Coherent Combination of Mismatched Fibre Lasers - Steps Towards Real World Applications

Peter Tudor

St. Hugh's College, Oxford





Thesis submitted in fulfilment of the requirements for the degree of Doctor of Philosophy at the University of Oxford

Trinity Term, 2018

Abstract

Large arrays of coherently-combined Quasi-Continuous-Wave (QCW) fibre lasers have been identified as a plausible solution to creating high-repetition rate (∼1kHz), high-wall-plug efficiency lasers capable of driving a wakefield for accelerator applications. The coherent combination of four near-identical Photonic Crystal Fibre (PCF) lasers to produce laser pulses with pulse energies of 5.7 mJ and peak powers of 22 GW has been successfully demonstrated by other groups. Extensions of the four PCF experiment also demonstrated the coherent combination of 8 photonic crystal fibres to attain pulse energies of 12 mJ and peak powers of 35 GW. However, these PCFs were near-identical in length, core size, fibre structure and dopant concentration. Work done in Oxford focused on the coherent combination of fibre lasers which were unmatched in length, core size, internal structure and dopant concentration. The aim was to identify the point at which the fibre differences caused the combining efficiency to degrade. Potentially, this information would be used to identify the tolerance levels required for the large-scale manufacture of the fibres needed to construct a facility-scale laser system for laser-wakefield acceleration.

This thesis covers the coherent combination of mismatched PCFs, a scenario which could occur during the construction of a facility-scale terawatt laser system. Here, hundreds or potentially thousands of PCFs could be purchased from different manufacturers, each of whom use their own methods of manufacture and tolerances. An investigation into the challenges associated with combining two highly-mismatched PCFs was performed, with the PCFs chosen to be very different to exaggerate any effects which could degrade their coherent combination. This work resulted in the successful coherent combination of these highly-mismatched PCFs, with a measured degree of linear polarisation of $85.9 \pm 0.5\%$ and a fraction of available beam power of $93.0 \pm 0.5\%$. Also, supporting work investigated the difference between simulation and measurement in explaining the observed degree of linear-polarisation of the combined-beam.

Acknowledgements

Firstly, I would like to thank my mother, Melanie Tudor, and my sisters: Lucy, Holly, Poppy and Lotty, and my girlfriend, Niloofar Karaji, for their unconditional support during my times of hardship and low mood. Also, I wish to thank my father, David, for giving up his time to help proof-read the final thesis draft.

My supervisors, Prof. Roman Walczak and Dr. Laura Corner, have had a very large influence on me during my time at Oxford and I am grateful for their time and patience. They have helped me to grow and mature into an independent scientist. I offer my thanks and enduring appreciation to them both.

I would like to thank Rob Shalloo for the seemingly limitless supply of coffee, friendship and moral support he has provided during my time at Oxford. Master students Tom Hird and Harry Jones have also been a source of friendship and humour during their time working in the same lab on their own projects and I thank them both. Also, I wish to add a thank you to Julian Gascoyne, University of Edinburgh, for giving his time to help read and spell check some of the chapters of this thesis.

For most of my adult life I have served in the military in a part-time capacity, enabling me to follow my aspirations of a degree and DPhil while still being able to rise up to the challenges offered by the military. This was only possible with the support of my squadron's commanding officer and so I wish to express my gratitude to Squadron Leader Steve Embleton, 2620 Sqn RAuxAF Regiment, for his patience and understanding while I conducted my studies.

I would like to thank Dr. Suzie Sheehy for her friendship and company during coffee breaks, out-of-work time and runs alike. She has been a source of Australian humour, interesting cooking choices and guilt-free coffee consumption.

Finally, I would like to thank Dr. Jimmy Holloway for his time and assistance in generating particle in cell simulations of laser plasma wakefield acceleration for presentation in this thesis.

Contents

Abstract **i**

Acknowledgements **i**

List of figures **x**

Glossary **xi**

1 Introduction **1**

1.1 Accelerators 1

1.2 Laser plasma wakefield acceleration 6

1.3 Laser wakefield acceleration with terawatt class lasers 12

1.4 Fibre lasers 14

1.5 Research in the field of fibre lasers 24

1.6 Quantifying coherent combination 28

1.7 Outline of the remainder of thesis 29

2 Fibre lasers **31**

2.1 Solid state ultrafast lasers 31

2.1.1 Drawbacks of solid state lasers for laser wakefield acceleration 32

2.2 Fibres 34

2.2.1 Fibre modes 34

2.3 Photonic Crystal Fibres 41

2.3.1 End caps 44

2.3.2 Advantages and disadvantages of photonic crystal fibres 45

2.4 Photonic Crystal Fibres used in this thesis 46
2.4.1 The 70 μm Photonic Crystal Fibre 47
2.4.2 The 85 μm Photonic Crystal Fibre 49
2.4.3 Pump coupling 51
2.5 Summary 54

