remains relatively unexplored. A long wavelength driving laser may be advantageous to a laser wakefield accelerator – this is something we want to explore firstly in simulation and latterly in experiment. Here we use the open-source particle-in-cell (PIC) code Smilei [172] to simulate a 3.7 µm central wavelength laser wakefield experiment using the Chimera laser system, currently under active development at Imperial College London.

In order to plan effectively for a first in-house LWFA experiment, it is important to try and understand not just the relationships that each of these parameters have with one another, but how the resulting effects might be experimentally diagnosed. This is not an easy task: the former requires exploring a multidimensional co-dependent parameter space, and the latter may require creative solutions where certain diagnostic equipment is unavailable. University-scale developmental laser systems, such as Chimera, may not provide perfect day-to-day reliability or stability, further complicating experimental planning. Accordingly, we seek to define the operational limits for a mid-infrared (MIR) LWFA, and to identify how other parameters might mitigate constraints in pulse duration, energy, and beam quality. In this chapter, we map the relevant parameter space for both best- and worst-case scenarios in terms of laser performance, providing a framework for the future planning and realisation of this experiment.

7.1 Laser Parameters

We begin by extending the plasma density scan introduced in Chapter 6. That scan, performed for a 25 mJ, 30 fs laser pulse, with $a_0 = 1.9$, indicated that the optimal starting plasma density for maximising electron energy gain is 2.8×10^{18} cm^{-3}. Building on this result, we now investigate how three key laser parameters – pulse energy E_{L}, pulse duration τ_{FWHM} and FWHM beam waist at best focus d_{FWHM} – affect the accelerator performance, keeping the plasma density fixed at this notionally optimal value.

Since the wakefield amplitude is governed not by energy or pulse duration alone, but by the dimensionless laser strength parameter a_0, for a Gaussian beam, we define I_0 as the peak laser intensity and λ_{L} the laser wavelength, to write:

$$a_0 = \frac{e}{\pi m_{\mathrm{e}} c^2} \sqrt{\frac{I_0}{2\epsilon_0 c}}\, \lambda_{\mathrm{L}} \tag{7.1}$$

$$\tag{7.2}$$

For a Gaussian spatial profile with a focal spot of w_0, $I(r) = I_0 e^{-2r^2/w_0^2}$ with $d_{\mathrm{FWHM}} = w_0\sqrt{2\ln 2}$ (see Appendix A.2 for the FWHM versus 1/e definitions) and a Gaussian temporal profile with FWHM τ_{FWHM}, the peak intensity relates to the pulse energy via:

$$I_0 = \frac{8E_{\mathrm{L}}(\ln 2)^{3/2}}{\pi^{3/2}\ \tau_{\mathrm{FWHM}}\ d_{\mathrm{FWHM}}^2} \tag{7.3}$$

Substituting this into the expression for a_0 gives:

$$a_0 = \frac{\lambda_{\mathrm{L}} e}{\pi m_{\mathrm{e}} c^2} \sqrt{\left(\frac{\ln 2}{\pi}\right)^{3/2} \left(\frac{4E_{\mathrm{L}}}{\epsilon_0\ c\ \tau_{\mathrm{FWHM}}\ d_{\mathrm{FWHM}}^2}\right)} \tag{7.4}$$

This scaling makes it clear that the strength of the LWFA interaction can be altered through a combination of the laser energy, its pulse duration, and its spot size at focus.

With energy scans taking place in the range 5 mJ - 55 mJ, pulse duration scans in the range 20 fs - 70 fs, (19.78 fs being the transform limit of a Gaussian 3.7 µm centred pulse with 1 µm of bandwidth), and focal spot sizes between 9.6 µm - 19.6 µm, we present the energy spectra for each 1 dimensional scan in Figure 7.1.

These plots show that for a fixed energy of 25 mJ and a FWHM beam waist of 13.6 µm, the maximum pulse duration that still allows self-injection and acceleration of

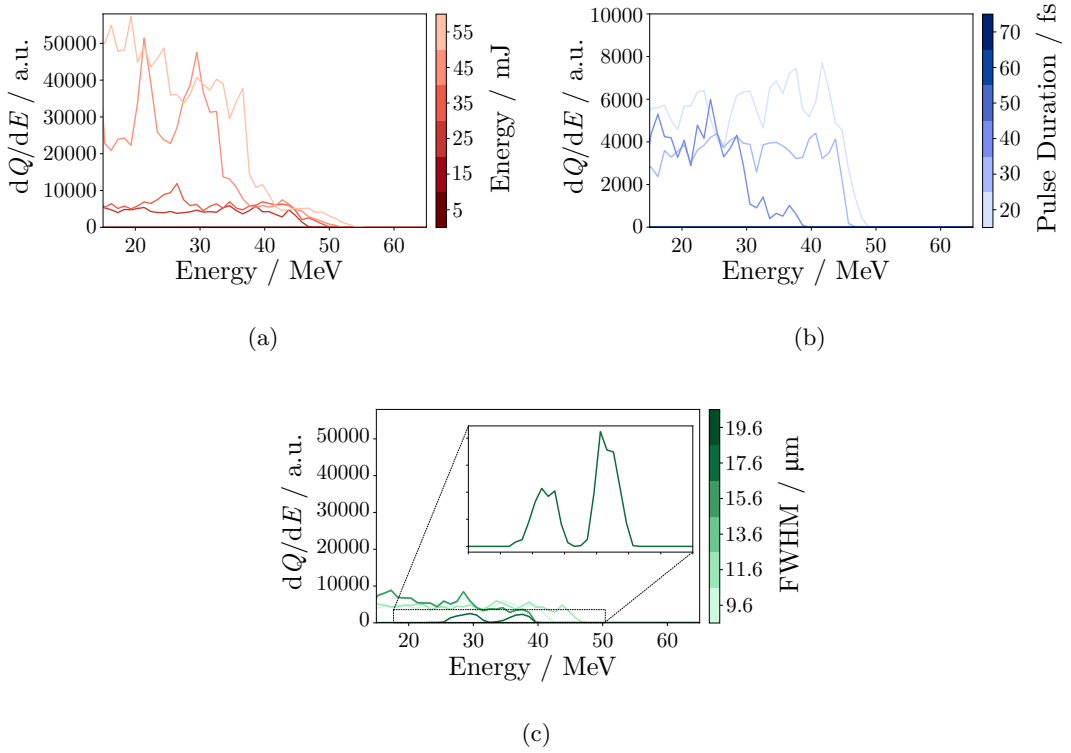


Figure 7.1: Electron energy spectra for 1D scans of (a) laser energy, E_{L}, (b) pulse duration, τ_{FWHM} and (c) focal spot, d_{FWHM}. The inset of (c) shows the quasi-monoenergetic spectrum achieved when using a spot size of $d_{\mathrm{FWHM}} = 17.6$ µm. These parameters are varied in 6 discrete steps between 5 - 55 mJ, 20 - 70 fs and for 9.6 - 19.6 µm. Within each scan, the parameters not being varied are $E_{\mathrm{L}} = 25$ mJ, $\tau_{\mathrm{FWHM}} = 30$ fs, $d_{\mathrm{FWHM}} = 13.6$ µm.

electrons is 40 fs (FWHM) – there is no available energy spectrum for the simulations above this. Similarly, for a fixed pulse duration of 30 fs and FWHM beam waist of 13.6 µm, 25 mJ appears to be the minimum laser energy capable of driving the same process. We note that for these simulations, we are operating with a perfect Gaussian pulse, meaning that half of the available energy is contained within the FWHM focal spot. This is unlikely to represent the pulse shape in an experiment, meaning that the total detected pulse energy will therefore need to be higher, depending on the focal spot quality. Around these limits, the spectrum shows much greater variation when adjusting pulse energy compared to pulse duration, suggesting that the system is more sensitive to energy fluctuations than we are to a tightly defined temporal pulse. A forgiving pulse duration is likely to be beneficial to the real system we aim to use, since, as discussed in Chapter 6, achieving clean temporal compression is often the most challenging aspect of operating a high-power CPA system. Although 40 fs remains a challenging goal for the system, it is possible that adjusting other parameters, such as the laser energy and spot size at focus, will allow more flexibility here.

When varying the FWHM beam waist, keeping pulse duration and laser energy fixed to 30 fs and 25 mJ, we find a relative tolerance. Whilst there is negligible injection seen for $d_{\mathrm{FWHM}} = 19.6$ µm, simulations using FWHM beam waists between 9.6 µm - 15.6 µm result in very similar electron energy spectra. In fact, in Figure 7.1(c), the spectra for

$d_{\mathrm{FWHM}} = 13.6$ µm and 15.6 µm are indistinguishable. This is not completely unexpected: as the laser propagates into the plasma, it will undergo some period of self-focusing as it converges towards its matched spot size. Somewhat interestingly, the simulation using $d_{\mathrm{FWHM}} = 17.6$ µm results in an electron spectrum with significantly less energy spread, with two, more monoenergetic peaks centred at 27.5 MeV and 36.5 MeV. This can be seen in more detail in the inset of Figure 7.1(c). Given previous work investigating long versus short focal length optics [225] (and here equating a smaller spot size to a shorter focal length lens), we might not be surprised that a larger FWHM spot size may lend itself to a more mono-energetic (albeit lower charge) energy spectrum.

Aside from the pronounced increase in accelerated charge observed for laser energies above 45 mJ, the resulting spectra are broadly similar across the parameter space. As we might expect for an unoptimised self-injection regime, the spectra have a wide distribution of electron energies, all showing similar amounts of injected charge. These similarities are indicative of the fact that the interaction is primarily a_0 driven, to which all three of these laser parameters will contribute. To this end, if a_0 is altered by adjusting any one of these parameters, we would expect equivalent effects on the acceleration process. Using Equation (7.4) to calculate the simulated a_0 for each parameter set, we extract the 90$^{\mathrm{th}}$ percentile of the electron energy distribution to compare how changes in each parameter influence the outcome. This is shown in Figure 7.2.

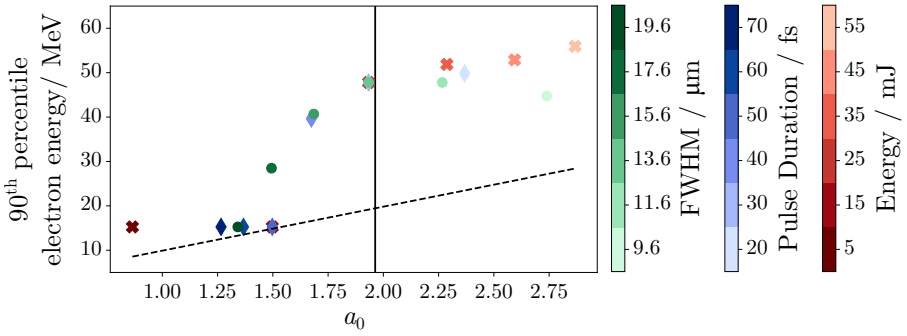
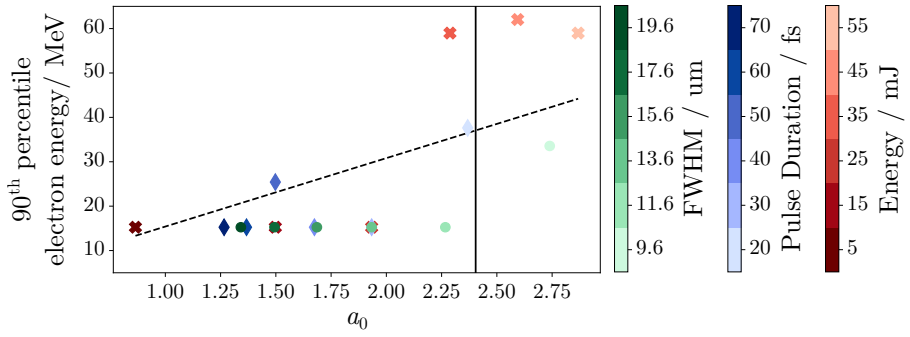


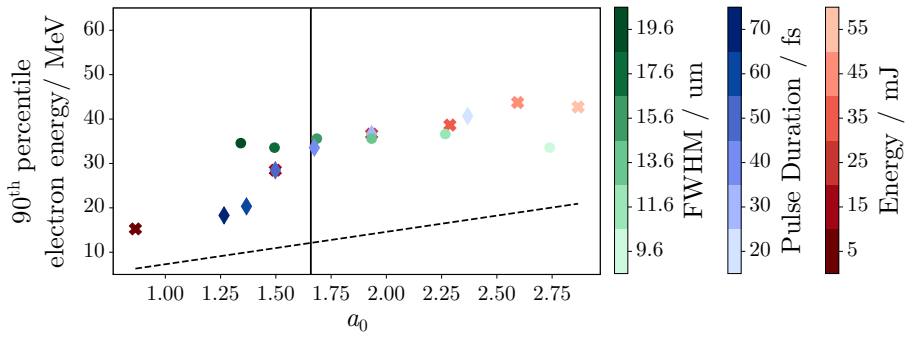
Figure 7.2: The 90$^{\mathrm{th}}$ percentile of accelerated electron energies plotted against the simulated normalised vector potential a_0. The initial plasma density is fixed at $n_{\mathrm{e}} = 2.8 \times 10^{18}$ cm^{-3}, while individual parameters (laser energy, E_{L}, pulse duration, τ_{FWHM} and focal spot, d_{FWHM}) are varied to alter a_0. The Lu *et al.* [74] energy scaling is shown as a black dashed line, and the Thomas [143] self-injection threshold as a solid black line.

Also included in the figure is the threshold for self-injection, using the model described by Thomas [143], as well as the Lu *et al.* 2007 energy scaling [74] considered in detail in Chapter 6. As perhaps we now might expect from the discussion of Chapter 6, the Lu scaling continues to fare poorly within our parameter space. The self-injection threshold, on the other hand, which aims to predict a lower bound on the a_0 required for electron injection and acceleration, seems to overestimate the threshold. Also discussed in Chapter 6, this is likely due to the evolution of the laser as it propagates through the dispersive plasma, meaning that the plotted a_0 (taken from the simulation starting parameters) is almost certainly an underestimate of the actual a_0 at injection.

This set of numerical tests provides a useful set of constraints on the laser parameters required in the lab. The next question is: if we find ourselves in a 'worst-case scenario' – for example, with a combination of parameters giving $a_0 < 1.4$ – can we manipulate the system to still achieve electron injection and energy gain? In 2012, Mangles *et al.* [144] investigated this problem on an 800 nm short-pulse system, showing that the limitations of low-energy laser beams could be offset by using a higher-density plasma. To explore the implications for our case, shown in Figure 7.3, we repeat the parameter scan at different plasma densities.



(a)
$n_e = 1.8 \times 10^{18}$ cm^{-3}



(b)
$n_e = 3.8 \times 10^{18}$ cm^{-3}

Figure 7.3: The 90$^{\text{th}}$ percentile of accelerated electron energies plotted against the simulated normalised vector potential a_0. The initial plasma density is fixed at (a) $n_e = 1.8 \times 10^{18}$ cm^{-3} and (b) $n_e = 3.8 \times 10^{18}$ cm^{-3}, while individual parameters (laser energy, E_L, pulse duration, τ_{FWHM} and focal spot, d_{FWHM}) are varied to alter a_0. The Lu *et al.* [74] energy scaling is shown as a black dashed line, and the Thomas [143] self-injection threshold as a solid black line.

Although the scans do show that, by using a higher plasma density, we would have perhaps a little more leeway for sub-optimal laser performance, they more importantly imply that the laser parameters we proposed as interchangeable show considerably more variation at both higher and lower densities. This is particularly evident in the lower density scan (1.8×10^{18} cm^{-3}, Figure 7.3a) where, for example, despite all three parameters producing $a_0 \simeq 2.25$, the resulting electron energies differ quite dramatically.

Extending this scan into a 3-dimensional parameter space (a computationally expensive exercise), Figure 7.4 shows how the 90$^{\text{th}}$ percentile electron energies are affected through combined changes in laser energy and pulse duration, for three different plasma densities.

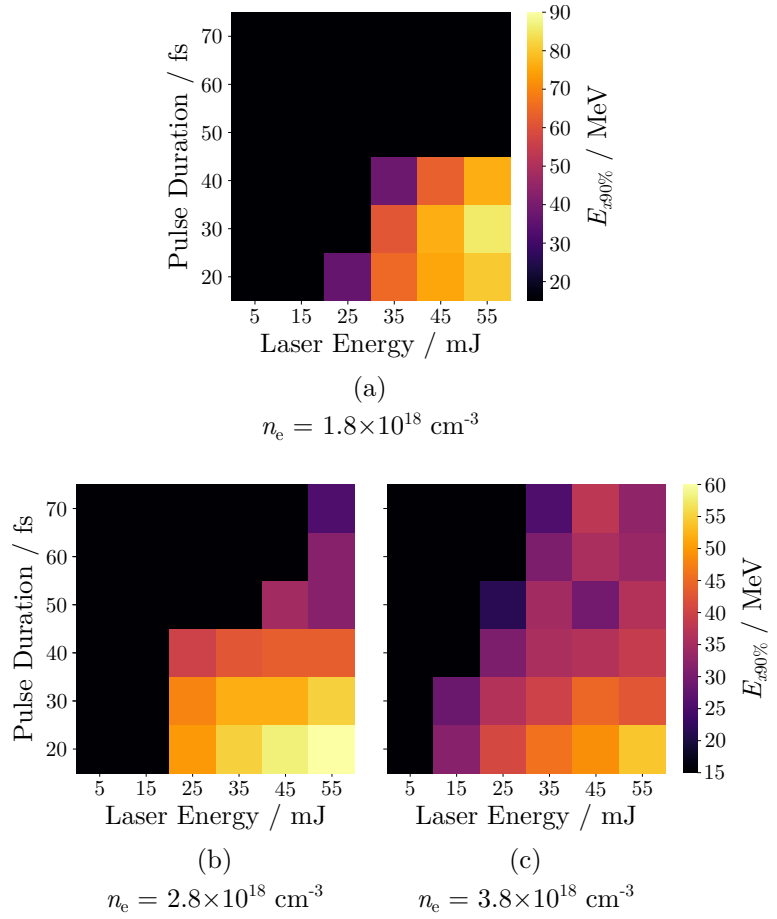


Figure 7.4: The 90$^{\text{th}}$ percentile electron energies extracted from a 3-dimensional parameter scan across laser energies $E_{\text{L}} = 5 - 55$ mJ and pulse durations $\tau_{\text{FWHM}} = 20 - 70$ fs for plasma densities (a) $n_{\text{e}} = 1.8 \times 10^{18}$ cm^{-3}, (b) $n_{\text{e}} = 2.8 \times 10^{18}$ cm^{-3} and (c) $n_{\text{e}} = 3.8 \times 10^{18}$ cm^{-3}. Note the different intensity scale for (a) compared to (b) and (c).