3 Coherent Combination of Fibre Lasers 56
3.1 Tiled aperture setups 57
3.2 Filled aperture setups 60
3.3 Passive stabilisation of beams 64
3.3.1 Sagnac interferometer 64
3.3.2 Divided Pulse Amplification 65
3.4 Active stabilisation of beams 66
3.4.1 Hänsch-Couillaud polarisation locking 68
3.5 Summary 71

4 Ultrafast laser theory 73
4.1 Outline 73
4.2 Ultrafast lasers 75
4.2.1 Pulse theory 75
4.2.2 Non-linear effects 79
4.3 Modelling coherent combination 81
4.3.1 Model of combination 81
4.3.2 Non-linear phase compensation 84
4.4 Coherent combination of N lasers 89
4.5 Summary 90

5 Experimental Setup 91
5.1 Overview of the experiment 91
5.2 Beam transport 93
5.3 The interferometer 94
5.4 Locking and stabilisation 96

5.5 Pulse compression . . . 97
5.5.1 Compressor model . . . 98
5.5.2 Compressor design . . . 99
5.5.3 Compressor testing . . . 100
5.6 Single shot autocorrelator . . . 104
5.6.1 Single shot autocorrelator theory . . . 104
5.6.2 Single shot autocorrelator setup . . . 109
5.7 Beam combination . . . 111
5.8 Diagnostics . . . 112
5.9 Summary . . . 114

6 Seed Laser Characterisation **116**
6.1 Seed laser system . . . 117
6.1.1 Seed laser diagram . . . 117
6.1.2 Oscillator . . . 117
6.1.3 Stretcher . . . 117
6.1.4 Acoustic Optical Modulator . . . 120
6.1.5 Amplifiers . . . 120
6.1.6 Electro Optical Modulator . . . 122
6.1.7 System pulse characteristics . . . 122
6.2 Split Step Fourier method . . . 127
6.3 Summary . . . 140

7 Wavefront and Thermal Effects **142**
7.1 Wavefront considerations . . . 142
7.2 Beam divergence matching . . . 153
7.3 Thermal effects . . . 159
7.4 Summary . . . 165

8 Results **166**
8.1 Low power combination Results . . . 166
8.1.1 Power . . . 167

8.1.2 Degree of linear polarisation and efficiencies . . . 168
8.1.3 Beam profiles . . . 169
8.1.4 Autocorrelations . . . 170
8.1.5 Summary of low power coherent combination . . . 172
8.2 High power combination results . . . 173
8.2.1 Experimental modification . . . 174
8.2.2 Power . . . 178
8.2.3 Degree of linear polarisation . . . 178
8.2.4 Efficiencies . . . 180
8.2.5 Beam profiles . . . 180
8.2.6 Autocorrelations . . . 181
8.2.7 Lock stability . . . 184
8.2.8 High power Experiment summary . . . 187
8.3 Summary . . . 188

9 Analysis of Combination **190**
9.1 Estimate of degree-of-linear-polarisation loss . . . 190
9.2 Endcap and path length analysis . . . 207
9.3 Summary . . . 209

10 Conclusions and Future work **211**
10.1 Conclusions . . . 211
10.2 Future work . . . 217
10.3 Closing remarks . . . 225

A Datasheets **227**
A.1 70 µm Photonic Crystal Fibre Datasheet . . . 228
A.2 85 µm Photonic Crystal Fibre Datasheet . . . 229

References **230**

List of Figures

1.1 Accelerator comparisons 2
1.2 Radio frequency cavity 3
1.3 Synchrotron radiation example 5
1.4 Laser-plasma wakefield acceleration simulation 7
1.5 Radio-Frequency-cavity/wakefield comparison 7
1.6 Ponderomotive force timelapse 11
1.7 Bessel beam 16
1.8 Chirped pulse amplification 19
1.9 Polarisation splitting 22
1.10 Standard interferometers 23

2.1 Terawatt laser architecture 32
2.2 Wave vector triangle 36
2.3 Fibre modes 40
2.4 Photonic crystal fibre internal structure 42
2.5 Photonic crystal fibre air hole definition 43
2.6 Photonic crystal fibre hole parameter plot 44
2.7 End cap example 45
2.8 Yb-doped silica gain spectrum 46
2.9 A 70 μm photonic crystal fibre schematic diagram 48
2.10 Beam focus diagram 48
2.11 A 70 μm photonic crystal fibre decoupled beam 50
2.12 A 85 μm photonic crystal fibre schematic diagram 50
2.13 A 85 μm photonic crystal fibre decoupled beam 51