The result of this foray into 3D parameter space is a relatively consistent one. The panels in Figure 7.4 show a predictable trend within the space: increasing the pulse energy is more or less synonymous with decreasing pulse duration, and corresponds to an increase in electron energies. Higher densities enable injection and acceleration with less demanding laser parameters, although, as expected, this comes at the cost of lower electron energies. At higher densities, more charge is injected into the wake, which broadens the spectrum and shifts a substantial fraction of electrons to lower energies. As a result, the 90$^{\text{th}}$ percentile energy is reduced, both due to maximum energies being lower, but also because the larger low-energy population pulls down the percentile threshold. A lower starting density operates much closer to the density threshold for self-injection, given by the Mangles *et al.* [144] model – there are fewer injected electrons and thus a

narrower distribution. For the parameter space explored here, the self-injection threshold spans $1.3 - 5.4 \times 10^{18}$ cm^{-3}. By contrast, at lower densities, injection is closer to the self-injection threshold predicted by Mangles *et al.* [144].

For the Chimera laser, these results suggest that the parameter space required for a first demonstration experiment is broadly flexible. If the eventual laser performance lies toward the lower end of the explored range, however, operation would require much higher plasma densities. Running a smaller scan for this eventuality, we see that even at a density of $n_e = 5.8 \times 10^{18}$ cm^{-3} – a density above the predicted self-injection threshold for the worst-case laser parameters – pulse energy is considerably more important than pulse duration, with a lower bound of around 15 mJ for electron acceleration. This is approximately half the current maximum pre-compression energy at 3.7 µm already demonstrated by Chimera.

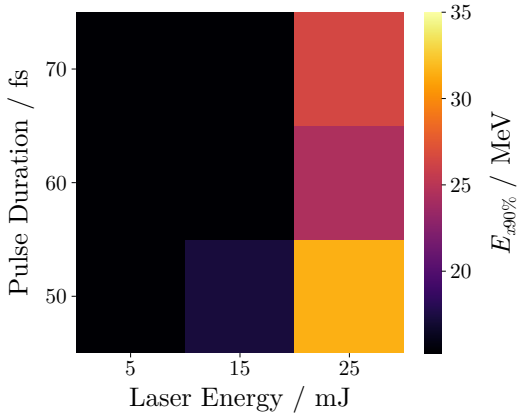


Figure 7.5: The results of a 3-dimensional parameter scan across conservative laser parameters of $E_L = 5 - 25$ mJ and pulse durations $\tau_{FWHM} = 50 - 70$ fs for a plasma density of $n_e = 5.8 \times 10^{18}$ cm^{-3}. This is just above the threshold for self-injection of $n_e = 5.4 \times 10^{18}$ cm^{-3} for 'worst-case' laser parameters $E_L = 15$ mJ, $\tau_{FWHM} = 70$ fs.

To explore whether increasing the laser a_0 via a smaller focal spot could partly compensate for suboptimal laser parameters, Figure 7.6 shows that, within this high-density regime, reducing the spot size cannot offset other limitations. This implies that the lower bound for laser energy to drive LWFA on the Chimera system is approximately 15 mJ, with a maximum compressed pulse duration of ~ 50 fs.

For the long-term planning and operation of Chimera, it would therefore be advantageous to ensure sufficient pulse energy (and focal spot quality), enabling operation at longer pulse durations and lower plasma densities. If the goal is to optimise electron energies, operating near $n_e = 2.8 \times 10^{18}$ cm^{-3} provides a stable regime: one more tolerant to laser parameter fluctuations, while still producing average electron energies around 40 MeV. Electron beams in this energy range are highly valuable for a broad range of applications. They are energetic enough to penetrate matter and drive secondary processes, while remaining more compact and practical than conventional RF-generated sources. In particular, 10 - 100 MeV electrons are well-suited for probing structural

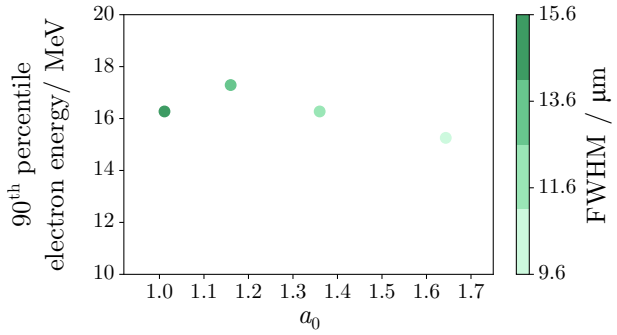


Figure 7.6: The 90$^{\text{th}}$ percentile of accelerated electron energies plotted against the simulated normalised vector potential a_0. The initial plasma density is fixed at $n_e = 5.8 \times 10^{18}$ cm^{-3}, laser energy, E_L and pulse duration, τ_{FWHM} are fixed to $E_L = 15$ mJ and $\tau_{FWHM} = 50$ fs, while the focal spot, d_{FWHM} is varied to alter a_0.

dynamics on the femtosecond timescale [248–250], and for generating secondary radiation such as bremsstrahlung [251, 252].

7.2 Plasma Considerations

Having previously explored the role of plasma density, concluding that $n_e = 2.8 \times 10^{18}$ cm^{-3} is likely the most stable region of accelerator operation for a first Chimera driven experiment, we now turn to the question of targetry. Although other injection schemes should be considered in future campaigns, particularly if optimisation of electron beam quality becomes a priority, in this section, we consider the simplest commissioning experiment, and therefore focus on self-injection only.

Here, we are looking to determine the optimum accelerator length. This is an important parameter when considering experimental design: if electron energies are relatively insensitive to accelerator length, a supersonic gas jet of fixed size, which is straightforward and reliable, would be sufficient. Gas jets also have the advantage of lower average gas load and higher compatibility with higher repetition rates, while gas cells impose much heavier vacuum demands, typically restricting operation to lower repetition rates despite offering superior stability. If, however, energy gain is highly sensitive to accelerator length, a gas cell would be more appropriate, as it allows *in situ* tuning of the interaction length, albeit at the cost of greater experimental complexity.

Across the same laser parameter space presented in Section 7.1, we extract the acceleration distance corresponding to the maximum (90$^{\text{th}}$ percentile) energies reached by injected electrons. Using Equation (7.4) to combine the laser parameters into a common a_0, this is shown for plasma densities $n_e = 1.8 \times 10^{18}$ cm^{-3}, 2.8×10^{18} cm^{-3} and 3.8×10^{18} cm^{-3} in Figure 7.7.

For this analysis, to compare this to the Lu scaling for electron dephasing length L_d (Equation (3.48)) [74] – the length scale which should set an approximate upper bound upon the accelerating distance – as in Chapter 6, we look not solely at the absolute

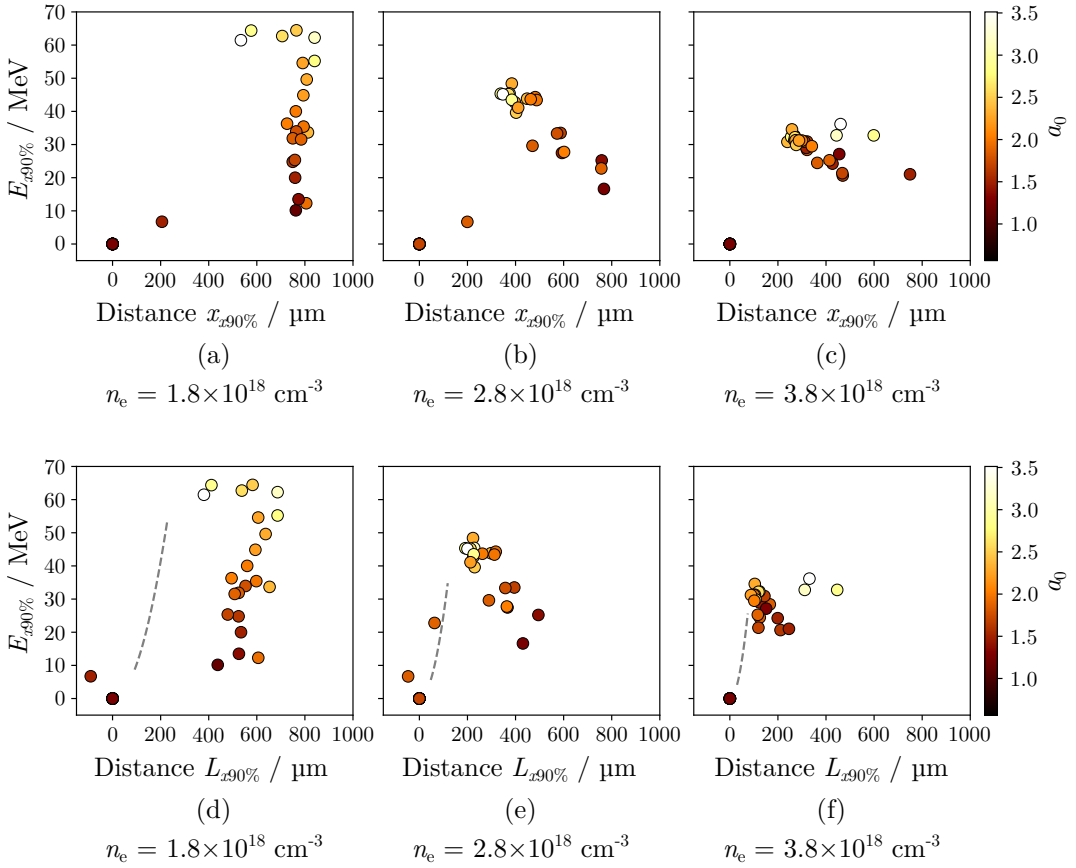


Figure 7.7: Acceleration distance corresponding to the maximum (90$^{\text{th}}$ percentile) energies reached by injected electrons. (a), (b) and (c) show the absolute position in the lab frame, whereas (d), (e) and (f) show the relative position over which the electron has been accelerating. All simulations within the 3D simulation space are included. The columns show simulations with differing plasma densities of $n_e = 1.8 \times 10^{18}$ cm^{-3}, $n_e = 2.8 \times 10^{18}$ cm^{-3} and $n_e = 3.8 \times 10^{18}$ cm^{-3}. The laser energy and pulse duration are combined into a common a_0 using Equation (7.4).

position of the electron in the lab frame, $x_{x90\%}$, (panels (a) - (c) in Figure 7.7), but the relative distance over which the electron has been accelerating in the wake, $L_{x90\%}$, (panels (d) - (f) in Figure 7.7). This subtracts the distance travelled by the laser before electron injection – subtracting the driver's motion isolates the effective acceleration length. For panels (d) - (f) in Figure 7.7, we include this scaling, given by Equation (6.18).

We find that, in most cases, the maximum electron energy is reached at approximately the same spatial position for each plasma density, suggesting that the dephasing length is more sensitive to the plasma density than to the laser parameters. Excluding data which sees electrons accelerated to less than 10 MeV, for the three density steps (increasing) we see average accelerating distances of $L_{x90\%}$ equal to 550(80) µm, 300(100) µm and 160(90) µm respectively, where the error represent the standard deviation of $L_{x90\%}$ across all laser parameter combinations at that particular density. Given the predicted $1/n_e$ scaling of the dephasing length compared with the weaker $\sqrt{a_0}$ dependence, this is not

unexpected, however, as with the electron energy scalings presented in Chapter 6, the Lu model appears to underperform when tested against our simulations.

While the absolute optimum accelerator length is a useful benchmark, in a real experiment, it is often more practical to understand the distance over which electrons retain their maximum energies before dephasing begins. Of greater interest, therefore, is the relative evolution of the electron spectra in time, and the characteristic length scale this corresponds to. Figure 7.8 shows six representative electron spectra across the parameter space. We consider three plasma densities, $n_e = 1.8,\ 2.8,\ 3.8 \times 10^{18}$ cm^{-3}, at pulse energies and durations of 25 mJ / 30 fs and 45 mJ / 40 fs.

These spectra are broadly representative of trends observed across the full parameter space. Panel (a) shows a regime insufficient for self-injection and acceleration. Panel (b) sees a small population of electrons accelerated to high peak energies with an energy spread of $\sim$ 20 MeV, while panel (c) shows greater injected charge but with a much larger spread. Panels (d), (e), and (f) all exhibit the parabolic signature characteristic of electrons gaining energy by acceleration, before curving over in the energy phase space as they outrun the wake and dephase. Also evident in these plots are tracks which appear to 'wiggle' as they accelerate and decelerate. These are consistent oscillatory modulations in the electron tracks consistent with beam loading and multiple injection effects, which become more pronounced with both increasing plasma density and increasing laser energy.

Returning to the matter of experimental design, what we see, consistent within almost all spectra within the parameter space (other than those closer to the self-injection threshold), is that once electrons reach their peak energies, there seems to be a very marginal decrease in energy throughout the remainder of the simulation, indicating that we are seemingly fairly insensitive to the length of the accelerator, and a supersonic gas jet would be sufficient for our purposes.

The reason for this behaviour is not immediately obvious, but it is likely also the result of beam loading. Once accelerated, electron bunches subsequently travel with velocities very close to c. The laser driver has been depleted, and because of the large amount of accelerated charge, the electrons begin to drive their own wake. As a result, the electrons do not lose energy through dephasing; instead, they propagate almost freely through the plasma. To illustrate this, Figure 7.9 shows an evolution in density space of the $n_e = 2.8 \times 10^{18}$ cm^{-3}, $E_L = 25$ mJ, $\tau_{FWHM} = 30$ fs case, up to a simulation time of 6 ps.

Four timesteps are shown. Panel (a), at 1.7 ps, corresponds to the point at which the electron energy peaks. The laser has not fully depleted, however, as evidenced by the low-density region following the electron bunch, the electrons have started to drive their own wake. Panel (b) is at 2.3 ps. At this stage, the electron bunch has caught up to what remains of the laser – the electrons have fully dephased; however, this has not come at a significant detriment to their energy. The bubble has elongated longitudinally

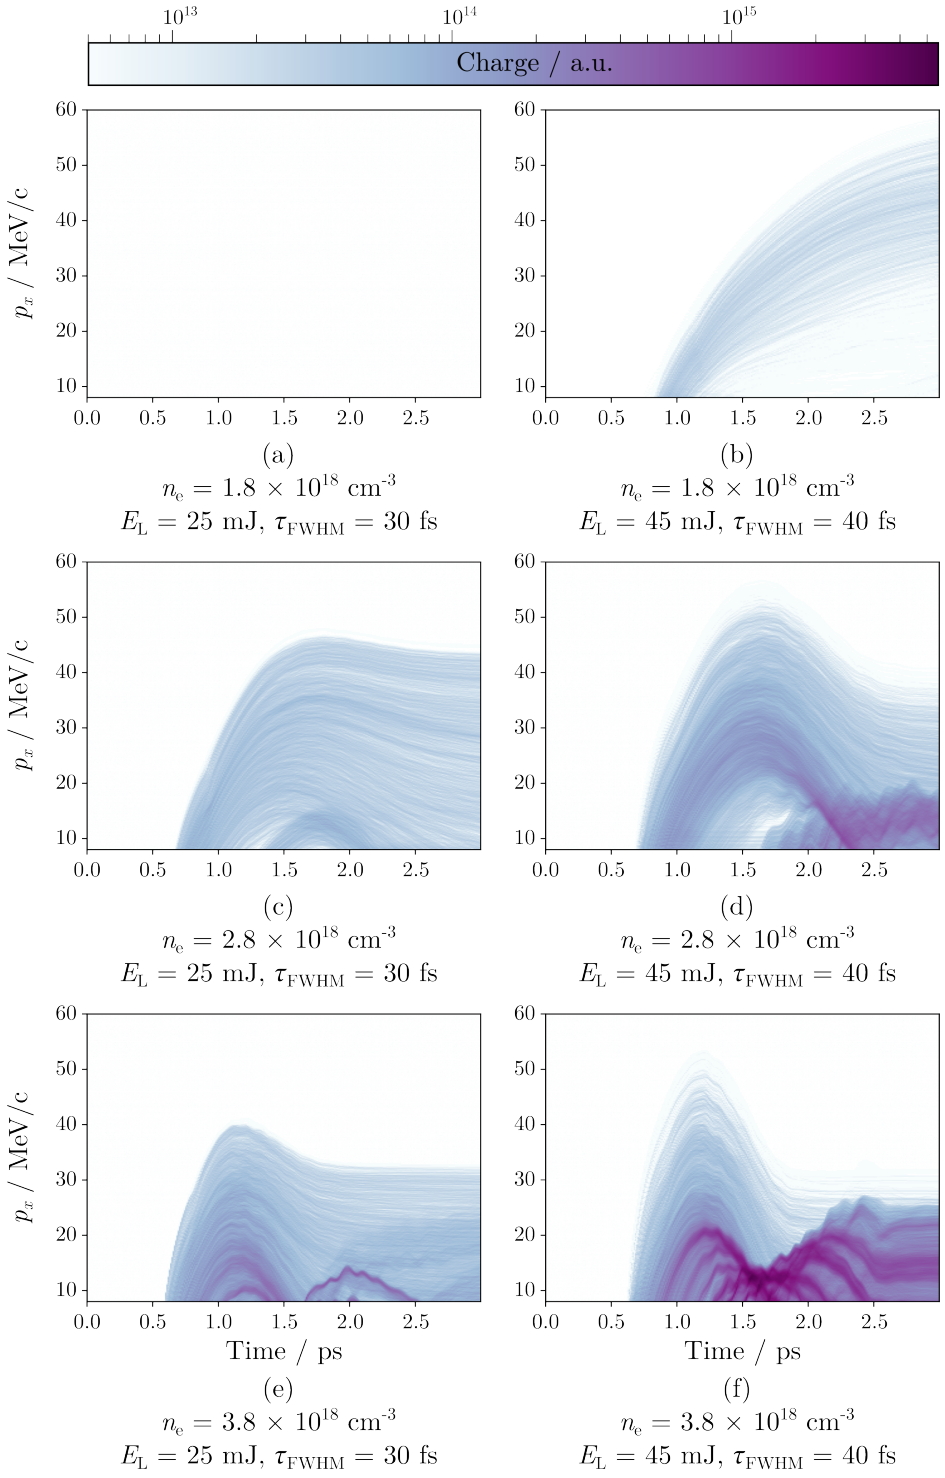


Figure 7.8: Representative electron spectra across the explored parameter space. The rows show simulations with differing plasma densities of $n_e = 1.8 \times 10^{18}$ cm^{-3}, $n_e = 2.8 \times 10^{18}$ cm^{-3} and $n_e = 3.8 \times 10^{18}$ cm^{-3}, while the columns show varying laser energies and pulse durations of 25 mJ, 30 fs and 45 mJ, 40 fs. The x- and y-axes are identical for all cases, as is the logarithmic intensity scale.