2.14 A 85 μm photonic crystal fibre Gaussian beam comparison 52
2.15 Pump laser spectra 53
2.16 Pump beam telescope 54

3.1 Tiled aperture array example 58
3.2 Tiled aperture array beam combination 59
3.3 Full aperture setup 61
3.4 Beam splitting ultrafast laser pulses using a diffractive optical element. 62
3.5 Polarisation splitting and recombination 63
3.6 Passively stabilised amplifier 65
3.7 Divided pulse amplification example 65
3.8 Hänsch-Couillaud phase locking example 69
3.9 Hänsch-Couillaud detector layout 70
3.10 Hänsch-Couillaud polarisation locking system example 70
3.11 Utilisation of Hänsch-Couillaud polarisation locking. 71

4.1 Compressor example 78
4.2 Example of stretcher 78
4.3 Combined beam visibility 84
4.4 Frequency phase difference example 85
4.5 Phase difference combination effect 86
4.6 Figure of merit diagrams 88

5.1 Locations of power measurements 92
5.2 Beam transport post-seed-laser 94
5.3 Mach-Zehnder interferometer diagram 95
5.4 Proportional integral differential controller, closed loop example 98
5.5 Compressor layout 100
5.6 Compressor annotated 101
5.7 Compressor optimisation curve 102
5.8 Scanning autocorrelator autocorrelation 102
5.9 Scanning autocorrelator statistics 105

5.10 Single shot autocorrelator layout . . . 105
5.11 Crossed beam geometry in the frequency doubling crystal. . . . 106
5.12 Index ellipsoid . . . 108
5.13 Single shot autocorrelator calibration . . . 110
5.14 Dataray lineout . . . 110
5.15 Free running stability . . . 112
5.16 Diagnostics layout overview . . . 113

6.1 Seed laser layout . . . 118
6.2 Spectral output of the oscillator . . . 119
6.3 Seed laser autocorrelation . . . 119
6.4 Stretcher layout . . . 121
6.5 pre-amplifier fibre schematic diagram . . . 122
6.6 Aplitude Systemé laser amplifiers . . . 123
6.7 Amplitude Systemé seed laser spectral evolution . . . 126
6.8 compressed seed laser pulse . . . 127
6.9 Non-Linear contributions to the model . . . 131
6.10 Split Step Fourier method model of whole laser system . . . 134
6.11 Spectral shaping demonstration . . . 135
6.12 Simulated and measured autocorrelation comparison . . . 136
6.13 Seed laser simulated intensity profile . . . 137
6.14 Split step Fourier method model visibility prediction . . . 138
6.15 600 μJ pulse autocorrelation . . . 139
6.16 100 μJ pulse autocorrelation . . . 140

7.1 Mismatched beams observed through polarising beam splitter . . . 143
7.2 Wavefront interaction schematic diagram . . . 144
7.3 Exaggerated wavefront mismatch . . . 145
7.4 No wavefront mismatch example . . . 147
7.5 Small wavefront mismatch example . . . 147
7.6 Large wavefront mismatch example . . . 148

7.7 70 and 85 μm photonic crystal fibre beams . . . 149
7.8 Combined beams comparison grid . . . 150
7.9 Beam angular mismatch grid . . . 151
7.10 Moderate angular beam separation and wavefront mismatch . . . 152
7.11 Actual combining beams with wavefront and angular mismatch . . . 152
7.12 Pump and seed beam coupling . . . 154
7.13 Beam divergence matching example . . . 155
7.14 Example telescope in beam line . . . 156
7.15 Effect of the telescope on the combined beam . . . 157
7.16 Comparison of beams before and after optimisation . . . 158
7.17 The 70 μm photonic crystal fibre beam size under pumping . . . 160
7.18 Beam pointing stability . . . 161
7.19 High power beam profiles after thermalisation . . . 162
7.20 Low power regime, beam power lost to the inner cladding . . . 163

8.1 Low power regime, beam profiles . . . 170
8.2 Low power regime, combined beam profile comparison . . . 171
8.3 Low power regime autocorrelation . . . 172
8.4 High power regime, 70 μm photonic crystal fibre thermal behaviour. . . . 176
8.5 High power regime, beam size under pumping . . . 178
8.6 High power regime, combined beam profile comparison . . . 181
8.7 High power regime, combined and rejected beams . . . 182
8.8 High power regime autocorrelations . . . 182
8.9 High power regime, simulate and measured autocorrelation comparison . . . 183
8.10 High power regime, lock stability . . . 184
8.11 Explantion of stability measurement . . . 186
8.12 High power regime, root-mean-square noise histogram . . . 186
8.13 High power regime, power spectral density . . . 187