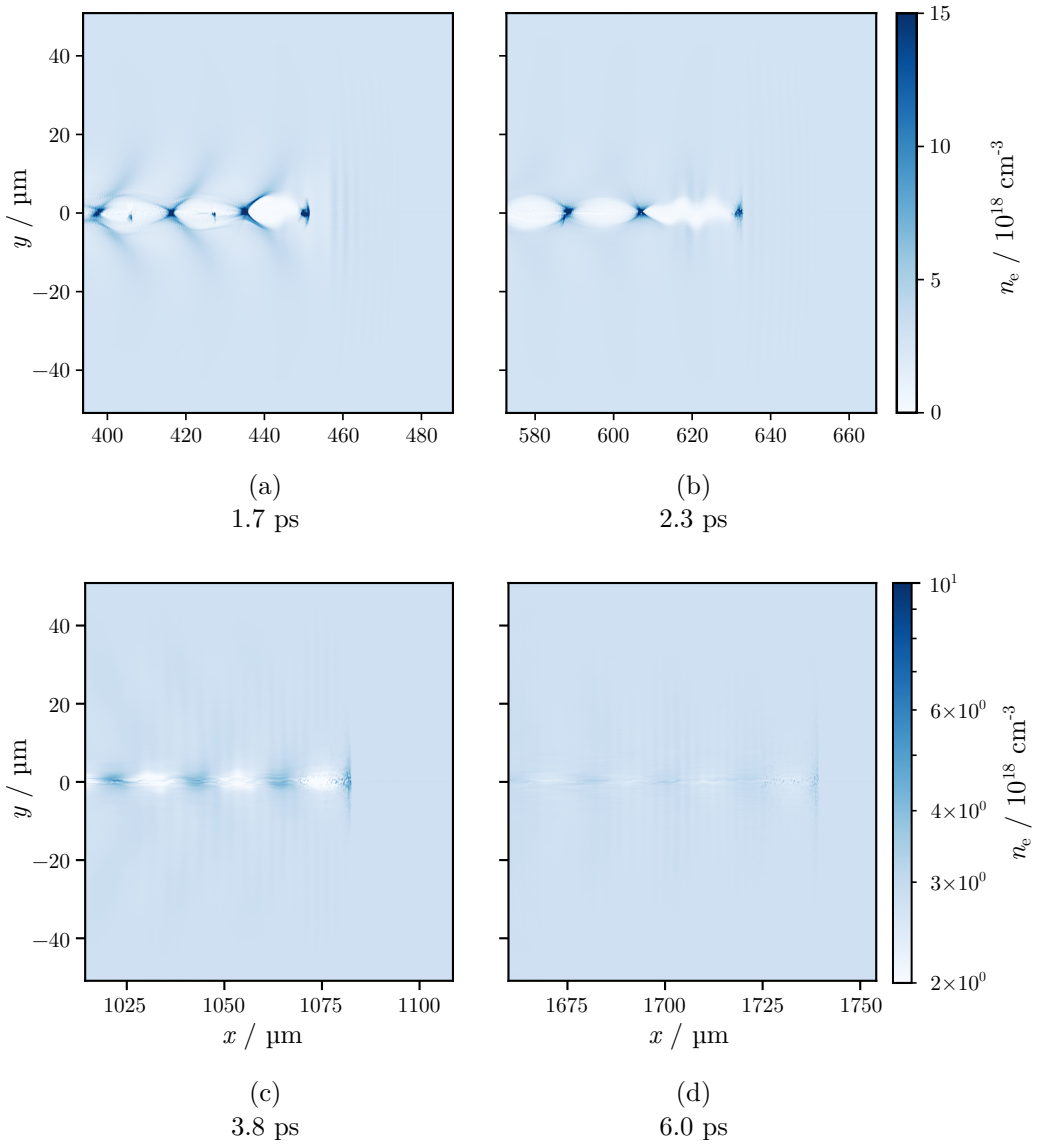


Figure 7.9: Evolution of in density space of a simulation at $n_e = 2.8 \times 10^{18}$ cm^{-3}, $E_L = 25$ mJ, $\tau_{FWHM} = 30$ fs, with snapshots taken at (a) 1.7 ps, (b) 2.3 ps, (c) 3.8 ps and (d) 6.0 ps. The y-axes are shared. (a) and (b) share a linear intensity scale, (c) and (d) share a logarithmic one. (a) At 1.7 ps, near the point of peak electron energy, the trailing low-density region indicates that the accelerated bunch has begun to drive its own wake. (b) At 2.3 ps, the bunch overlaps with the depleted laser field and has fully dephased. The wake cavity shows signs of longitudinal elongation and transverse contraction, consistent with the onset of a beam-driven regime. (c - d) At later times (3.8 ps and 6.0 ps), the bunch is further depleted, though residual density modulations remain in its wake.

and contracted transversely, as expected when there is a large amount of injected charge. Panel (c) at 3.8 ps and panel (d) at 6.0 ps show the much depleted electron beam, although the density modulation can still be seen in its trail.

This case provides a realistic example of what a LWFA experiment driven by the Chimera laser could look like within the explored parameter space. In almost all cases, even those with substantially different laser parameters, we see strong evidence of electron bunches beginning to drive their own wakes. At higher plasma densities and larger laser

a_0, this frequently coincides with a secondary injection into the particle-driven wake, opening the possibility of a full transition from beam-loaded LWFA to plasma wakefield acceleration (PWFA). Seen experimentally within 800 nm driver systems [242, 243, 245], this hybrid LWFA/PWFA scheme is discussed in more detail in Chapter 6. From the perspective of experimental commissioning, the implication is reassuring: beam-driven effects suggest that the system is likely to be relatively insensitive to accelerator length, even under conditions of variable or imperfect laser performance.

7.3 Discussion

The results from this study are encouraging. Contrary to what might have been expected from its contribution to the normalised vector potential, the relative insensitivity to the pulse duration appears to offer considerable flexibility for a LWFA commissioning experiment on the Chimera laser. From the perspective of laser development, this is a very positive outcome: as discussed in Chapter 5, achieving and measuring clean, short, temporal compression is particularly challenging in the mid-infrared spectral range. The ability to drive a self-injecting wakefield accelerator instead shows a stronger dependence on pulse energy – a parameter that is comparatively easier to measure with confidence and optimise experimentally. There also appears to be considerable scope to compensate for poor laser performance: increasing the plasma density relaxes laser requirements, while tighter focusing – although detrimental to optimising electron spectra – raises a_0 at focus and can, in regimes of stable operation, partially offset limitations in pulse energy or duration. The challenge for the Chimera laser will therefore be to maximise the amplified mid-infrared energy through the OPA chain, ensure highly efficient subsequent compression, and maintain a spot quality that confines a large fraction of the available energy within the FWHM focal spot.

The laser parameters aside, what is particularly interesting amongst these results is the seeming insensitivity to acceleration length. This will provide a significant advantage for a commissioning experiment as it allows the removal of a degree of freedom: if electrons are injected at any point in the accelerator, it is unlikely that they will lose energy as a result of pump depletion or dephasing before we are able to detect them. Such robustness is particularly valuable in the mid-infrared regime, where experimental diagnostics are challenging, and the parameter space is relatively unexplored.

In this analysis, we have consistently excluded tracking of electrons with energies below 10 - 15 MeV. This was done to minimise noise from background plasma electrons, but it likely means that a mid-infrared LWFA experiment could, in fact, tolerate laser parameters less optimised than the 15 mJ, 50 fs suggested here as a lower bound. Indeed, the only MIR LWFA experiment that has, to the best of our knowledge, been undertaken, was driven with considerably less favourable parameters: $E_\mathrm{L} = 23$ mJ, a minimum $\tau_\mathrm{FWHM} = 105$ fs, and $d_\mathrm{FWHM} = 30$ µm, corresponding to an $a_0 = 0.47$, scanning densities from $n_\mathrm{e} = 1.8 \times 10^{18}$ - 8.0×10^{19} cm^{-3} [185]. The highest electron energy they report is 12 MeV, indicating that their experiment probed well below the energy range we have

considered here. The main conclusion from this work is that a MIR driver is capable of accelerating significantly more charge than its 800 nm counterparts. For many potential applications of laser-plasma acceleration, the importance of having a high charge (and therefore high current) electron beam cannot be understated. Applications such as driving an irradiation source for radiobiological research [253, 254], the generation of betatron x-rays for phase-contrast imaging and tomography [255, 256] and the efficient production of nuclear isomers [257] have much less stringent requirements on electron energy, with the key performance metric instead being towards high beam currents. Since near-infrared LWFA systems are often limited by achievable charge and peak current, mid-infrared systems such as Chimera are particularly well suited to driving wakefield accelerators for this class of applications. Although our MIR simulations focus on electron energies above those seen by Woodbury *et al.* [185], we too observe an excess of accelerated charge, something which may be of great significance for many of these applications. Figure 7.10 shows the total charge seen within the parameter space explored in this chapter, extracted at the point where the average energy (as opposed to the 90$^{\text{th}}$ percentile energy) was maximal.

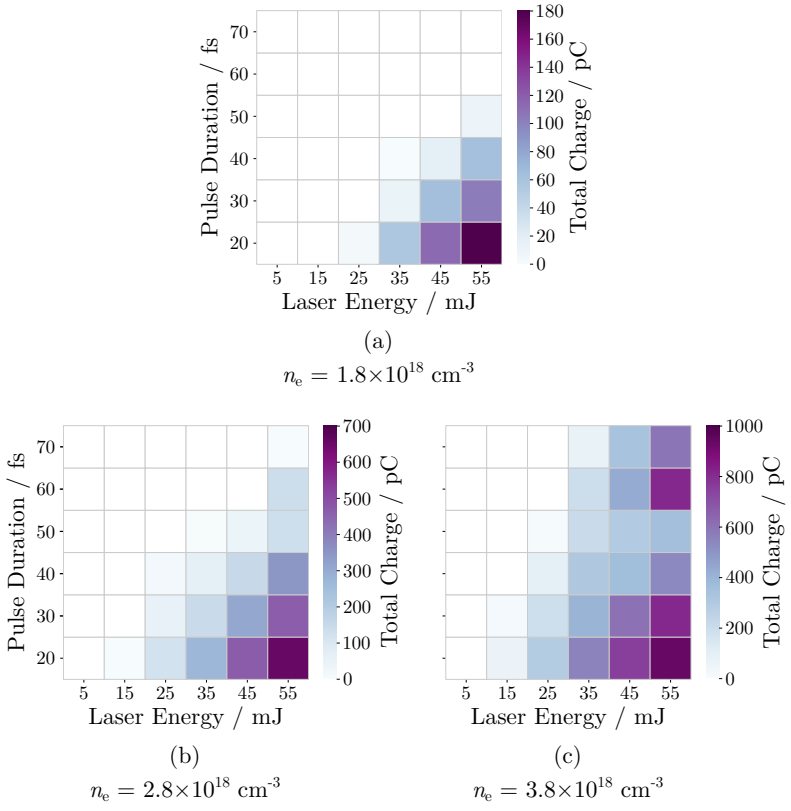


Figure 7.10: The total charge at maximum average energy extracted over a 3-dimensional parameter scan across laser energies $E_L = 5 - 55$ mJ and pulse durations $\tau_{FWHM} = 20 - 70$ fs for plasma densities (a) $n_e = 1.8 \times 10^{18}$ cm^{-3}, (b) $n_e = 2.8 \times 10^{18}$ cm^{-3} and (c) $n_e = 3.8 \times 10^{18}$ cm^{-3}.

For applications requiring only modest electron energies, this quantification of accelerated charge should be considered a conservative underestimate, as the peak charge does

not coincide in time with the maximum average energy. Even so, the values presented here already demonstrate extremely high laser-to-beam conversion efficiencies, particularly at higher densities, where electron beam energies of ~ 20 mJ correspond to $\sim 40\%$ laser-to-energy conversion. If realised experimentally, this would imply roughly 22 nC per joule of laser energy and, using the longitudinal bunch length ($\sim$10 µm) to estimate duration ($\sim$33 fs), peak beam currents in the 10s of kiloamperes – well-suited to a wide range of applications. This likely underestimated metric would match (and potentially exceed) a recent result from LOA, who report a maximum 20.5 nC/J at $E_{\mathrm{L}} = 0.12$ J, $\lambda = 800$ nm, with a conversion efficiency of 6% [258].

Taken together, these results suggest that a Chimera-driven LWFA experiment should be experimentally achievable with a broad tolerance to pulse duration, a manageable lower bound on pulse energy (of order a factor of two lower than already demonstrated), and a reduced sensitivity to accelerator length. Although optimising compression and spot quality will remain central technical challenges, the simulations indicate that even modestly performing mid-infrared pulses are capable of producing high-charge electron bunches at efficiencies far exceeding those typically seen at 800 nm (when in a sub-terawatt regime of laser power) [259, 260]. This combination of experimental robustness and access to high-current beams makes the Chimera system not only a viable driver for first-of-its-kind MIR LWFA commissioning but also a uniquely promising platform for exploring applications where charge and beam current are of primary importance.

7.4 Summary

In this chapter, we have explored a parameter space representative of the Chimera laser system's expected performance, with the aim of informing the design of a laser wakefield accelerator driven by its mid-infrared beam. Above a minimum requirement of ~ 15 mJ of laser energy, the space was found to be broadly flexible. Pulse duration requirements are less stringent, although if the energy is limited to 15 mJ, a maximum of ~ 50 fs is needed to sustain self-injection. Under these 'worst-case' parameters, operating at a higher plasma density of $n_{\mathrm{e}} = 5.8 \times 10^{18}$ cm^{-3} would nevertheless enable a commissioning experiment, while optimisation of electron energies points instead to $n_{\mathrm{e}} = 2.8 \times 10^{18}$ cm^{-3} as a more stable operating regime.

The optimum accelerating distance was also found to be relatively forgiving. A supersonic gas jet of length 450 - 1000 µm should be sufficient for commissioning, although this will ultimately depend on the accessible laser parameters. Interestingly, this insensitivity to acceleration length appears to result from strong beam loading: the electron bunches carry enough charge to drive their own wake, enabling continued propagation beyond the expected dephasing and pump depletion lengths.

In addition to electron energy, these simulations highlight the striking capacity of mid-infrared drivers to generate high charge beams. Consistent with experimental observations by Woodbury *et al.* [185], our results show that, even under relatively modest laser

parameters, MIR-driven accelerators can support an order of magnitude more accelerated charge than their 800 nm counterparts. This comes with high laser-to-beam conversion efficiencies, in some cases approaching 40%, and corresponds to kiloampere-level beam currents. For many applications of laser-plasma acceleration – from radiobiological irradiation to betatron x-ray generation and nuclear isomer production – high current, and therefore charge, is of greater importance than high electron energy. As such, this property of MIR-driven wakefields provides a compelling case for their experimental development on Chimera, even before optimisation of beam quality or energy is achieved.

8 Conclusions and Future Outlook

This thesis has explored the capabilities of the Chimera short-pulse, high-power mid-infrared laser system at Imperial College London, with the goal of enabling routine experimental operation and assessing its suitability for driving a laser wakefield acceleration (LWFA) experiment. Supported by detailed numerical simulations, the work has combined technical development, experimental commissioning, and theoretical investigation.

The technical report in Chapter 5 highlighted both the novelty of Chimera within the landscape of high-power MIR systems and the day-to-day practicalities of operating a complex university-scale system. With nine cascaded nonlinear stages, the system demands careful diagnosis and optimisation at every stage to ensure reliable performance. Particular progress has been made in improving the front end, including the Ti:Sa oscillator and Nd:YLF ring booster, as well as in developing a suite of simple, low-cost diagnostics to better characterise the MIR output. The system's first commissioning experiment – a MIR-driven (3700 nm, $\sim$ 1000 nm bandwidth) solid-target high harmonic generation study using crystalline ZnSe – produced a signal consistent with harmonics up to the 8^{th} order, an encouraging benchmark for future MIR-based experiments. As the output of Chimera's first MIR OPA already exceeds reported pulse energies in the literature for many systems, this motivates a further exploration of materials, sample geometries, and post-compression performance.

Chapter 6 presented a plasma density-scan simulation campaign using parameters consistent with Chimera's (optimistically) predicted MIR output. Surprisingly, the resulting LWFA electron energies significantly outperformed the widely used scaling laws of Lu *et al.* [74]. This potential overperformance suggests that MIR drivers may access a previously underexplored regime of parameter space. Although alternative explanations such as enhanced a_0 scaling and direct laser acceleration were considered, the most plausible mechanism appears to be the early onset of a hybrid beam-driven phase due to the unusually high n_e/n_c ratio accessible at MIR wavelengths. This highlights the limitations of applying standard NIR-based models to MIR-driven systems and points to the need for new scaling frameworks tailored to this regime, informed by experiment where possible.

In Chapter 7, the operational requirements for an LWFA experiment driven by Chimera were evaluated, with the results indicating that the constraints on laser performance are less restrictive than initially anticipated. Self-injection can be achieved with a minimum of ~ 15 mJ of laser energy (half that already demonstrated experimentally) and a pulse duration below ~ 50 fs at a plasma density of $n_e = 5.8 \times 10^{18}$ cm^{-3}. Higher laser energies allow for effective accelerator operation for longer pulse durations (up to ~ 70 fs) while still reaching electron energies of ~ 30 MeV. Unexpectedly, the optimum acceleration conditions were found to be relatively forgiving – a significant advantage, particularly for a commissioning experiment, as it effectively removes one degree of freedom in the initial experimental setup and tuning. This tolerance also suggests that a supersonic gas jet will likely be sufficient, avoiding the added complexity of gas-cell targets. The most likely explanation for this relative insensitivity to acceleration length is the onset of beam loading effects, which enable propagation well beyond the usual dephasing and depletion limits of a more standard 800 nm driven system. A further consequence of beam loading is the acceleration of a substantial charge across the parameter space, with high laser-plasma energy efficiency of $\sim 40\%$ and correspondingly high beam currents in the 10s of kiloamperes. These results imply that MIR drivers such as Chimera could serve as robust sources of high-charge electron beams driven by a tabletop-scale system, a feature of particular importance for industrial and medical applications. More broadly, these results provide a practical framework for planning future LWFA experiments with Chimera, identifying clear performance thresholds that the laser needs to reach, while highlighting opportunities to exploit the unique advantages of MIR-driven acceleration.