9.1 Photonic crystal fibre wavefront comparison . . . 192
9.2 Phase difference colourmap . . . 193

9.3 Zernike polynomials . . . 194
9.4 Zernike mode contributions . . . 195
9.5 Phase difference comparison . . . 196
9.6 Fine structure phase difference map . . . 197
9.7 Zernike mode, fraction of available beam power effects . . . 198
9.8 Zernike mode fraction of available beam power effects, omitting mode 5 . . . 199
9.9 Beam cross-hair comparison . . . 201
9.10 Phase difference at one beam radius . . . 201
9.11 Phase difference given contribution from tilt . . . 202
9.12 Visibility scan over defocus term . . . 204
9.13 Phase Difference vs Defocus coefficient . . . 205
9.14 Radius of curvature comparison . . . 206
9.15 Telescope length scan . . . 206

10.1 Beam divergence modification . . . 218
10.2 Spatial light modulator used for wavefront correction . . . 219
10.3 Staged locking domains . . . 223

Glossary

BBO Barium Borate (BaB_2O_4)

CCD Charge-Coupled Device

CPA Chirped Pulse Amplification

CW Continuous Wave

DOE Dispersive Optical Element

FOM Figure Of Merit

FWHM Full Width Half Maximum

HWP Half Wave Plate

ICAN International Coherent Amplification Network

LBO Lithium Triborate (LiB_3O_5)

LOCSET Locking of Optical Coherence by Single detector Electronic frequency Tagging technology

MFD Mode Field Diameter

PBS Polarising Beam Splitter

PCF Photonic Crystal Fibre

PID Proportional Integral Differential

QCW Quasi Continuous Wave

QWP Quarter Wave Plate

RF Radio Frequency

RTP Room temperature and pressure

SPIDER Spectral Phase Interferometry for Direct Electric-field Reconstruction

TTL Transistor-Transistor Logic

Chapter 1

Introduction

1.1 Accelerators

Acceleration of particles has found broad application in today's world. From a medical perspective, this ranges from the creation of X-Rays for medical diagnostics by firing electrons at an anode, to delivering high-speed particles into patients to help to kill cancer. Industrial applications range from the use of particle beams for ion implantation for making semiconductors to electron beams used to weld, machine and even form metals into complicated shapes using a 3D printer. High-energy physics has made extensive use of these machines also, culminating with the large hadron collider at the Conseil Européen pour la Recherche Nucléaire (CERN). The large hadron collider enables particles to collide with a centre-of-mass energy reaching 13 TeV [1]. It is hoped that these high energies will enable the discovery of physics beyond the realms of the standard model, which may help to explain the presence, for example, of dark matter in the universe.

The large hadron collider is an accelerator known as a synchrotron. The circular construction of this accelerator allows the particles effectively to see an infinitely-long accelerator. The circular design also enables the radio frequency (RF) cavities used to accelerate the particles to be used over and over again on the same particle bunch. Linear accelerators are less complicated in their design, but are single pass systems and their finite length limits the maximum achievable beam energy. Schematic examples of a linear accelerator and a synchrotron can be seen in Figure 1.1

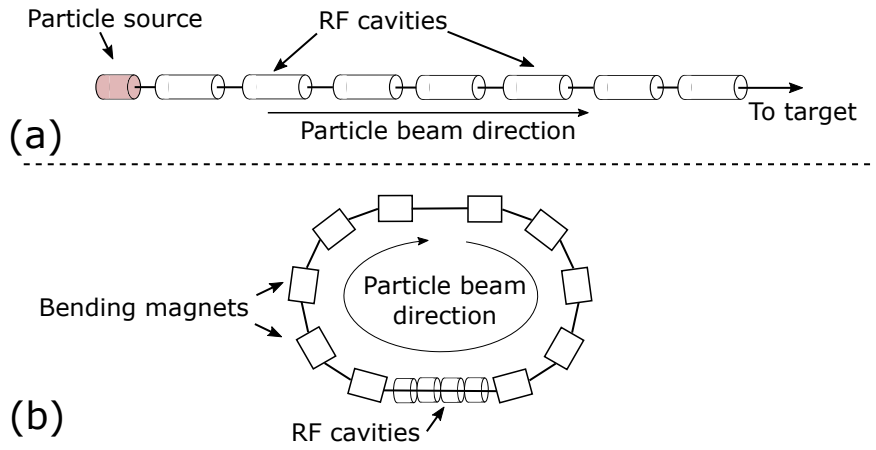


Figure 1.1: A comparison of (a) a linear accelerator and (b) a synchrotron. The linear accelerator has a finite length and a simpler design whereas the synchrotron has an effectively infinite length at the expense of more complexity. This complexity comes from the extra magnetic dipoles and quadrupoles needed to steer and focus the beam as it propagates around the synchrotron.

The circular nature of a synchrotron trades this simplicity for greater beam energy, and so requires additional dipole and quadrupole magnets to help to steer and focus the beam. However, even synchrotrons have limits; the first limitation is the strength of the magnets required to keep the beam on its circular trajectory. The magnetic rigidity [2] is described by an equation which relates the maximum beam energy to the strength of the dipole magnets and the size of the ring. It is defined as

$$\mathrm{B}\rho = 3.3356 \times \frac{\mathrm{P}}{\mathrm{Q}}\ \mathrm{Tm}, \tag{1.1}$$

where B is the dipole magnet's field strength, ρ is the bending radius of the ring, P is the particle's momentum given in GeV/c and Q is the particle's charge in coulombs. It can be seen from Equation (1.1) that by increasing the magnetic field, the particle energy for a given ring size can also be increased. However, the magnetic field strength is limited by the materials used for magnet construction. So to reach higher particle energies, given the limits imposed by the magnetic field strength, the bending radius of the ring will have to increase. Hence, the proposals for the future circular collider (electron collider, proton collider or electron-positron collider) assume a 100 km circumference ring seeded by the current large hadron collider [3, 4].

The second limitation of a synchrotron is the accelerating gradient provided by the RF cavities. An RF cavity, shown in Figure 1.2a, is a metallic cavity that is designed to cause

RF waves, delivered via a waveguide or an antenna, to resonantly excite a standing wave. For this standing wave to be effective, the direction of the electric field from the entrance to the exit of the cavity should be parallel to the intended direction of beam propagation. Attention must also be paid to the phase of the RF wave, as the electric field direction will periodically change between the forward and reverse directions. The desired mode of operation, for protons, is a forward facing electric field. Protons entering a parallel electric field will extract energy from the field while protons entering an anti-parallel electric field will lose energy to the field. Protons which have already extracted the maximum energy possible from a single cavity (a single cell) can see further energy gains if multiple RF cavities are daisy-chained together (multiple cells), as seen in Figure 1.2b. However, the direction of the electric field between adjacent cavities will not be the same. As a proton travels through the multiple cells, the direction of the electric field should be parallel with it for the duration of time taken for the proton to traverse the cavity. If the proton is still traversing the cavity when the field direction changes, then the net energy gain will be reduced.

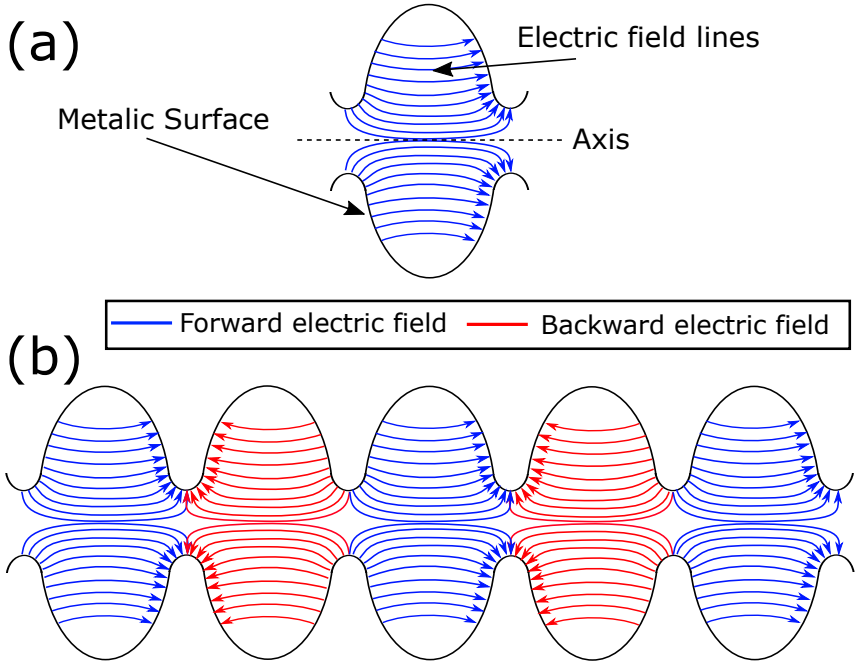


Figure 1.2: (a) A schematic cut-away view of a single RF cavity. There is rotational symmetry around the axis of the cavity. The blue lines represent the standing wave electric field established within the cavity. The standing wave can be established by use of a waveguide or an RF antenna. (b) A chain of RF cavities allows a particle to extract more energy than from a single cavity. However, the electric field alternates between pairs of RF cavities. Therefore, it is necessary to ensure that the particles arrive at the next cavity when the electric field has changed phase from the reverse to the forward direction. Failure to match properly the particles to the phase of the standing wave will cause the particles to lose energy to the electric field, rather than gain energy.