Although these results have been discussed in isolation, it is important to recognise that they are intrinsically linked – not only do these results help us guide and understand the direction of future MIR-driven LWFA experiments, they also inform the ongoing design and optimisation of the Chimera laser system itself. Motivated by the hypothesis that longer-wavelength drivers may offer significant advantages (supported by wavelength scaling arguments), the detailed numerical simulation results presented in this thesis suggest that this is indeed a promising avenue for experimental pursuit. From an applications perspective, charge and beam current are of paramount importance – particularly if both the charge and the higher-than-expected electron energies observed in simulation are confirmed experimentally.

In terms of the current laser system status, although the output bandwidth from the DFG stage is sufficient to support a 30 fs pulse, the present design indicates that a more realistic minimum pulse duration, after subsequent OPA stages and compression, will be closer to 70 fs. While the compressed pulse duration has yet to be measured, this sets a clear design goal: although the simulation results predictably motivate a reduction of pulse duration, if the pulse energy is sufficiently high (> 25 mJ), a 70 fs pulse remains entirely adequate for driving the planned experiment. This stronger dependence on energy rather than pulse duration plays to Chimera's strengths: its scalability through the Öffner stretcher, coupled to a cascaded series of 1053 nm pumped OPA stages,

provides greater operational flexibility compared to alternative systems, many of which are more tightly constrained by both pulse energy and duration limits.

Looking beyond an initial LWFA commissioning experiment, Chimera's multi-beam capabilities open further opportunities for more complex experimental geometries. Multi-wavelength pump–probe experiments are a natural extension, but the potential becomes even more compelling in the context of wakefield acceleration. A long-wavelength driver produces a comparatively large plasma bubble, offering the possibility for a probe beam to resolve much finer structural detail of the acceleration process. For instance, an 800 nm probe alongside a long-wavelength driver would deliver nearly three times the resolution of a 500 nm probe in a conventional 800/1000 nm-driven experiment, providing unprecedented insight into bubble and electron injection dynamics.

Even more ambitious is the prospect of using the MIR beam to generate a large, quasi-linear bubble while employing the 800 nm beam to drive the actual acceleration process – an experimental regime that, to the best of our knowledge, has never been realised. If the mid-infrared pulse were used primarily to shape and drive the large bubble, while the 800 nm pulse was injected at the correct phase within that bubble, the two could in principle act together: the MIR setting the stage and the NIR delivering additional drive or controlling injection. This hybrid scheme closely follows the concept of two-colour injection proposed by Yu *et al.* [261], in which a long-wavelength driver creates a suitable wakefield structure and a secondary pulse is used for controlled electron injection. In our case, both pulses would originate from the same laser system, which could offer very robust intrinsic synchronisation compared to two independent sources.

The practical challenges would, however, remain substantial. The pulses must be synchronised to within a fraction of the plasma period, demanding timing stability at the few-femtosecond or sub-femtosecond level. The plasma density must be chosen carefully so that the MIR pulse can drive a strong wake while the 800 nm pulse propagates as intended. Ionisation dynamics add further complexity, as the MIR front edge may pre-ionise the medium and alter the interaction of the subsequent 800 nm pulse. Overlapping two high-intensity pulses also raises the possibility of cross-instabilities and nonlinear coupling that could degrade the wake.

Nonetheless, the potential advantages make this an intriguing idea. A large bubble driven by the MIR pulse could provide a structure into which the 800 nm pulse can inject or further accelerate electrons, combining the scalability of long-wavelength drivers with the precision and higher-energy benefits of shorter wavelengths. Detailed particle-in-cell simulations scanning density, pulse delay, and relative intensities would be the logical next step before considering an experimental demonstration.

A Appendix A

A.1 Fundamental Constants

Table A.1 defines all fundamental constants used within this thesis.

Symbol	Name	Value	Unit
c	Speed of light	2.998×10^{8}	m s^{-1}
e	Elementary charge	1.602×10^{-19}	C
m_e	Electron rest mass	9.109×10^{-31}	kg
$m_e c^2$	Electron rest energy	0.5110	MeV
ϵ_0	Vacuum permittivity	8.854×10^{-12}	F m^{-1}
μ_0	Vacuum permeability	1.257×10^{-6}	H m^{-1} or N A^{-2}
h	Planck's constant	6.626×10^{-34}	m^2 kg s^{-1}
$\hbar$	Reduced Planck's constant (hbar)	1.055×10^{-34}	m^2 kg s^{-1}

Table A.1: To 4 significant figures, the fundamental constants used within this thesis.

A.2 The definition of the full-width-at-half-maximum

Across laser and plasma research, we often refer to the full-width-at-half-maximum (FWHM) of various quantities. While this at first glance is straightforward, we must be precisely clear as to what we mean here, particularly as discrepancies in definition often exist between experimentalists and theoreticians. The difference, often undefined, is whether we refer to the FWHM of the field or of the intensity. An experimentalist will typically use the FWHM of the intensity, as this quantity is directly measurable – for example, the FWHM of a spatial profile of a laser beam, as measured on some detector. Theoreticians, in their confounding wisdom, prefer to use the FWHM of the field. It should also be noted that one might additionally refer to the 1/e width or radius of either the field of the intensity. This has, naturally, a separate definition. Here, then, we will extract the relationship between the field and intensity FWHM.

Using a Gaussian distribution, let us define a spatio-temporal intensity profile of a propagating laser pulse as:

$$E(r,z,t) = E_0 \frac{w_0}{w(z)} e^{-\frac{r^2}{w(z)^2}} e^{-\frac{(t-z/c)^2}{\tau_0^2}} e^{i(kz - w_0 t + \Psi(z))} \tag{A.1}$$

In this description E_0 is the peak intensity, $r = \sqrt{x^2 + y^2}$ is the radial distance in the transverse plane, w_0 is the beam waist (the minimum 1/e radius of the electric field), $w(z)$ is the transverse 1/e beam radius (of the electric field) at distance z, and τ_0 is the pulse duration defined at 1/e of the electric field amplitude. The first exponential term describes the radial envelope of the pulse, the second represents the temporal envelope, and the third contains the carrier wave information of the electric field. The final term $\Psi(z)$ is the Gouy phase shift, an additional phase of:

$$\Psi(z) = \arctan\left(\frac{z}{z_\mathrm{R}}\right) \tag{A.2}$$

Which arises when a Gaussian beam is focused. Here, z_R represents the Rayleigh length. Unlike a plane wave, a Gaussian beam accumulates an extra π radians of phase as it propagates through its focus (from $-\pi/2$ to $+\pi/2$). This effect is a direct consequence of diffraction.

Isolating the component describing the temporal envelope only, let us simplify at $z = 0$ and solve for t. We also assign the variables t_FWHM and τ_0 in to the respective field (E) and intensity (I) variables using $t_{E/2}$ and $t_{I/2}$, and τ_E and τ_I.

$$E(t) = E_0 e^{-\frac{t^2}{\tau_E^2}} \tag{A.3}$$

$$t = \tau_E \sqrt{\ln\left(E_0/E(t)\right)} \tag{A.4}$$

If we define the FWHM pulse duration of the field as $t_{E/2}$, we can identify that $t = t_{E/2}$ will occur at $E(t_{E/2})/E_0 = 0.5$ and can thus extract the relationship between the FWHM pulse duration and the 1/e pulse duration:

$$t_{E/2} = \tau_E \sqrt{\ln 2} \tag{A.5}$$

To now extract the relationship between the field and intensity quantities, we can find an equivalent expression for intensity simply by using $I \propto |E|^2$.

$$E(t)^2 = E_0^2 e^{-\frac{2t^2}{\tau_E^2}} \tag{A.6}$$

$$I(t) = I_0 e^{-\frac{2t^2}{\tau_E^2}} \tag{A.7}$$

$$\tag{A.8}$$

We may now extract the relationship between the two expressions. Thus, for both 1/e and FWHM variables, the field and intensity definitions can be related by:

$$t_E = \sqrt{2} t_I \tag{A.9}$$

The respective FWHM and 1/e variables taken from the intensity distribution will thus be smaller than those calculated from the field. This is shown graphically in Figure A.1.

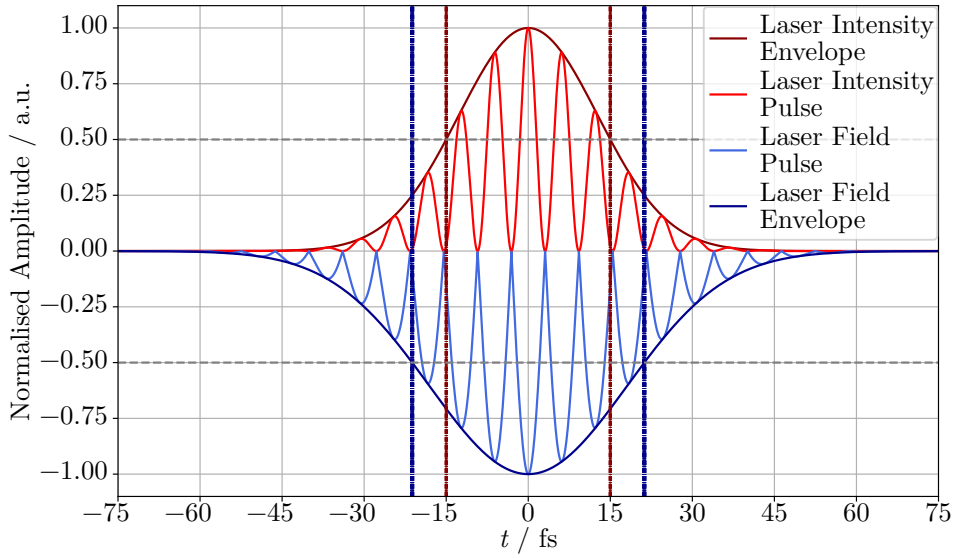


Figure A.1: The temporal intensity and (normalised) field Gaussian representations of the 3.7 µm central wavelength Chimera laser pulse. Here we see plainly that these two representations will produce different FWHM values of the pulse duration. The grey dashed line shows the half maximum positions, and the blue and red dot-dashed lines show the half-width positions for both the field and intensity envelopes. The intensity distribution returns a FWHM of 30 fs, while the field distribution returns a FWHM of 42 fs. This is in keeping with the derivation above.

B Appendix B

B.1 Relativistic Single Particle Motion

Following Gibbon [121], we derive the coordinates and equations of motion for an electron in a relativistic field (with $a_0 > 1$).

Expressing the electric and magnetic fields in terms of the vector potential $\mathbf{A}$, and assuming a vanishing scalar potential $\phi = 0$, we have $\mathbf{E} = -\frac{\partial}{\partial t}\mathbf{A}$ and $\mathbf{B} = \nabla \times \mathbf{A}$. Using the normalised quantities $\mathbf{a} = e\mathbf{A}/m_{\mathrm{e,i}}c$, $\boldsymbol{\beta} = \mathbf{v}/c$ and $\mathbf{p} = \mathbf{p}/m_{\mathrm{e,i}}c$, we choose the vector potential $\mathbf{a} = a_x\hat{\mathbf{x}} = a_0 \sin\theta\hat{\mathbf{x}}$, with $\theta = kz - \omega t$. This corresponds to a linearly polarised plane wave (in the x-direction) propagating along z. This allows the Lorentz force to be expressed as:

$$\frac{\mathrm{d}\mathbf{p}}{\mathrm{d}t} = -\frac{\partial \mathbf{a}}{\partial t} + [\boldsymbol{\beta} \times (\nabla \times \mathbf{a})] \tag{B.1}$$

$$= -\frac{\partial a_x}{\partial t}\hat{\mathbf{x}} + \left(-\beta_z \frac{\partial a_x}{\partial z}\hat{\mathbf{x}} + \beta_x \frac{\partial a_x}{\partial z}\hat{\mathbf{z}}\right) \tag{B.2}$$

Separating into transverse and longitudinal components gives:

$$\frac{\mathrm{d}p_x}{\mathrm{d}t} = -\frac{\partial a_x}{\partial t} - \beta_z \frac{\partial a_x}{\partial z} \tag{B.3}$$

$$\frac{\mathrm{d}p_z}{\mathrm{d}t} = \beta_x \frac{\partial a_x}{\partial z} \tag{B.4}$$

Using the convective derivative identity $\mathrm{d}/\mathrm{d}t = \partial/\partial t + \mathbf{v} \cdot \nabla$ we can write:

$$\frac{\mathrm{d}p_x}{\mathrm{d}t} = \frac{\mathrm{d}a_x}{\mathrm{d}t} \tag{B.5}$$

Integrating with respect to time, we find that the transverse momentum is equal simply to the sum of the potential and the initial momentum:

$$p_x = a_x + p_{x,0} \tag{B.6}$$

To determine the longitudinal momentum, we consider the rate of change of energy γ, derived from the work done by the fields, $\boldsymbol{\beta} \cdot \mathbf{F}$:

$$\frac{\mathrm{d}\gamma}{\mathrm{d}t} = -\beta_x \frac{\mathrm{d}a_x}{\mathrm{d}t} \tag{B.7}$$

Comparing this with the $\hat{\mathbf{z}}$-component of Equation (B.2), we find:

$$\frac{\mathrm{d}p_z}{\mathrm{d}t} - \frac{\mathrm{d}\gamma}{\mathrm{d}t} = \beta_x \left(\frac{\mathrm{d}a_x}{\mathrm{d}z} + \frac{\mathrm{d}a_x}{\mathrm{d}t} \right) = 0 \tag{B.8}$$

Integrating this with respect to time gives:

$$\gamma - p_z = \alpha \tag{B.9}$$

Where α is a constant of integration. Using the definition $\gamma = \sqrt{1 + |\mathbf{p}|^2}$, we can rearrage this to show:

$$p_z = \frac{1 + p_x^2 - \alpha^2}{2\alpha} \tag{B.10}$$

As we are assuming the electron starts from rest, we take $p_{x,0} = 0$ and $\alpha = 1$, to reveal the equations of motion as:

$$p_x = a_0 \sin\theta \tag{B.11}$$

$$p_z = \frac{a_0^2}{2} \sin^2\theta \tag{B.12}$$

At this stage, it becomes convenient to shift from working in terms of time t or position z, to using the phase variable $\theta = kz - \omega t$. In normalised units $k = \omega = 1$, so this becomes $\theta = z - t$.

Using the identity $p_z = \gamma \dot{z}$, and recalling $\gamma - p_z = \alpha = 1$, we start by computing the total derivative of θ:

$$\frac{\mathrm{d}\theta}{\mathrm{d}t} = \frac{\mathrm{d}z}{\mathrm{d}t} - 1 \tag{B.13}$$

$$= (\beta_z - 1) \tag{B.14}$$

$$= -\frac{\alpha}{\gamma} = -\frac{1}{\gamma} \tag{B.15}$$

We now rewrite the transverse momentum as

$$p_x = \gamma \frac{\mathrm{d}\theta}{\mathrm{d}t} \frac{\mathrm{d}x}{\mathrm{d}\theta} = -\frac{\mathrm{d}x}{\mathrm{d}\theta} \tag{B.16}$$

Integrating with respect to θ (and choosing the integration constant $C = 0$), the coordinates are then given by:

$$x = a_0 \cos\theta \tag{B.17}$$

$$z = \frac{a_0^4}{2}\left(\theta - \frac{1}{2}\sin 2\theta\right) \tag{B.18}$$

Bibliography

[1] J. C. Wilcke. “Ytterligare rön och försök om contraira electriciteterne vid laddningen och därtil hörande delar”. In: *Kongliga Svenska Vetenskapsakademiens Handlingar* 23, 1762, pp. 206–229, 245–266. (See p. 17)

[2] A. Volta. “On the electricity excited by the mere contact of conducting substances of different kinds. In a letter from Mr. Alexander Volta, F. R. S. Professor of Natural Philosophy in the University of Pavia, to the Rt. Hon. Sir Joseph Banks, Bart. K. B. P. R. S”. In: *Abstracts of the Papers Printed in the Philosophical Transactions of the Royal Society of London* 1, 1832. (See p. 17)

[3] E. Rutherford. “LXXIX. The scattering of α and β particles by matter and the structure of the atom”. In: *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* 21, 1911, pp. 669–688. (See p. 17)

[4] P. W. Higgs. “Broken Symmetries and the Masses of Gauge Bosons”. In: *Physical Review Letters* 13, 1964, pp. 508–509. (See p. 17)

[5] ATLAS Collaboration. “Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC”. In: *Physics Letters B* 716, 2012, pp. 1–29. (See p. 17)

[6] T. Shears. “The Standard Model”. In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 370, 2012, pp. 805–817. (See p. 17)

[7] W. Levin, H. Kooy, J. S. Loeffler, and T. F. Delaney. “Proton beam therapy”. In: *British journal of Cancer* 93, 2005, pp. 849–854. (See p. 17)

[8] O. Jacobson, D. O. Kiesewetter, and X. Chen. “Fluorine-18 Radiochemistry, Labeling Strategies and Synthetic Routes”. In: *Bioconjugate Chemistry* 26, 2015, pp. 1–18. (See p. 17)

[9] R. E. M. Hedges, W. D. Mook, and H. Margaritz. “Radiocarbon dating by accelerator mass spectrometry: some recent results and applications”. In: *Philosophical Transactions of the Royal Society of London. Series A, Mathematical and Physical Sciences* 323, 1987, pp. 57–73. (See p. 17)

[10] J. Williams. “Ion implantation of semiconductors”. In: *Materials Science and Engineering: A* 253, 1998, pp. 8–15. (See p. 17)

[11] A. Mousavi Khaneghah, M. Hashemi Moosavi, C. A. Oliveira, F. Vanin, and A. S. Sant'Ana. "Electron beam irradiation to reduce the mycotoxin and microbial contaminations of cereal-based products: An overview". In: *Food and Chemical Toxicology* 143, 2020, p. 111557. (See p. 17)

[12] J. Dik, K. Janssens, G. Van Der Snickt, L. van der Loeff, K. Rickers, and M. Cotte. "Visualization of a Lost Painting by Vincent van Gogh Using Synchrotron Radiation Based X-ray Fluorescence Elemental Mapping". In: *Analytical Chemistry* 80, 2008, pp. 6436–6442. (See p. 17)

[13] J. D. Cockcroft and E. T. S. Walton. "Experiments with high velocity positive ions.—(I) Further developments in the method of obtaining high velocity positive ions". In: *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character* 136, 1932, pp. 619–630. (See p. 17)

[14] G. Ising. "Prinzip einer Methode zur Herstellung von Kanalstrahlen hoher Voltzahl". In: *Arkiv för Matematik, Astronomi och Fysik* 18, 1924, pp. 1–4. (See p. 17)