State of the art RF cavities use superconducting materials and are capable of producing accelerating gradients between 10 and 100 MeV/m. The field strengths inside these cavities are high enough for imperfections on the inside surface of these cavities to cause local enhancements of the electric field. These enhancements can result in RF breakdown and cause irreparable damage to the RF cavity [5]. This breakdown point is known as the Kilpatrick limit. The 100 MeV accelerating gradient upper limit would not be a problem if the particle beam did not suffer any losses as it propagated around the ring. However, this is not the case. Synchrotron radiation is one such loss mechanism, which is caused when charged particles are bent by a magnetic field. The power lost by synchrotron radiation is given by [2]

$$\mathrm{P}_\gamma = \frac{\mathrm{e}^2}{6\pi\epsilon_0 \mathrm{r}^2}\frac{1}{\mathrm{m}^4\mathrm{c}^7}\mathrm{U}^4\beta^4, \tag{1.2}$$

where m is the particle mass, c is the speed of light, r is the bending radius of the particle track, U is the particle energy, β is the particle velocity divided by the speed of light. The mechanism of synchrotron radiation is schematically represented in Figure 1.3. Power lost by synchrotron radiation scales as $1/\mathrm{m}^4$, U^4 and $1/\mathrm{r}^2$. The fact that heavier particles radiate less than lighter ones is one of the reasons why CERN moved away from an e^- e^+ machine (the large electron-positron collider) to one which used protons (the large hadron collider). Particle energy is also a strong driver of synchrotron radiation, with the U^4 term leading to an upper limit on the achievable particle energy in a circular accelerator. This limit stems from the power loss due to the synchrotron radiation becoming equal to the power added by the RF cavities on every turn. Increasing the bending radius can reduce the power lost to synchrotron radiation and is another reason for the proposal to build the 100 km circumference future circular collider. However, even with its $\sim$ 100 km circumference, the power lost to synchrotron radiation for the entire ring, whilst operating with a pair of counter-propagating electron beams, is estimated to be $\sim$ 100 MW (50 MW per beam at 350 GeV) [4, 6].

Synchrotron radiation was first seen as a detrimental effect that needed to be minimised, until it was realised that this could be used as an ultrafast, wavelength-tuneable light source capable of probing materials and biological samples. This led to the development of un-

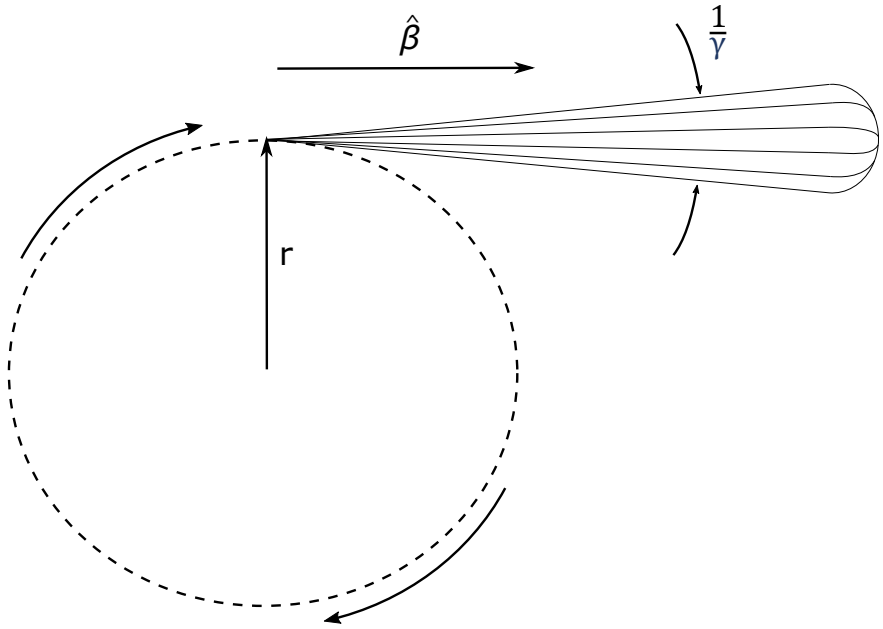


Figure 1.3: For particles travelling at speeds close to the speed of light, the synchrotron radiation is produced in a cone with an opening angle $\phi = \frac{1}{\gamma}$. Here $\hat{\beta}$ is the direction of the particle bunch and r is the bending radius of the particle bunch.

dulators and wigglers. These consist of an array of magnets of alternating polarity, which produce a sinusoidally-varying magnetic field along the direction of propagation. The sinusoidal magnetic field seen by the particles causes them to follow a sinusoidal trajectory through the device. These oscillations enable substantial increases in the brightness of the radiation emitted.

Synchrotron radiation is of interest to many different fields of research. Some of these uses are ultrafast imaging of chemical reactions, crystallography and medical diagnostics. The advantage of synchrotron radiation in medical imaging comes from its point-like source which enables higher-brightness X-Rays. Radiation emitted from a point-like source can be used to create images with higher resolutions, such as in techniques like phase contrast imaging. In addition, the higher source-brightness means a lower overall X-Ray dose is delivered to the patient during the imaging process, in comparison to that of standard cold-anode devices. The major hurdle preventing the mass proliferation of this technology is the size and cost. A typical synchrotron used for a light source is about the size of a hospital and, with space being a precious commodity in some population-dense regions, it is infeasible to build such an accelerator on site. To enable the size of a synchrotron to be reduced, stronger magnets are required to contain the beam in a smaller ring as per Equation (1.1).

However, as mentioned previously, material limitations prevent the required field strengths from being realised, which puts a lower limit on the size of a synchrotron-based device. A linear accelerator can be used to bypass the limitation imposed by the magnetic rigidity, but again, the space required becomes a factor. RF cavities cannot provide enough energy in one pass and in a short-enough distance without breaking down. Rather than working to avoid this material breakdown, one solution is to embrace it and use a material which has already been broken down, i.e. a plasma.

1.2 Laser plasma wakefield acceleration

A plasma is a gas that has been ionised. Ionisation can be achieved by supplying a large quantity of energy either through electrical discharge, a particle beam or a laser pulse. In all of these cases, the energy delivered to a gas initially causes the ionisation. Once the plasma has been formed, the free plasma-electrons are then subject to forces caused by the driver. If an electron bunch is the driver, its electric charge causes repulsion, while in the case of the laser, the ponderomotive force [7] displaces plasma-electrons out of its path. Both of these cause the plasma-electrons to oscillate in a Langmuir wave resulting in a wakefield being formed behind the driver. A simulation of a laser-driven wakefield can be seen in Figure 1.4.

To save confusion between electrons being accelerated and those that exist within the plasma, plasma-electrons will be the term used to describe the electrons which exist naturally within the plasma. Terms such as accelerated electrons or electron bunch will be used when describing those electrons undergoing acceleration within the wakefield. From Figure 1.4, the displacement of plasma-electrons by the laser pulse produces a region in the wake of the laser pulse which has an electron density lower than that of the plasma density. Around this void of plasma-electrons is an electron shell with an electron density higher than the plasma density. Between the area of highest electron density (white) and lowest electron density (black), there exists an electric field. The similarities between the wakefield structure and RF cavities are shown in Figure 1.5.

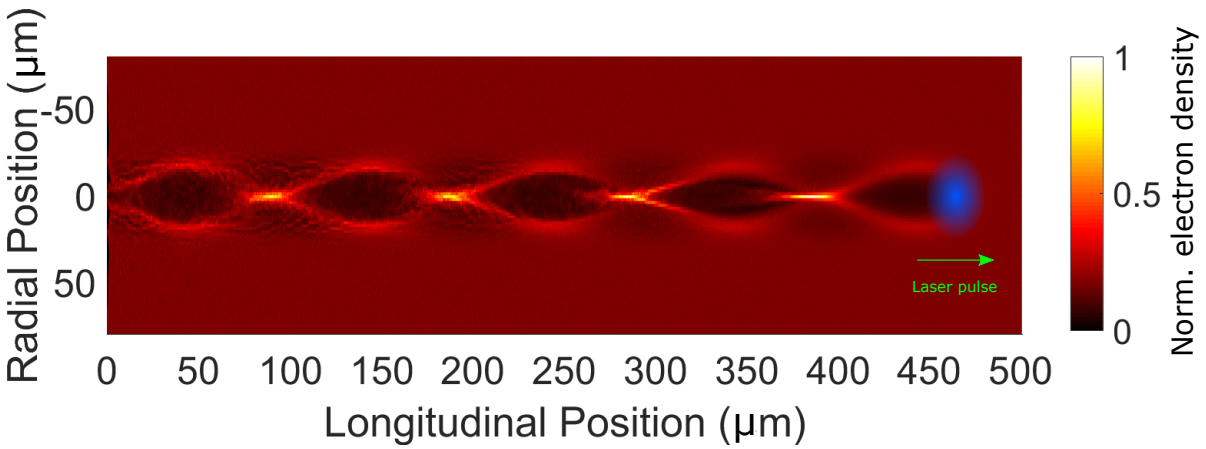


Figure 1.4: A particle in cell simulation of a laser pulse (blue) travelling left to right and driving a wakefield by displacing the plasma-electrons in its path. The frame of reference shown is one co-propagating with the laser pulse. The colour map describes the normalised electron density in the plasma.