[15] P. J. Bryant. *A brief history and review of accelerators*. CERN Report CERN-1994-001. CERN, 1994 (See p. 17)

[16] H. Padamsee. "50 years of success for SRF accelerators—a review". In: *Superconductor Science and Technology* 30, 2017, p. 053003. (See p. 18)

[17] The New York Times. *Developer of the Laser Calls It 'A Solution Seeking a Problem'; President of Korad Spends Spare Time Gardening and Fixing TV Sets*. May 1964 (See p. 19)

[18] T. H. Maiman. "Stimulated Optical Radiation in Ruby". In: *Nature* 187, 1960, pp. 493–494. (See p. 19)

[19] F. J. McClung and R. W. Hellwarth. "Giant Optical Pulsations from Ruby". In: *Journal of Applied Physics* 33, 1962, pp. 828–829. (See p. 19)

[20] L. E. Hargrove, R. L. Fork, and M. A. Pollack. "Locking of He-Ne Laser Modes Induced by Synchronous Intracavity Modulation". In: *Applied Physics Letters* 5, 1964, pp. 4–5. (See p. 19)

[21] J. DiDomenico M. "Small-Signal Analysis of Internal (Coupling-Type) Modulation of Lasers". In: *Journal of Applied Physics* 35, 1964, pp. 2870–2876. (See p. 19)

[22] A. Yariv. "Internal Modulation in Multimode Laser Oscillators". In: *Journal of Applied Physics* 36, 1965, pp. 388–391. (See p. 19)

[23] H. W. Mocker and R. J. Collins. "Mode Competition and Self-Locking Effects in a Q-Switched Ruby Laser". In: *Applied Physics Letters* 7, 1965, pp. 270–273. (See p. 19)

[24] A. J. DeMaria, D. A. Stetser, and H. Heynau. "Self Mode-Locking of Lasers with Saturable Absorbers". In: *Applied Physics Letters* 8, 1966, pp. 174–176. (See p. 19)

[25] D. Strickland and G. Mourou. "Compression of amplified chirped optical pulses". In: *Optics Communications* 56, 1985, pp. 219–221. (See pp. 19, 22, 39)

[26] M. Bellini, A. Bartoli, and T. W. Hänsch. “Two-photon Fourier spectroscopy with femtosecond light pulses”. In: *Optics Letters* 22, 1997, pp. 540–542. (See p. 20)

[27] T. Brabec and F. Krausz. “Intense few-cycle laser fields: Frontiers of nonlinear optics”. In: *Reviews of Modern Physics* 72, 2000, pp. 545–591. (See p. 20)

[28] T. Tajima and J. M. Dawson. “Laser Electron Accelerator”. In: *Physical Review Letters* 43, 1979, pp. 267–270. (See pp. 20, 81)

[29] S. P. Mangles, C. D. Murphy, Z. Najmudin, A. G. Thomas, J. L. Collier, A. E. Dangor, E. J. Divall, P. S. Foster, J. G. Gallacher, C. J. Hooker, D. A. Jaroszynski, A. J. Langley, W. B. Mori, P. A. Norreys, F. S. Tsung, R. Viskup, B. R. Walton, and K. Krushelnick. “Monoenergetic beams of relativistic electrons from intense laser–plasma interactions”. In: *Nature* 431, 2004, pp. 535–538. (See pp. 20, 22, 71, 81)

[30] C. G. Geddes, C. Toth, J. V. Tilborg, E. Esarey, C. B. Schroeder, D. Bruhwiler, C. Nieter, J. Cary, and W. P. Leemans. “High-quality electron beams from a laser wakefield accelerator using plasma-channel guiding”. In: *Nature* 431, 2004, pp. 538–541. (See pp. 20, 22, 71, 81)

[31] J. Faure, Y. Glinec, A. Pukhov, S. Klselev, S. Gordienko, E. Lefebvre, J. P. Rousseau, F. Burgy, and V. Malka. “A laser-plasma accelerator producing monoenergetic electron beams”. In: *Nature* 431, 2004, pp. 541–544. (See pp. 20, 22, 81)

[32] M. E. Fermann and I. Hartl. “Ultrafast Fiber Laser Technology”. In: *IEEE Journal of Selected Topics in Quantum Electronics* 15, 2009, pp. 191–206. (See pp. 20, 23)

[33] A. Khilo, S. J. Spector, M. E. Grein, A. H. Nejadmalayeri, C. W. Holzwarth, M. Y. Sander, M. S. Dahlem, M. Y. Peng, M. W. Geis, N. A. DiLello, J. U. Yoon, A. Motamedi, J. S. Orcutt, J. P. Wang, C. M. Sorace-Agaskar, M. A. Popović, J. Sun, G.-R. Zhou, H. Byun, J. Chen, J. L. Hoyt, H. I. Smith, R. J. Ram, M. Perrott, T. M. Lyszczarz, E. P. Ippen, and F. X. Kärtner. “Photonic ADC: overcoming the bottleneck of electronic jitter”. In: *Optics Express* 20, 2012, pp. 4454–4469. (See p. 20)

[34] S. H. Chung and E. Mazur. “Surgical applications of femtosecond lasers”. In: *Journal of Biophotonics* 2, 2009, pp. 557–572. (See p. 20)

[35] S. Siewert, E. Bohne, S. Grossmann, W. Schmidt, N. Grabow, T. Stahnke, R. Guthoff, M. Stiehm, and K.-P. Schmitz. “Femtosecond laser manufacturing technology for microstents as innovative medical devices to treat open-angle glaucoma”. In: *Current Directions in Biomedical Engineering* 8, 2022, pp. 592–595. (See p. 20)

[36] C. N. Danson, C. Haefner, J. Bromage, T. Butcher, J. C. F. Chanteloup, E. A. Chowdhury, A. Galvanauskas, L. A. Gizzi, J. Hein, D. I. Hillier, N. W. Hopps, Y. Kato, E. A. Khazanov, R. Kodama, G. Korn, R. Li, Y. Li, J. Limpert, J. Ma, C. H. Nam, D. Neely, D. Papadopoulos, R. R. Penman, L. Qian, J. J. Rocca, A. A. Shaykin, C. W. Siders, C. Spindloe, S. Szatmári, *et al.* “Petawatt and

exawatt class lasers worldwide". In: *High Power Laser Science and Engineering* 7, 2019, e54. (See pp. 20, 40)

[37] A. Dangor, A. Dymoke-Bradshaw, and A. Dyson. "Observation of relativistic plasma waves generated by the beat-wave with 1 µm lasers". In: *Physica Scripta* 1990, 1990, p. 107. (See p. 21)

[38] Y. Kitagawa, T. Matsumoto, T. Minamihata, K. Sawai, K. Matsuo, K. Mima, K. Nishihara, H. Azechi, K. A. Tanaka, H. Takabe, and S. Nakai. "Beat-wave excitation of plasma wave and observation of accelerated electrons". In: *Physical Review Letters* 68, 1992, pp. 48–51. (See p. 21)

[39] M. Everett, A. Lal, D. Gordon, C. Clayton, K. Marsh, and C. Joshi. "Trapped electron acceleration by a laser-driven relativistic plasma wave". In: *Nature* 368, 1994, pp. 527–529. (See p. 21)

[40] A. Modena, Z. Najmudin, A. E. Dangor, C. E. Clayton, K. A. Marsh, C. Joshi, V. Malka, C. B. Darrow, C. Danson, D. Neely, and F. N. Walsh. "Electron Acceleration from the Breaking of Relativistic Plasma Waves". In: *Nature* 377, 1995, pp. 606–608. (See p. 22)

[41] K. Nakajima, D. Fisher, T. Kawakubo, H. Nakanishi, A. Ogata, Y. Kato, Y. Kitagawa, R. Kodama, K. Mima, H. Shiraga, K. Suzuki, K. Yamakawa, T. Zhang, Y. Sakawa, T. Shoji, Y. Nishida, N. Yugami, M. Downer, and T. Tajima. "Observation of Ultrahigh Gradient Electron Acceleration by a Self-Modulated Intense Short Laser Pulse". In: *Physical Review Letters* 74, 1995, pp. 4428–4431. (See p. 22)

[42] D. Gordon, K. C. Tzeng, C. E. Clayton, A. E. Dangor, V. Malka, K. A. Marsh, A. Modena, W. B. Mori, P. Muggli, Z. Najmudin, D. Neely, C. Danson, and C. Joshi. "Observation of Electron Energies Beyond the Linear Dephasing Limit from a Laser-Excited Relativistic Plasma Wave". In: *Physical Review Letters* 80, 1998, pp. 2133–2136. (See p. 22)

[43] F. Amiranoff, S. Baton, D. Bernard, B. Cros, D. Descamps, F. Dorchies, F. Jacquet, V. Malka, J. R. Marquès, G. Matthieussent, P. Miné, A. Modena, P. Mora, J. Morillo, and Z. Najmudin. "Observation of Laser Wakefield Acceleration of Electrons". In: *Physical Review Letters* 81, 1998, pp. 995–998. (See p. 22)

[44] E. L. Clark, K. Krushelnick, J. R. Davies, M. Zepf, M. Tatarakis, F. N. Beg, A. Machacek, P. A. Norreys, M. I. K. Santala, I. Watts, and A. E. Dangor. "Measurements of Energetic Proton Transport through Magnetized Plasma from Intense Laser Interactions with Solids". In: *Physical Review Letters* 84, 2000, pp. 670–673. (See p. 22)

[45] E. L. Clark, K. Krushelnick, M. Zepf, F. N. Beg, M. Tatarakis, A. Machacek, M. I. K. Santala, I. Watts, P. A. Norreys, and A. E. Dangor. "Energetic Heavy-Ion and Proton Generation from Ultraintense Laser-Plasma Interactions with Solids". In: *Physical Review Letters* 85, 2000, pp. 1654–1657. (See p. 22)

[46] S. P. Hatchett, C. G. Brown, T. E. Cowan, E. A. Henry, J. S. Johnson, M. H. Key, J. A. Koch, A. B. Langdon, B. F. Lasinski, R. W. Lee, A. J. Mackinnon, D. M. Pennington, M. D. Perry, T. W. Phillips, M. Roth, T. C. Sangster, M. S. Singh, R. A. Snavely, M. A. Stoyer, S. C. Wilks, and K. Yasuike. "Electron, photon, and ion beams from the relativistic interaction of Petawatt laser pulses with solid targets". In: *Physics of Plasmas* 7, 2000, pp. 2076–2082. (See p. 22)

[47] R. A. Snavely, M. H. Key, S. P. Hatchett, T. E. Cowan, M. Roth, T. W. Phillips, M. A. Stoyer, E. A. Henry, T. C. Sangster, M. S. Singh, S. C. Wilks, A. MacKinnon, A. Offenberger, D. M. Pennington, K. Yasuike, A. B. Langdon, B. F. Lasinski, J. Johnson, M. D. Perry, and E. M. Campbell. "Intense High-Energy Proton Beams from Petawatt-Laser Irradiation of Solids". In: *Physical Review Letters* 85, 2000, pp. 2945–2948. (See p. 22)

[48] L. M. Chen, H. Kotaki, K. Nakajima, J. Koga, S. V. Bulanov, T. Tajima, Y. Q. Gu, H. S. Peng, X. X. Wang, T. S. Wen, H. J. Liu, C. Y. Jiao, C. G. Zhang, X. J. Huang, Y. Guo, K. N. Zhou, J. F. Hua, W. M. An, C. X. Tang, and Y. Z. Lin. "Self-guiding of 100 TW femtosecond laser pulses in centimeter-scale underdense plasma". In: *Physics of Plasmas* 14, 2007. (See p. 22)

[49] A. Zigler, Y. Ehrlich, C. Cohen, J. Krall, and P. Sprangle. "Optical guiding of high-intensity laser pulses in a long plasma channel formed by a slow capillary discharge". In: *Journal of the Optical Society of America B* 13, 1996, pp. 68–71. (See p. 22)

[50] Y. Ehrlich, C. Cohen, D. Kaganovich, A. Zigler, R. F. Hubbard, P. Sprangle, and E. Esarey. "Guiding and damping of high-intensity laser pulses in long plasma channels". In: *Journal of the Optical Society of America B* 15, 1998, pp. 2416–2423. (See p. 22)

[51] A. Butler, D. J. Spence, and S. M. Hooker. "Guiding of High-Intensity Laser Pulses with a Hydrogen-Filled Capillary Discharge Waveguide". In: *Physical Review Letters* 89, 2002, p. 185003. (See p. 22)

[52] D. J. Spence, A. Butler, and S. M. Hooker. "Gas-filled capillary discharge waveguides". In: *Journal of the Optical Society of America B* 20, 2003, pp. 138–151. (See p. 22)

[53] W. P. Leemans, B. Nagler, A. J. Gonsalves, C. Tóth, K. Nakamura, C. G. Geddes, E. Esarey, C. B. Schroeder, and S. M. Hooker. "GeV electron beams from a centimetre-scale accelerator". In: *Nature Physics* 2, 2006, pp. 696–699. (See p. 22)

[54] W. P. Leemans, A. J. Gonsalves, H.-S. Mao, K. Nakamura, C. Benedetti, C. B. Schroeder, C. Tóth, J. Daniels, D. E. Mittelberger, S. S. Bulanov, J.-L. Vay, C. G. R. Geddes, and E. Esarey. "Multi-GeV Electron Beams from Capillary-Discharge-Guided Subpetawatt Laser Pulses in the Self-Trapping Regime". In: *Physical Review Letters* 113, 2014, p. 245002. (See p. 22)

[55] A. J. Gonsalves, K. Nakamura, J. Daniels, C. Benedetti, C. Pieronek, T. C. H. de Raadt, S. Steinke, J. H. Bin, S. S. Bulanov, J. van Tilborg, C. G. R. Geddes, C. B. Schroeder, C. Tóth, E. Esarey, K. Swanson, L. Fan-Chiang, G. Bagdasarov, N. Bobrova, V. Gasilov, G. Korn, P. Sasorov, and W. P. Leemans. "Petawatt Laser Guiding and Electron Beam Acceleration to 8 GeV in a Laser-Heated Capillary Discharge Waveguide". In: *Physical Review Letters* 122, 2019, p. 084801. (See p. 22)

[56] C. G. Durfee and H. M. Milchberg. "Light pipe for high intensity laser pulses". In: *Physical Review Letters* 71, 1993, pp. 2409–2412. (See p. 22)

[57] H. M. Milchberg, T. R. Clark, I. Durfee C. G., T. M. Antonsen, and P. Mora. "Development and applications of a plasma waveguide for intense laser pulses". In: *Physics of Plasmas* 3, 1996, pp. 2149–2155. (See p. 22)

[58] R. J. Shalloo, C. Arran, L. Corner, J. Holloway, J. Jonnerby, R. Walczak, H. M. Milchberg, and S. M. Hooker. "Hydrodynamic optical-field-ionized plasma channels". In: *Physical Review E* 97, 2018, p. 053203. (See p. 22)

[59] A. Picksley, A. Alejo, R. J. Shalloo, C. Arran, A. von Boetticher, L. Corner, J. A. Holloway, J. Jonnerby, O. Jakobsson, C. Thornton, R. Walczak, and S. M. Hooker. "Meter-scale conditioned hydrodynamic optical-field-ionized plasma channels". In: *Physical Review E* 102, 2020, p. 053201. (See p. 22)

[60] B. Miao, L. Feder, J. E. Shrock, A. Goffin, and H. M. Milchberg. "Optical Guiding in Meter-Scale Plasma Waveguides". In: *Physical Review Letters* 125, 2020, p. 074801. (See p. 22)

[61] B. Miao, J. E. Shrock, L. Feder, R. C. Hollinger, J. Morrison, R. Nedbailo, A. Picksley, H. Song, S. Wang, J. J. Rocca, and H. M. Milchberg. "Multi-GeV Electron Bunches from an All-Optical Laser Wakefield Accelerator". In: *Physical Review X* 12, 2022, p. 031038. (See p. 22)

[62] A. Picksley, J. Stackhouse, C. Benedetti, K. Nakamura, H. E. Tsai, R. Li, B. Miao, J. E. Shrock, E. Rockafellow, H. M. Milchberg, C. B. Schroeder, J. van Tilborg, E. Esarey, C. G. R. Geddes, and A. J. Gonsalves. "Matched Guiding and Controlled Injection in Dark-Current-Free, 10-GeV-Class, Channel-Guided Laser-Plasma Accelerators". In: *Physical Review Letters* 133, 2024, p. 255001. (See p. 22)

[63] S. Steinke, J. van Tilborg, C. Benedetti, C. G. Geddes, C. B. Schroeder, J. Daniels, K. K. Swanson, A. J. Gonsalves, K. Nakamura, N. H. Matlis, B. H. Shaw, E. Esarey, and W. P. Leemans. "Multistage coupling of independent laser-plasma accelerators". In: *Nature* 530, 2016, pp. 190–3. (See p. 22)

[64] G. Mourou, B. Brocklesby, T. Tajima, and J. Limpert. "The future is fibre accelerators". In: *Nature Photonics* 7, 2013, pp. 258–261. (See p. 22)

[65] J. J. van de Wetering, S. M. Hooker, and R. Walczak. “Multi-GeV wakefield acceleration in a plasma-modulated plasma accelerator”. In: *Physical Review E* 109, 2024, p. 025206. (See p. 22)

[66] J. Faure, D. Gustas, D. Guénot, A. Vernier, F. Böhle, M. Ouillé, S. Haessler, R. Lopez-Martens, and A. Lifschitz. “A review of recent progress on laser-plasma acceleration at kHz repetition rate”. In: *Plasma Physics and Controlled Fusion* 61, 2018, p. 014012. (See pp. 22, 59, 138, 139, 151, 164)

[67] C. G. R. Geddes, K. Nakamura, G. R. Plateau, C. Toth, E. Cormier-Michel, E. Esarey, C. B. Schroeder, J. R. Cary, and W. P. Leemans. “Plasma-Density-Gradient Injection of Low Absolute-Momentum-Spread Electron Bunches”. In: *Physical Review Letters* 100, 2008, p. 215004. (See pp. 22, 68)

[68] K. Schmid, A. Buck, C. M. S. Sears, J. M. Mikhailova, R. Tautz, D. Herrmann, M. Geissler, F. Krausz, and L. Veisz. “Density-transition based electron injector for laser driven wakefield accelerators”. In: *Physical Review Special Topics - Accelerators and Beams* 13, 2010, p. 091301. (See p. 22)

[69] J. Faure, C. Rechatin, A. Norlin, A. Lifschitz, Y. Glinec, and V. Malka. “Controlled injection and acceleration of electrons in plasma wakefields by colliding laser pulses”. In: *Nature* 444, 2006, pp. 737–739. (See p. 22)

[70] A. Pak, K. A. Marsh, S. F. Martins, W. Lu, W. B. Mori, and C. Joshi. “Injection and Trapping of Tunnel-Ionized Electrons into Laser-Produced Wakes”. In: *Physical Review Letters* 104, 2010, p. 025003. (See pp. 22, 68)

[71] C. McGuffey, A. G. R. Thomas, W. Schumaker, T. Matsuoka, V. Chvykov, F. J. Dollar, G. Kalintchenko, V. Yanovsky, A. Maksimchuk, K. Krushelnick, V. Y. Bychenkov, I. V. Glazyrin, and A. V. Karpeev. “Ionization Induced Trapping in a Laser Wakefield Accelerator”. In: *Physical Review Letters* 104, 2010, p. 025004. (See pp. 22, 68, 69)

[72] P. Sprangle, B. Hafizi, J. R. Peñano, R. F. Hubbard, A. Ting, C. I. Moore, D. F. Gordon, A. Zigler, D. Kaganovich, and T. M. Antonsen. “Wakefield generation and GeV acceleration in tapered plasma channels”. In: *Physical Review E* 63, 2001, p. 056405. (See p. 22)

[73] A. Pukhov and I. Kostyukov. “Control of laser-wakefield acceleration by the plasma-density profile”. In: *Physical Review E* 77, 2008, p. 025401. (See p. 22)

[74] W. Lu, M. Tzoufras, C. Joshi, F. S. Tsung, W. B. Mori, J. Vieira, R. A. Fonseca, and L. O. Silva. “Generating multi-GeV electron bunches using single stage laser wakefield acceleration in a 3D nonlinear regime”. In: *Physical Review Special Topics - Accelerators and Beams* 10, 2007, p. 061301. (See pp. 22, 54, 66, 70–72, 81, 140, 144, 155, 164, 165, 170, 171, 174, 183)

[75] A. M. Larson. *Multiphoton microscopy*. 2011 (See p. 23)

[76] E. E. Hoover and J. A. Squier. “Advances in multiphoton microscopy technology”. In: *Nature Photonics* 7, 2013, pp. 93–101. (See p. 23)

[77] Indirect Drive ICF Collaboration. “Lawson Criterion for Ignition Exceeded in an Inertial Fusion Experiment”. In: *Physical Review Letters* 129, 2022, p. 075001. (See p. 23)

[78] H. Takabe and Y. Kuramitsu. “Recent progress of laboratory astrophysics with intense lasers”. In: *High Power Laser Science and Engineering* 9, 2021, e49. (See p. 23)

[79] M. G. Allen. “Diode laser absorption sensors for gas-dynamic and combustion flows”. In: *Measurement Science and technology* 9, 1998, p. 545. (See p. 23)

[80] A. Schliesser, N. Picqué, and T. W. Hänsch. “Mid-infrared frequency combs”. In: *Nature Photonics* 6, 2012, pp. 440–449. (See p. 23)

[81] H. Clark, L. Corner, W. Denzer, G. Hancock, A. Hutchinson, M. Islam, R. Peverall, and G. Ritchie. “Difference frequency generation in periodically poled lithium niobate and its use in the detection of atmospheric methane”. In: *Chemical Physics Letters* 399, 2004, pp. 102–108. (See p. 23)

[82] K. C. Cossel, E. M. Waxman, I. A. Finneran, G. A. Blake, J. Ye, and N. R. Newbury. “Gas-phase broadband spectroscopy using active sources: progress, status, and applications”. In: *Journal of the Optical Society of America B* 34, 2017, pp. 104–129. (See p. 23)

[83] U. Willer, M. Saraji, A. Khorsandi, P. Geiser, and W. Schade. “Near- and mid-infrared laser monitoring of industrial processes, environment and security applications”. In: *Optics and Lasers in Engineering* 44, 2006, pp. 699–710. (See p. 23)

[84] C. Bauer, A. Sharma, U. Willer, J. Burgmeier, B. Braunschweig, W. Schade, S. Blaser, L. Hvozdara, A. Müller, and G. Holl. “Potentials and limits of mid-infrared laser spectroscopy for the detection of explosives”. In: *Applied Physics B* 92, 2008, pp. 327–333. (See p. 23)

[85] B. Jean and T. Bende. In: “Mid-IR Laser Applications in Medicine”. I. T. Sorokina and K. L. Vodopyanov, eds. Vol. 89, pp. 530–565. Springer Science, 2003. (See p. 23)

[86] L. Flannigan, L. Yoell, and C.-q. Xu. “Mid-wave and long-wave infrared transmitters and detectors for optical satellite communications — a review”. In: *Journal of Optics* 24, 2022, p. 043002. (See p. 23)

[87] A. D. Shiner, C. Trallero-Herrero, N. Kajumba, H.-C. Bandulet, D. Comtois, F. Légaré, M. Giguère, J.-C. Kieffer, P. B. Corkum, and D. M. Villeneuve. “Wavelength Scaling of High Harmonic Generation Efficiency”. In: *Physical Review Letters* 103, 2009, p. 073902. (See p. 23)

[88] G. New. *Introduction to Nonlinear Optics*. Cambridge University Press, 2011. (See pp. 26, 32, 37, 45, 51)

[89] R. W. Boyd. *Nonlinear optics*. Fourth edition. Elsevier, 2020. (See p. 29)

[90] P. A. Franken, A. E. Hill, C. W. Peters, and G. Weinreich. “Generation of Optical Harmonics”. In: *Physical Review Letters* 7, 1961, pp. 118–119. (See p. 31)

[91] M. Göppert-Mayer. “Über Elementarakte mit zwei Quantensprüngen”. In: *Annalen der Physik* 401, 1931, pp. 273–294. (See p. 31)

[92] P. D. Maker, R. W. Terhune, M. Nisenoff, and C. M. Savage. “Effects of Dispersion and Focusing on the Production of Optical Harmonics”. In: *Physical Review Letters* 8, 1962, pp. 21–22. (See p. 31)

[93] A. Aghigh, S. Bancelin, M. Rivard, M. Pinsard, H. Ibrahim, and F. Légaré. “Second harmonic generation microscopy: a powerful tool for bio-imaging”. In: *Biophysical Reviews* 15, 2023, pp. 43–70. (See p. 31)

[94] E. Bartholini. *Experimenta Crystalli Islandici disdiaclastici, quibus mira et insolita refractio detegitur*. Landmarks of Science. Sumptibus D. Paulli, 1669. (See p. 33)

[95] J. A. Armstrong, N. Bloembergen, J. Ducuing, and P. S. Pershan. “Interactions between Light Waves in a Nonlinear Dielectric”. In: *Physical Review* 127, 1962, pp. 1918–1939. (See p. 33)

[96] K. Nassau and H. J. Levinstein. “Ferroelectric Behavior of Lithium Niobate”. In: *Applied Physics Letters* 7, 1965, pp. 69–70. (See p. 33)

[97] J. A. Giordmaine and R. C. Miller. “Tunable coherent parametric oscillation in $LiNbO_3$ at optical frequencies”. In: *Physical Review Letters* 14, 1965, p. 973. (See p. 33)

[98] J. M. Manley and H. E. Rowe. “Some General Properties of Nonlinear Elements-Part I. General Energy Relations”. In: *Proceedings of the IRE* 44, 1956, pp. 904–913. (See p. 33)

[99] G. Cerullo and S. D. Silvestri. “Ultrafast optical parametric amplifiers”. In: *Review of Scientific Instruments* 74, 2003, pp. 1–18. (See p. 37)

[100] P. K. Bates. “Optical Parametric Chirped Pulse Amplification in the Few-Cycle Regime”. PhD thesis. Imperial College London, 2008. (See pp. 37, 107, 108)

[101] Y. Chen, Y. Li, W. Li, X. Guo, and Y. Leng. “Generation of high beam quality, high-energy and broadband tunable mid-infrared pulse from a KTA optical parametric amplifier”. In: *Optics Communications* 365, 2016, pp. 7–13. (See p. 37)

[102] A. Dubietis, R. Butkus, and A. P. Piskarskas. “Trends in chirped pulse optical parametric amplification”. In: *IEEE Journal on Selected Topics in Quantum Electronics* 12, 2006, pp. 163–172. (See p. 38)

[103] D. A. Kleinman. “Theory of Optical Parametric Noise”. In: *Physical Review* 174, 1968, pp. 1027–1041. (See p. 38)

[104] N. Stuart, D. Bigourd, R. Hill, T. Robinson, K. Mecseki, S. Patankar, G. New, and R. Smith. “Direct fluorescence characterisation of a picosecond seeded optical parametric amplifier”. In: *Optics Communications* 336, 2015, pp. 319–325. (See p. 38)

[105] S. De Silvestri, P. Laporta, and O. Svelto. “The role of cavity dispersion in CW mode-locked dye lasers”. In: *IEEE Journal of Quantum Electronics* 20, 1984, pp. 533–539. (See p. 41)

[106] E. Treacy. “Optical pulse compression with diffraction gratings”. In: *IEEE Journal of Quantum Electronics* 5, 1969, pp. 454–458. (See pp. 44, 45)

[107] A. Öffner. “Unit Power Imaging Catoptric Anastigmat”. Patent US3748015A. U. S. Patent and T. Office. 1973 (See p. 44)

[108] O. Martinez. “3000 times grating compressor with positive group velocity dispersion: Application to fiber compensation in 1.3-1.6 µm region”. In: *IEEE Journal of Quantum Electronics* 23, 1987, pp. 59–64. (See p. 44)

[109] A. Dubietis, G. Jonušauskas, and A. Piskarskas. “Powerful femtosecond pulse generation by chirped and stretched pulse parametric amplification in BBO crystal”. In: *Optics Communications* 88, 1992, pp. 437–440. (See p. 47)

[110] I. Ross, P. Matousek, M. Towrie, A. Langley, and J. Collier. “The prospects for ultrashort pulse duration and ultrahigh intensity using optical parametric chirped pulse amplifiers”. In: *Optics Communications* 144, 1997, pp. 125–133. (See p. 47)

[111] C. Danson, P. Brummitt, R. Clarke, J. Collier, B. Fell, A. Frackiewicz, S. Hawkes, C. Hernandez-Gomez, P. Holligan, M. Hutchinson, and et al. “Vulcan petawatt: Design, operation and interactions at 5×10^{20} Wcm^{-2}”. In: *Laser and Particle Beams* 23, 2005, pp. 87–93. (See p. 47)

[112] L. Waxer, D. Maywar, J. Kelly, T. Kessler, B. Kruschwitz, S. Loucks, R. McCrory, D. Meyerhofer, S. Morse, C. Stoeckl, and J. Zuegel. “High-Energy Petawatt Capability for the Omega Laser”. In: *Optics and Photonics News* 16, 2005, pp. 30–36. (See p. 47)

[113] P. L. Kelley. “Self-Focusing of Optical Beams”. In: *Physical Review Letters* 15, 1965, pp. 1005–1008. (See p. 50)

[114] R. R. Alfano and S. L. Shapiro. “Observation of Self-Phase Modulation and Small-Scale Filaments in Crystals and Glasses”. In: *Physical Review Letters* 24, 1970, pp. 592–594. (See p. 50)

[115] S. Hooker and C. Webb. *Laser Physics*. Oxford Master Series in Atomic, Optical and Laser Physics. Oxford University Press, 2010. (See pp. 51, 75)

[116] N. Stuart. “Generation of High-Contrast, Terawatt to Petawatt OPCPA Laser Pulses”. PhD thesis. Imperial College London, 2017. (See pp. 51, 77–79)

[117] K. Mecseki. “An Ultrafast Optical Parametric Laser for Driving High Energy Density Science”. PhD thesis. Imperial College London, 2013. (See pp. 51, 89, 95, 96, 99, 104)

[118] D. E. Spence, P. N. Kean, and W. Sibbett. “60-fsec pulse generation from a self-mode-locked Ti:sapphire laser”. In: *Optics Letters* 16, 1991, pp. 42–44. (See p. 51)

[119] S. Yefet and A. Pe'er. "A Review of Cavity Design for Kerr Lens Mode-Locked Solid-State Lasers". In: *Applied Sciences* 3, 2013, pp. 694–724. (See pp. 51–53, 94, 95)

[120] U. Morgner, F. X. Kärtner, S. H. Cho, Y. Chen, H. A. Haus, J. G. Fujimoto, E. P. Ippen, V. Scheuer, G. Angelow, and T. Tschudi. "Sub-two-cycle pulses from a Kerr-lens mode-locked Ti:sapphire laser". In: *Optics Letters* 24, 1999, pp. 411–413. (See p. 51)

[121] P. Gibbon. *Short Pulse Laser Interactions with Matter: An Introduction*. Imperial College Press, 2005. (See pp. 56, 189)

[122] A. Thomas. "Studies of laser propagation and mono-energetic electron beam injection in laser wakefield accelerators". PhD thesis. Imperial College London, 2007. (See p. 56)

[123] P. Mora and T. M. Antonsen Jr. "Kinetic modeling of intense, short laser pulses propagating in tenuous plasmas". In: *Physics of Plasmas* 4, 1997, pp. 217–229. (See p. 58)

[124] A. Zhidkov, T. Fujii, and K. Nemoto. "Electron self-injection during interaction of tightly focused few-cycle laser pulses with underdense plasma". In: *Physical Review E* 78, 2008, p. 036406. (See p. 59)

[125] K. Schmid, L. Veisz, F. Tavella, S. Benavides, R. Tautz, D. Herrmann, A. Buck, B. Hidding, A. Marcinkevicius, U. Schramm, M. Geissler, J. Meyer-ter-Vehn, D. Habs, and F. Krausz. "Few-Cycle Laser-Driven Electron Acceleration". In: *Physical Review Letters* 102, 2009, p. 124801. (See p. 59)

[126] E. N. Nerush and I. Y. Kostyukov. "Carrier-Envelope Phase Effects in Plasma-Based Electron Acceleration with Few-Cycle Laser Pulses". In: *Physical Review Letters* 103, 2009, p. 035001. (See p. 59)

[127] J. Huijts, I. A. Andriyash, L. Rovige, A. Vernier, and J. Faure. "Identifying observable carrier-envelope phase effects in laser wakefield acceleration with near-single-cycle pulses". In: *Physics of Plasmas* 28, 2021, p. 043101. (See pp. 59, 151, 164)

[128] P. M. Woodward. "A method of calculating the field over a plane aperture required to produce a given polar diagram". In: *Journal of the Institution of Electrical Engineers* 93, 1947, pp. 1554–1558. (See p. 60)

[129] J. D. Lawson. "Lasers and Accelerators". In: *IEEE Transactions on Nuclear Science* 26, 1979, pp. 4217–4219. (See p. 60)

[130] E. Esarey, C. B. Schroeder, and W. P. Leemans. "Physics of laser-driven plasma-based electron accelerators". In: *Reviews of Modern Physics* 81, 2009, pp. 1229–1285. (See p. 60)

[131] F. F. Chen. *Introduction to plasma physics and controlled fusion*. 2nd ed. Vol. 1. Springer, 1984. (See p. 61)

[132] W. Mori. "The physics of the nonlinear optics of plasmas at relativistic intensities for short-pulse lasers". In: *IEEE Journal of Quantum Electronics* 33, 1997, pp. 1942–1953. (See pp. 63, 64)

[133] G.-Z. Sun, E. Ott, Y. C. Lee, and P. Guzdar. "Self-focusing of short intense pulses in plasmas". In: *The Physics of Fluids* 30, 1987, pp. 526–532. (See p. 64)

[134] A. G. R. Thomas, Z. Najmudin, S. P. D. Mangles, C. D. Murphy, A. E. Dangor, C. Kamperidis, K. L. Lancaster, W. B. Mori, P. A. Norreys, W. Rozmus, and K. Krushelnick. "Effect of Laser-Focusing Conditions on Propagation and Monoenergetic Electron Production in Laser-Wakefield Accelerators". In: *Physical Review Letters* 98, 2007, p. 095004. (See pp. 64, 70)

[135] S. C. Wilks, J. M. Dawson, W. B. Mori, T. Katsouleas, and M. E. Jones. "Photon accelerator". In: *Physical Review Letters* 62, 1989, pp. 2600–2603. (See pp. 66, 147)

[136] C. D. Murphy, R. Trines, J. Vieira, A. J. W. Reitsma, R. Bingham, J. L. Collier, E. J. Divall, P. S. Foster, C. J. Hooker, A. J. Langley, P. A. Norreys, R. A. Fonseca, F. Fiuza, L. O. Silva, J. T. Mendonça, W. B. Mori, J. G. Gallacher, R. Viskup, D. A. Jaroszynski, S. P. D. Mangles, A. G. R. Thomas, K. Krushelnick, and Z. Najmudin. "Evidence of photon acceleration by laser wake fields". In: *Physics of Plasmas* 13, 2006, p. 033108. (See pp. 66, 147)

[137] J. M. Dawson. "Nonlinear Electron Oscillations in a Cold Plasma". In: *The Physical Review* 113, 1959, pp. 383–387. (See p. 67)

[138] P. Sprangle, E. Esarey, and A. Ting. "Nonlinear interaction of intense laser pulses in plasmas". In: *Physical Review A* 41, 1990, pp. 4463–4469. (See p. 67)

[139] E. Esarey, P. Sprangle, J. Krall, and A. Ting. "Overview of plasma-based accelerator concepts". In: *IEEE Transactions on plasma science* 24, 2002, pp. 252–288. (See p. 67)

[140] E. Esarey and M. Pilloff. "Trapping and acceleration in nonlinear plasma waves". In: *Physics of Plasmas* 2, 1995, pp. 1432–1436. (See p. 67)

[141] V. Berezhiani and I. Murusidze. "Relativistic wake-field generation by an intense laser pulse in a plasma". In: *Physics Letters A* 148, 1990, pp. 338–340. (See p. 68)

[142] C. A. Coverdale, C. B. Darrow, C. D. Decker, W. B. Mori, K.-C. Tzeng, K. A. Marsh, C. E. Clayton, and C. Joshi. "Propagation of Intense Subpicosecond Laser Pulses through Underdense Plasmas". In: *Physical Review Letters* 74, 1995, pp. 4659–4662. (See p. 68)

[143] A. G. R. Thomas. "Scalings for radiation from plasma bubbles". In: *Physics of Plasmas* 17, 2010. (See pp. 68, 150, 170, 171)

[144] S. P. Mangles, G. Genoud, M. S. Bloom, M. Burza, Z. Najmudin, A. Persson, K. Svensson, A. G. Thomas, and C. G. Wahlström. "Self-injection threshold in self-guided laser wakefield accelerators". In: *Physical Review Special Topics - Accelerators and Beams* 15, 2012. (See pp. 68, 171–173)

[145] T. P. Rowlands-Rees, C. Kamperidis, S. Kneip, A. J. Gonsalves, S. P. D. Mangles, J. G. Gallacher, E. Brunetti, T. Ibbotson, C. D. Murphy, P. S. Foster, M. J. V. Streeter, F. Budde, P. A. Norreys, D. A. Jaroszynski, K. Krushelnick, Z. Najmudin, and S. M. Hooker. "Laser-Driven Acceleration of Electrons in a Partially Ionized Plasma Channel". In: *Physical Review Letters* 100, 2008, p. 105005. (See p. 69)

[146] D. L. Bruhwiler, D. A. Dimitrov, J. R. Cary, E. Esarey, W. Leemans, and R. E. Giacone. "Particle-in-cell simulations of tunneling ionization effects in plasma-based accelerators". In: *Physics of Plasmas* 10, 2003, pp. 2022–2030. (See p. 69)

[147] A. Pukhov and J. Meyer-ter-Vehn. "Laser wake field acceleration: the highly non-linear broken-wave regime". In: *Applied Physics B* 74, 2002, pp. 355–361. (See p. 69)

[148] S. P. D. Mangles, A. G. R. Thomas, C. Bellei, A. E. Dangor, C. Kamperidis, S. Kneip, S. R. Nagel, L. Willingale, and Z. Najmudin. "Self-Guided Wakefield Experiments Driven by Petawatt-Class Ultrashort Laser Pulses". In: *IEEE Transactions on Plasma Science* 36, 2008, pp. 1715–1721. (See p. 71)

[149] C. Decker, W. Mori, K.-C. Tzeng, and T. Katsouleas. "The evolution of ultra-intense, short-pulse lasers in underdense plasmas". In: *Physics of Plasmas* 3, 1996, pp. 2047–2056. (See pp. 72, 162)

[150] M. S. Bloom, M. J. V. Streeter, S. Kneip, R. A. Bendoyro, O. Cheklov, J. M. Cole, A. Döpp, C. J. Hooker, J. Holloway, J. Jiang, N. C. Lopes, H. Nakamura, P. A. Norreys, P. P. Rajeev, D. R. Symes, J. Schreiber, J. C. Wood, M. Wing, Z. Najmudin, and S. P. D. Mangles. "Bright X-ray radiation from plasma bubbles in an evolving laser wakefield accelerator". In: *Physical Review Accelerators and Beams* 23, 2020, p. 061301. (See pp. 73, 144)

[151] A. Pukhov, Z.-M. Sheng, and J. Meyer-ter-Vehn. "Particle acceleration in relativistic laser channels". In: *Physics of Plasmas* 6, 1999, pp. 2847–2854. (See p. 74)

[152] J. Shaw, N. Lemos, L. Amorim, N. Vafaei-Najafabadi, K. Marsh, F. Tsung, W. Mori, and C. Joshi. "Role of Direct Laser Acceleration of Electrons in a Laser Wakefield Accelerator with Ionization Injection". In: *Physical Review Letters* 118, 2017, p. 064801. (See pp. 74, 154)

[153] W. Rudolph, M. Sheik-Bahae, A. Bernstein, and L. Lester. "Femtosecond autocorrelation measurements based on two-photon photoconductivity in ZnSe". In: *Optics Letters* 22, 1997, pp. 313–315. (See p. 75)

[154] J. K. Ranka, A. L. Gaeta, A. Baltuska, M. S. Pshenichnikov, and D. A. Wiersma. "Autocorrelation measurement of 6-fs pulses based on the two-photon-induced photocurrent in a GaAsP photodiode". In: *Optics Letters* 22, 1997, pp. 1344–1346. (See p. 75)

[155] V. Baraulya, S. Kobtsev, and A. Korablev. “Use of AL307 light-emitting diodes as photodetectors for diagnostics of femtosecond light pulses”. In: *Technical Physics Letters* 24, 1998, pp. 28–29. (See p. 75)

[156] D. Kane and R. Trebino. “Characterization of arbitrary femtosecond pulses using frequency-resolved optical gating”. In: *IEEE Journal of Quantum Electronics* 29, 1993, pp. 571–579. (See p. 76)

[157] C. Iaconis and I. A. Walmsley. “Spectral phase interferometry for direct electric-field reconstruction of ultrashort optical pulses”. In: *Optics Letters* 23, 1998, pp. 792–794. (See p. 76)

[158] M. Miranda, T. Fordell, C. Arnold, A. L’Huillier, and H. Crespo. “Simultaneous compression and characterization of ultrashort laser pulses using chirped mirrors and glass wedges”. In: *Optics Express* 20, 2012, pp. 688–697. (See p. 76)

[159] M. Miranda, C. L. Arnold, T. Fordell, F. Silva, B. Alonso, R. Weigand, A. L’Huillier, and H. Crespo. “Characterization of broadband few-cycle laser pulses with the d-scan technique”. In: *Optics Express* 20, 2012, pp. 18732–18743. (See p. 76)

[160] S. Luan, M. H. R. Hutchinson, R. A. Smith, and F. Zhou. “High dynamic range third-order correlation measurement of picosecond laser pulse shapes”. In: *Measurement Science and Technology* 4, 1993, p. 1426. (See p. 77)

[161] N. Stuart, T. Robinson, D. Hillier, N. Hopps, B. Parry, I. Musgrave, G. Nersisyan, A. Sharba, M. Zepf, and R. A. Smith. “Comparative study on the temporal contrast of femtosecond mode-locked laser oscillators”. In: *Optics Letters* 41, 2016, pp. 3221–3224. (See p. 77)

[162] Amplitude Laser. *Sequoia - Ultrafast Pulse Characterization*. https://amplitude-laser.com/wp-content/uploads/2023/07/Sequoia-Ref-1211-rev-E.pdf. Accessed: 2025-05-28. Datasheet downloaded from https://amplitude-laser.com/add_ons/metrology/sequoia/. 2023 (See p. 77)

[163] J. A. Fullalove. “Characterisation, Enhancement and Impact of Temporal Intensity Contrast in Multi-Terawatt Laser Systems”. PhD thesis. Imperial College London, 2023. (See p. 78)

[164] S. Jalas, M. Kirchen, P. Messner, P. Winkler, L. Hübner, J. Dirkwinkel, M. Schnepp, R. Lehe, and A. R. Maier. “Bayesian Optimization of a Laser-Plasma Accelerator”. In: *Phyical Review Letters* 126, 2021, p. 104801. (See p. 81)

[165] G. J. Cao, C. A. Lindstrøm, E. Adli, S. Corde, and S. Gessner. “Positron acceleration in plasma wakefields”. In: *Physical Review Accelerators & Beams* 27, 2024, p. 034801. (See p. 81)

[166] A. Taflove, S. C. Hagness, and M. Piket-May. “9 - Computational Electromagnetics: The Finite-Difference Time-Domain Method”. In: *The Electrical Engineering Handbook*. Burlington: Academic Press, 2005. Pp. 629–670. (See p. 82)

[167] K. Yee. "Numerical Solution of Initial Boundary Value Problems Involving Maxwell's Equations in Isotropic Media". In: *IEEE Transactions on Antennas and Propagation* 14, 1966, pp. 302–307. (See p. 82)

[168] J. P. Boris. "Relativistic plasma simulation-optimization of a hybrid code". In: *Proceeding of the 4th Conference on Numerical Simulation of Plasmas*. 1970. Pp. 3–67. (See p. 82)

[169] J.-L. Vay. "Simulation of beams or plasmas crossing at relativistic velocity". In: *Physics of Plasmas* 15, 2008. (See p. 82)

[170] B. B. Godfrey. "Numerical Cherenkov instabilities in electromagnetic particle codes". In: *Journal of Computational Physics* 15, 1974, pp. 504–521. (See p. 82)

[171] R. Nuter, M. Grech, P. Gonzalez de Alaiza Martinez, G. Bonnaud, and E. d'Humieres. "Maxwell solvers for the simulations of the laser-matter interaction". In: *The European Physical Journal D* 68, 2014, pp. 1–9. (See p. 82)

[172] J. Derouillat, A. Beck, F. Pérez, T. Vinci, M. Chiaramello, A. Grassi, M. Flé, G. Bouchard, I. Plotnikov, N. Aunai, J. Dargent, C. Riconda, and M. Grech. "Smilei : A collaborative, open-source, multi-purpose particle-in-cell code for plasma simulation". In: *Computer Physics Communications* 222, 2018, pp. 351–373. (See pp. 83, 87, 167)

[173] P.-L. Bourgeois and X. Davoine. "Improved modellisation of laser–particle interaction in particle-in-cell simulations". In: *Journal of Plasma Physics* 89, 2023, p. 905890206. (See p. 83)

[174] A. Lifschitz, X. Davoine, E. Lefebvre, J. Faure, C. Rechatin, and V. Malka. "Particle-in-Cell modelling of laser–plasma interaction using Fourier decomposition". In: *Journal of Computational Physics* 228, 2009, pp. 1803–1814. (See p. 83)

[175] X. Davoine, E. Lefebvre, J. Faure, C. Rechatin, A. Lifschitz, and V. Malka. "Simulation of quasimonoenergetic electron beams produced by colliding pulse wakefield acceleration". In: *Physics of Plasmas* 15, 2008, p. 113102. (See p. 83)

[176] A. A. Golovanov and I. Y. Kostyukov. "Bubble regime of plasma wakefield in 2D and 3D geometries". In: *Physics of Plasmas* 25, 2018, p. 103107. (See p. 83)

[177] I. Zemzemi, F. Massimo, and A. Beck. "Azimuthal decomposition study of a realistic laser profile for efficient modeling of Laser WakeField Acceleration". In: *Journal of Physics: Conference Series* 1596, 2020, p. 012054. (See p. 83)

[178] R. Courant, K. Friedrichs, and H. Lewy. "Über die partiellen Differenzengleichungen der mathematischen Physik". In: *Mathematische Annalen* 100, 1928, pp. 32–74. (See p. 84)

[179] R. W. Hockney and J. W. Eastwood. *Computer simulation using particles*. McGraw-Hill, 1981. (See p. 85)

[180] Smilei Collaboration. *benchmarks*. `https://github.com/SmileiPIC/Smilei/tree/master/benchmarks`. Accessed: 2025-05-16. 2024 (See p. 87)

[181] Homer. *The Iliad*. Trans. by A. T. Murray. Loeb Classical Library. Book 6, lines 179–182. Cambridge, MA: Harvard University Press, 1924. (See p. 89)

[182] G. Andriukaitis, T. Balčiūnas, S. Ališauskas, A. Pugžlys, A. Baltuška, T. Popmintchev, M.-C. Chen, M. M. Murnane, and H. C. Kapteyn. "90 GW peak power few-cycle mid-infrared pulses from an optical parametric amplifier". In: *Optics Letters* 36, 2011, pp. 2755–2757. (See pp. 91, 93)

[183] V. Shumakova, P. Malevich, S. Ališauskas, A. Voronin, A. M. Zheltikov, D. Faccio, D. Kartashov, A. Baltuška, and A. Pugžlys. "Multi-millijoule few-cycle mid-infrared pulses through nonlinear self-compression in bulk". In: *Nature Communications* 7, 2016, p. 12877. (See pp. 91, 93)

[184] V. Shumakova, S. Ališauskas, A. Baltuška, P. Malevich, A. Voronin, A. Mitrofanov, D. Sidorov-Biryukov, A. Zheltikov, D. Kartashov, and A. Pugžlys. "Multi-mJ mid-IR light bullets in air". In: *EPJ Web of Conferences* 205, 2019, p. 01004. (See pp. 91, 93)

[185] D. Woodbury, L. Feder, V. Shumakova, C. Gollner, R. Schwartz, B. Miao, F. Salehi, A. Korolov, A. Pugžlys, A. Baltuška, and H. M. Milchberg. "Laser wakefield acceleration with mid-IR laser pulses". In: *Optics Letters* 43, 2018, pp. 1131–1134. (See pp. 91, 93, 145, 146, 165, 179–181)

[186] D. Woodbury. "Applications of Intense Mid-Infrared Laser-Plasma Interactions". PhD thesis. University of Maryland, 2020. (See pp. 91, 93)

[187] Y. Fu, B. Xue, K. Midorikawa, and E. J. Takahashi. "TW-scale mid-infrared pulses near 3.3 µm directly generated by dual-chirped optical parametric amplification". In: *Applied Physics Letters* 112, 2018, p. 241105. (See pp. 91, 93, 123, 124)

[188] P. Wang, Y. Li, W. Li, H. Su, B. Shao, S. Li, C. Wang, D. Wang, R. Zhao, Y. Peng, Y. Leng, R. Li, and Z. Xu. "2.6 mJ/100 Hz CEP-stable near-single-cycle 4 µm laser based on OPCPA and hollow-core fiber compression". In: *Optics Letters* 43, 2018, pp. 2197–2200. (See pp. 91, 93)

[189] G. Fan, T. Balčiūnas, T. Kanai, T. Flöry, G. Andriukaitis, B. E. Schmidt, F. Légaré, and A. Baltuška. "Hollow-core-waveguide compression of multi-millijoule CEP-stable 3.2 µm pulses". In: *Optica* 3, 2016, pp. 1308–1311. (See pp. 91, 93)

[190] P. Dey, L. Ehrentraut, J. Tümmler, M. Schnürer, and S. Eisebitt. "Above-mJ optical parametric chirped pulse amplifier at 3 µm for laser-driven coherent soft x-ray generation beyond the water window". In: *APL Photonics* 10, 2025, p. 040807. (See pp. 91, 93)

[191] A. Tarasevitch, M. Bridger, O. A. Naranjo-Montoya, U. Bovensiepen, A. Tarasevitch, and U. Bovensiepen. "Towards high power broad-band OPCPA at 3000 nm". In: *Optics Express, Vol. 27, Issue 22, pp. 31330-31337* 27, 2019, pp. 31330–31337. (See pp. 91, 93)

[192] C. Heese, C. Phillips, L. Gallmann, M. Fejer, and U. Keller. "Ultrabroadband, highly flexible amplifier for ultrashort midinfrared laser pulses based on aperiodically poled Mg:$LiNbO_3$". In: *Optics Letters* 35, 2010, pp. 2340–2342. (See pp. 91, 93)

[193] B. W. Mayer, C. R. Phillips, L. Gallmann, M. M. Fejer, and U. Keller. "Sub-four-cycle laser pulses directly from a high-repetition-rate optical parametric chirped-pulse amplifier at 3.4 µm". In: *Optics Letters* 38, 2013, pp. 4265–4268. (See pp. 91, 93)

[194] C. R. Phillips, B. W. Mayer, L. Gallmann, M. M. Fejer, and U. Keller. "Design constraints of optical parametric chirped pulse amplification based on chirped quasi-phase-matching gratings". In: *Optics Express* 22, 2014, pp. 9627–9658. (See pp. 91, 93)

[195] N. Thiré, R. Maksimenka, B. Kiss, C. Ferchaud, P. Bizouard, E. Cormier, K. Osvay, and N. Forget. "4-W, 100-kHz, few-cycle mid-infrared source with sub-100-mrad carrier-envelope phase noise". In: *Optics Express* 25, 2017, pp. 1505–1514. (See pp. 91, 93)

[196] N. Thiré, R. Maksimenka, B. Kiss, C. Ferchaud, G. Gitzinger, T. Pinoteau, H. Jousselin, S. Jarosch, P. Bizouard, V. D. Pietro, E. Cormier, K. Osvay, and N. Forget. "Highly stable, 15 W, few-cycle, 65 mrad CEP-noise mid-IR OPCPA for statistical physics". In: *Optics Express* 26, 2018, pp. 26907–26915. (See pp. 91, 93)

[197] T. Fujii, Y. Taguchi, T. Saiki, and Y. Nagasaka. "A Fusion-Spliced Near-Field Optical Fiber Probe Using Photonic Crystal Fiber for Nanoscale Thermometry Based on Fluorescence-Lifetime Measurement of Quantum Dots". In: *Sensors* 11, 2011, pp. 8358–8369. (See p. 99)

[198] Thorlabs. *NL-PM-750 - Highly Nonlinear PM PCF, Supercontinuum Generation, 750 nm ZD, 1.8 µm Core*. `https://www.thorlabs.us/thorproduct.cfm?partnumber=NL-PM-750`. 2017 (See p. 99)

[199] J. E. Murray. "Pulsed gain and thermal lensing of Nd:LiYF4". In: *IEEE Journal of Quantum Electronics* 19, 1983, pp. 488–491. (See p. 102)

[200] G. D. Goodno, Z. Guo, R. J. D. Miller, I. J. Miller, J. W. Montgomery, S. R. Adhav, and R. S. Adhav. "Investigation of β-BaB2O4 as a Q switch for high power applications". In: *Applied Physics Letters* 66, 1995, pp. 1575–1577. (See p. 102)

[201] K. Mecseki, D. Bigourd, S. Patankar, N. H. Stuart, and R. A. Smith. "Flat-top picosecond pulses generated by chirped spectral modulation from a Nd:YLF regenerative amplifier for pumping few-cycle optical parametric amplifiers". In: *Applied Optics* 53, 2014, pp. 2229–2235. (See p. 104)

[202] N. Ishii, C. Teisset, T. Fuji, S. Kohler, K. Schmid, L. Veisz, A. Baltuska, and F. Krausz. “Seeding of an eleven femtosecond optical parametric chirped pulse amplifier and its Nd/sup 3+/ picosecond pump laser from a single broadband Ti:Sapphire oscillator”. In: *IEEE Journal of Selected Topics in Quantum Electronics* 12, 2006, pp. 173–180. (See p. 104)

[203] Cutting Edge Optronics, Northrop Grumman Corporation. *RB Plus User Manual*. http://www.ngceoservice.com/data/doc57/ceo-uman-0004j.pdf. 2012 (See p. 104)

[204] H. Y. Fan. “Temperature Dependence of the Energy Gap in Semiconductors”. In: *Physical Review* 82, 1951, pp. 900–905. (See p. 104)

[205] I. N. Ross, M. Csatári, and S. Hutchins. “High-performance diode-pumped Nd:YLF amplifier”. In: *Applied Optics* 42, 2003, pp. 1040–1047. (See p. 107)

[206] I. N. Ross, P. Matousek, G. H. C. New, and K. Osvay. “Analysis and optimization of optical parametric chirped pulse amplification”. In: *Journal of the Optical Society of America B* 19, 2002, pp. 2945–2956. (See p. 115)

[207] T.-J. Wang, Z. Major, I. Ahmad, S. A. Trushin, F. Krausz, and S. Karsch. “Ultrabroadband near-infrared pulse generation by noncollinear OPA with angular dispersion compensation”. In: *Applied Physics B* 100, 2010, pp. 207–214. (See p. 116)

[208] S.-W. Huang, J. Moses, and F. X. Kärtner. “Broadband noncollinear optical parametric amplification without angularly dispersed idler”. In: *Optics Letters* 37, 2012, pp. 2796–2798. (See p. 116)

[209] M. Mero and V. Petrov. “High-Power, Few-Cycle, Angular Dispersion Compensated Mid-Infrared Pulses From a Noncollinear Optical Parametric Amplifier”. In: *IEEE Photonics Journal* 9, 2017, pp. 1–8. (See p. 116)

[210] S. Wandel, G. Xu, Y. Yin, a. -, A. Grupp, A. Budweg, M. P. Fischer, C. Blancard, P. Cossé, G. Faussurier -, C. Manzoni, and G. Cerullo. “Design criteria for ultrafast optical parametric amplifiers”. In: *Journal of Optics* 18, 2016, p. 103501. (See p. 116)

[211] H. Furuya, A. Morikawa, K. Mizuuchi, and K. Yamamoto. “High-beam-quality continuous wave 3 W green-light generation in bulk periodically poled MgO: LiNbO3”. In: *Japanese Journal of Applied Physics* 45, 2006, p. 6704. (See p. 117)

[212] K. Mizuuchi, K. Yamamoto, and M. Kato. “Harmonic blue light generation in bulk periodically poled MgO:$LiNbO_3$”. In: *Electronics Letters* 32, 1996, pp. 2091–2092. (See p. 117)

[213] Amplitude / Fastlite. *Dazzler: ultrafast acousto-optic programmable dispersive filter*. https://amplitude-laser.com/products/femtosecond-lasers/instrumentation-lasers-femtosecondes/dazzler/. Accessed: 2025-07-30. Amplitude / Fastlite, 2023 (See p. 123)

[214] W. C. Hurlbut, Y.-S. Lee, K. Vodopyanov, P. Kuo, and M. Fejer. "Multiphoton absorption and nonlinear refraction of GaAs in the mid-infrared". In: *Optics Letters* 32, 2007, pp. 668–670. (See p. 125)

[215] B.-U. Sohn, C. Monmeyran, L. Kimerling, A. Agarwal, and D. Tan. "Kerr nonlinearity and multi-photon absorption in germanium at mid-infrared wavelengths". In: *Applied Physics Letters* 111, 2017. (See p. 125)

[216] B. Qiao, S. Dai, Y. Xu, P. Zhang, X. Shen, T. Xu, Q. Nie, W. Ji, and F. Chen. "Third-order optical nonlinearities of chalcogenide glasses within Ge-Sn-Se ternary system at a mid-infrared window". In: *Optical Materials Express* 5, 2015, pp. 2359–2365. (See p. 125)

[217] R. Lin, F. Chen, X. Zhang, Y. Huang, B. Song, S. Dai, X. Zhang, and W. Ji. "Mid-infrared optical properties of chalcogenide glasses within tin-antimony-selenium ternary system". In: *Optics Express* 25, 2017, pp. 25674–25688. (See p. 125)

[218] F. Théberge, P. Mathieu, N. Thiré, J.-F. Daigle, B. E. Schmidt, J. Fortin, R. Vallée, Y. Messaddeq, and F. Légaré. "Mid-infrared nonlinear absorption in As_2S_3 chalcogenide glass". In: *Optics Express* 24, 2016, pp. 24600–24610. (See p. 125)

[219] Y. Petrov. *Optometrika*. `https://github.com/caiuspetronius/Optometrika`. GitHub repository. Retrieved 18 March, 2024. 2025 (See p. 125)

[220] N. Flemens, D. Heberle, J. Zheng, D. J. Dean, C. Davis, K. Zawilski, P. G. Schunemann, and J. Moses. *High quantum efficiency parametric amplification via hybridized nonlinear optics*. 2022. arXiv: `2207.04147 [physics.optics]` (See p. 128)

[221] Y. Sharma, P. Ansaria, R. Sharma, D. Mathur, and R. Dard. "Bandgap tuning of optical and electrical properties of zinc selenide". In: *Chalcogenide Lett* 18, 2021, pp. 183–189. (See p. 133)

[222] A. H. Chin, O. G. Calderón, and J. Kono. "Extreme Midinfrared Nonlinear Optics in Semiconductors". In: *Physical Review Letters* 86, 2001, pp. 3292–3295. (See p. 134)

[223] K. Werner, M. G. Hastings, A. Schweinsberg, B. L. Wilmer, D. Austin, C. M. Wolfe, M. Kolesik, T. R. Ensley, L. Vanderhoef, A. Valenzuela, *et al.* "Ultrafast mid-infrared high harmonic and supercontinuum generation with n_2 characterization in zinc selenide". In: *Optics express* 27, 2019, pp. 2867–2885. (See p. 134)

[224] M. Awad, A. Manna, S. Hell, B. Ying, L. Ábrók, Z. Divéki, E. Cormier, B. Kiss, J. Böhmer, C. Ronning, *et al.* "Few-cycle laser pulse characterization on-target using high-harmonic generation from nano-scale solids". In: *Optics Express* 32, 2024, pp. 1325–1333. (See p. 135)

[225] K. Põder, J. C. Wood, N. C. Lopes, J. M. Cole, S. Alatabi, M. P. Backhouse, P. S. Foster, A. J. Hughes, C. Kamperidis, O. Kononenko, S. P. D. Mangles, C. A. J. Palmer, D. Rusby, A. Sahai, G. Sarri, D. R. Symes, J. R. Warwick, and Z. Najmudin. "Multi-GeV Electron Acceleration in Wakefields Strongly Driven by Oversized Laser Spots". In: *Physical Review Letters* 132, 2024, p. 195001. (See pp. 138, 139, 146, 147, 170)

[226] W. Lu, C. Huang, M. Zhou, W. B. Mori, and T. Katsouleas. "Nonlinear Theory for Relativistic Plasma Wakefields in the Blowout Regime". In: *Physical Review Letters* 96, 2006, p. 165002. (See p. 141)

[227] A. E. Gunn. *Thesis_LWFA_database*. https://github.com/agunn21/Thesis_LWFA_database. 2025 (See p. 145)

[228] R. Zgadzaj, J. Welch, Y. Cao, D. Amorim L, A. Cheng, A. Gaikwad, P. Iapozzutto, P. Kumar, V. N. Litvinenko, I. Petrushina, R. Samulyak, N. Vafaei-Najafabadi, C. Joshi, C. Zhang, M. Babzien, M. Fedurin, R. Kupfer, K. Kusche, M. A. Palmer, I. V. Pogorelsky, M. N. Polyanskiy, C. Swinson, and M. C. Downer. "Plasma electron acceleration driven by a long-wave-infrared laser". In: *Nature Communications* 15, 2024, p. 4037. (See p. 146)

[229] Y. Cao, R. Zgadzaj, N. Vafaei-Najafabadi, I. Petrushina, R. Samulyak, V. Litvinenko, and M. Downer. "Emittance Preservation and Optimization in an Externally-injected CO_2-laser-driven Wakefield Accelerator". In: *2022 IEEE Advanced Accelerator Concepts Workshop (AAC)*, 2022, pp. 1–5. (See p. 146)

[230] J. Kim, T. Wang, V. Khudik, and G. Shvets. "Subfemtosecond Wakefield Injector and Accelerator Based on an Undulating Plasma Bubble Controlled by a Laser Phase". In: *Physical Review Letters* 127, 2021, p. 164801. (See pp. 151, 164)

[231] A. F. Lifschitz and V. Malka. "Optical phase effects in electron wakefield acceleration using few-cycle laser pulses". In: *New Journal of Physics* 14, 2012, p. 053045. (See pp. 151, 164)

[232] S. Gordienko and A. Pukhov. "Scalings for ultrarelativistic laser plasmas and quasimonoenergetic electrons". In: *Physics of Plasmas* 12, 2005. (See p. 155)

[233] S. Wilks, T. Katsouleas, J. M. Dawson, P. Chen, and J. J. Su. "Beam Loading in Plasma Waves". In: *IEEE Transactions on Plasma Science* 15, 1987, pp. 210–217. (See p. 158)

[234] P. Chen, J. M. Dawson, R. W. Huff, and T. Katsouleas. "Acceleration of Electrons by the Interaction of a Bunched Electron Beam with a Plasma". In: *Physical Review Letters* 54, 1985, pp. 693–696. (See p. 158)

[235] P. Muggli, B. E. Blue, C. E. Clayton, S. Deng, F.-J. Decker, M. J. Hogan, C. Huang, R. Iverson, C. Joshi, T. C. Katsouleas, S. Lee, W. Lu, K. A. Marsh, W. B. Mori, C. L. O'Connell, P. Raimondi, R. Siemann, and D. Walz. "Meter-Scale Plasma-Wakefield Accelerator Driven by a Matched Electron Beam". In: *Physical Review Letters* 93, 2004, p. 014802. (See p. 158)

[236] C. L. O'Connell, C. D. Barnes, F.-J. Decker, M. J. Hogan, R. Iverson, P. Krejcik, R. Siemann, D. R. Walz, C. E. Clayton, C. Huang, D. K. Johnson, C. Joshi, W. Lu, K. A. Marsh, W. Mori, M. Zhou, S. Deng, T. Katsouleas, P. Muggli, and E. Oz. "Plasma production via field ionization". In: *Physical Review Special Topics - Accelerators and Beams* 9, 2006, p. 101301. (See p. 158)

[237] M. Tzoufras, W. Lu, F. S. Tsung, C. Huang, W. B. Mori, T. Katsouleas, J. Vieira, R. A. Fonseca, and L. O. Silva. "Beam Loading in the Nonlinear Regime of Plasma-Based Acceleration". In: *Physical Review Letters* 101, 2008, p. 145002. (See p. 158)

[238] M. Tzoufras, W. Lu, F. S. Tsung, C. Huang, W. B. Mori, T. Katsouleas, J. Vieira, R. A. Fonseca, and L. O. Silva. "Beam loading by electrons in nonlinear plasma wakesa)". In: *Physics of Plasmas* 16, 2009, p. 056705. (See p. 158)

[239] C. Rechatin, X. Davoine, A. Lifschitz, A. B. Ismail, J. Lim, E. Lefebvre, J. Faure, and V. Malka. "Observation of Beam Loading in a Laser-Plasma Accelerator". In: *Physical Review Letters* 103, 2009, p. 194804. (See pp. 158, 164)

[240] C. Rechatin, J. Faure, X. Davoine, O. Lundh, J. Lim, A. Ben-Ismaïl, F. Burgy, A. Tafzi, A. Lifschitz, E. Lefebvre, and V. Malka. "Characterization of the beam loading effects in a laser plasma accelerator". In: *New Journal of Physics* 12, 2010, p. 045023. (See pp. 158, 164)

[241] J. Couperus, R. Pausch, A. Köhler, O. Zarini, J. Krämer, M. Garten, A. Huebl, R. Gebhardt, U. Helbig, S. Bock, K. Zeil, A. Debus, M. Bussmann, U. Schramm, and A. Irman. "Demonstration of a beam loaded nanocoulomb-class laser wakefield accelerator". In: *Nature Communications* 8, 2017, p. 487. (See p. 158)

[242] K. H. Pae, I. W. Choi, and J. Lee. "Self-mode-transition from laser wakefield accelerator to plasma wakefield accelerator of laser-driven plasma-based electron acceleration". In: *Physics of Plasmas* 17, 2010. (See pp. 159, 179)

[243] P. Masson-Laborde, M. Mo, A. Ali, S. Fourmaux, P. Lassonde, J. Kieffer, W. Rozmus, D. Teychenné, and R. Fedosejevs. "Giga-electronvolt electrons due to a transition from laser wakefield acceleration to plasma wakefield acceleration". In: *Physics of Plasmas* 21, 2014. (See pp. 159, 161, 179)

[244] X. Chang, M. Zeng, J. Wang, and D. Li. "Energy boosting and scaling of self-guided hybrid laser-plasma wakefield acceleration in a single uniform plasma". In: *New Journal of Physics* 27, 2025, p. 023021. (See p. 159)

[245] J. Götzfried, A. Döpp, M. F. Gilljohann, F. M. Foerster, H. Ding, S. Schindler, G. Schilling, A. Buck, L. Veisz, and S. Karsch. "Physics of High-Charge Electron Beams in Laser-Plasma Wakefields". In: *Physical Review X* 10, 2020, p. 041015. (See pp. 159, 161, 179)

[246] Y. Wu, C. Yu, Z. Qin, W. Wang, Z. Zhang, R. Qi, K. Feng, L. Ke, Y. Chen, C. Wang, X. Yang, Y. Leng, J. Liu, R. Li, and Z. Xu. "Energy Enhancement and Energy Spread Compression of Electron Beams in a Hybrid Laser-Plasma Wakefield Accelerator". In: *Applied Sciences* 9, 2019. (See p. 161)

[247] E. Guillaume, A. Döpp, C. Thaury, A. Lifschitz, J.-P. Goddet, A. Tafzi, F. Sylla, G. Iaquanello, T. Lefrou, P. Rousseau, K. T. Phuoc, and V. Malka. "Physics of fully-loaded laser-plasma accelerators". In: *Physical Review Special Topics - Accelerators and Beams* 18, 2015, p. 061301. (See pp. 161, 164)

[248] G. Sciaini and R. J. D. Miller. "Femtosecond electron diffraction: heralding the era of atomically resolved dynamics". In: *Reports on Progress in Physics* 74, 2011, p. 096101. (See p. 174)

[249] Z.-H. He, B. Beaurepaire, J. Nees, G. Gallé, S. Scott, J. S. Pérez, M. Lagally, K. Krushelnick, A. Thomas, and J. Faure. "Capturing structural dynamics in crystalline silicon using chirped electrons from a laser wakefield accelerator". In: *Scientific Reports* 6, 2016, p. 36224. (See p. 174)

[250] W. Schumaker, N. Nakanii, C. McGuffey, C. Zulick, V. Chyvkov, F. Dollar, H. Habara, G. Kalintchenko, A. Maksimchuk, K. A. Tanaka, A. G. R. Thomas, V. Yanovsky, and K. Krushelnick. "Ultrafast Electron Radiography of Magnetic Fields in High-Intensity Laser-Solid Interactions". In: *Physical Review Letters* 110, 2013, p. 015003. (See p. 174)

[251] Y. Glinec, J. Faure, L. L. Dain, S. Darbon, T. Hosokai, J. J. Santos, E. Lefebvre, J. P. Rousseau, F. Burgy, B. Mercier, and V. Malka. "High-Resolution γ-Ray Radiography Produced by a Laser-Plasma Driven Electron Source". In: *Physical Review Letters* 94, 2005, p. 025003. (See p. 174)

[252] F. Mirani, D. Calzolari, A. Formenti, and M. Passoni. "Superintense laser-driven photon activation analysis". In: *Communications Physics* 4, 2021, p. 185. (See p. 174)

[253] K. K. Kainz, K. R. Hogstrom, J. A. Antolak, P. R. Almond, C. D. Bloch, C. Chiu, M. Fomytskyi, F. Raischel, M. Downer, and T. Tajima. "Dose properties of a laser accelerated electron beam and prospects for clinical application". In: *Medical Physics* 31, 2004, pp. 2053–2067. (See p. 180)

[254] C. Chiu, M. Fomytskyi, F. Grigsby, F. Raischel, M. C. Downer, and T. Tajima. "Laser electron accelerators for radiation medicine: A feasibility study". In: *Medical Physics* 31, 2004, pp. 2042–2052. (See p. 180)

[255] S. Kneip, C. McGuffey, F. Dollar, M. S. Bloom, V. Chvykov, G. Kalintchenko, K. Krushelnick, A. Maksimchuk, S. P. D. Mangles, T. Matsuoka, Z. Najmudin, C. A. J. Palmer, J. Schreiber, W. Schumaker, A. G. R. Thomas, and V. Yanovsky. "X-ray phase contrast imaging of biological specimens with femtosecond pulses of betatron radiation from a compact laser plasma wakefield accelerator". In: *Applied Physics Letters* 99, 2011, p. 093701. (See p. 180)

[256] J. M. Cole, J. C. Wood, N. C. Lopes, K. Põder, R. L. Abel, S. Alatabi, J. S. J. Bryant, A. Jin, S. Kneip, K. Mecseki, D. R. Symes, S. P. D. Mangles, and Z. Najmudin. "Laser-wakefield accelerators as hard x-ray sources for 3D medical imaging of human bone". In: *Scientific Reports* 5, 2015, p. 13244. (See p. 180)

[257] J. Feng, Y. Li, J. Tan, W. Wang, Y. Li, X. Zhang, Y. Meng, X. Ge, F. Liu, W. Yan, C. Fu, L. Chen, and J. Zhang. "Laser Plasma-Accelerated Ultra-Intense Electron Beam for Efficiently Exciting Nuclear Isomers". In: *Laser & Photonics Reviews* 17, 2023, p. 2300514. (See p. 180)

[258] L. Martelli, O. Kononenko, I. Andriyash, J. Wheeler, J. Gautier, J.-P. Goddet, A. Tafzi, R. Lahaye, C. Giaccaglia, A. Flacco, V. Tomkus, M. Mackevi či ūt ė, J. Dudutis, V. Stankevic, P. Ge čys, G. Ra čiukaitis, H. Kraft, X. Dinh, and C. Thaury. "Physics of high-charge laser-plasma accelerators for few-MeV applications". In: *Physical Review Applied* 23, 2025, p. 034033. (See p. 181)

[259] A. J. Goers, G. A. Hine, L. Feder, B. Miao, F. Salehi, J. K. Wahlstrand, and H. M. Milchberg. "Multi-MeV Electron Acceleration by Subterawatt Laser Pulses". In: *Physical Review Letters* 115, 2015, p. 194802. (See p. 181)

[260] D. Guénot, D. Gustas, A. Vernier, B. Beaurepaire, F. Böhle, M. Bocoum, M. Lozano, A. Jullien, R. Lopez-Martens, A. Lifschitz, *et al.* "Relativistic electron beams driven by kHz single-cycle light pulses". In: *Nature Photonics* 11, 2017, pp. 293–296. (See p. 181)

[261] L.-L. Yu, E. Esarey, C. B. Schroeder, J.-L. Vay, C. Benedetti, C. G. R. Geddes, M. Chen, and W. P. Leemans. "Two-Color Laser-Ionization Injection". In: *Physical Review Letters* 112, 2014, p. 125001. (See p. 185)