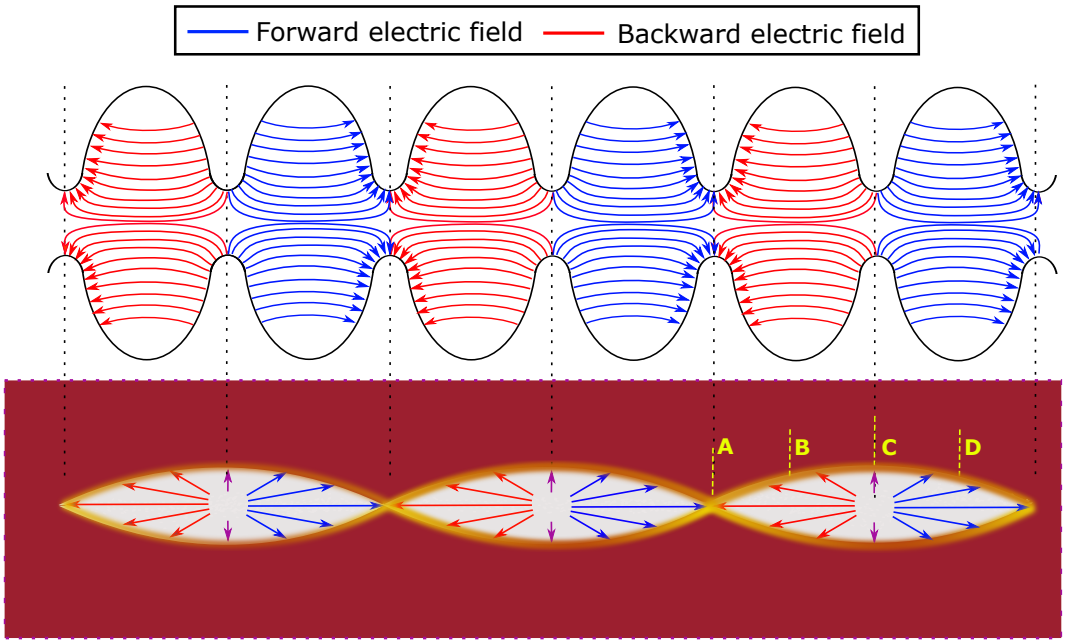


Figure 1.5: The comparison between a multiple cell RF cavity and a wakefield. Blue arrows represent forward facing electric fields while red arrows represent backwards facing electric fields. The wakefield has some similarities with RF cavities, including the need to place the particles to be accelerated in the correct phase to see acceleration. To accelerate electrons using the wakefield, the electron bunch needs to be placed at longitudinal position B to experience acceleration. An electron bunch positioned at longitudinal position C will feel no net force while a particle bunch positioned at longitudinal position D will experience deceleration.

As with an RF cavity, particles need to be in the correct position within the wakefield to be accelerated. Referring to Figure 1.5, an electron bunch should be placed at position B in order to feel a net accelerating force. An electron bunch found at position C will feel no net force and an electron bunch positioned at position D in the wakefield will feel a decelerating force. The three major differences between an RF cavity and a wakefield are: firstly, the wakefield structure is made up of a broken-down material; secondly, only negatively-charged

particles will feel a focusing force because of the direction of the electric field vectors; and thirdly, the accelerating gradient is much larger. Between the point of highest electron density (position A) and the centre of the void (position C), there can exist an accelerating gradient up to 3 orders of magnitude greater than that seen in standard RF cavities. This large accelerating gradient has been successfully used to demonstrate the acceleration of electrons. Electron bunches have been reported at giga-electron-volt energies from a single plasma stage a few centimetres long [8, 9, 10] with electric field strengths of $E_0 \approx 100$ GeV/m [7]. With the large accelerating gradients reported for laser-plasma wakefield accelerators, they are very suitable for producing high energy electron beams from very short length plasma cells. This has motivated further research from many different groups into laser wakefield acceleration. From here onwards, the discussion regarding wakefields will be limited to that of laser-based drivers. Wakefields using proton and electron beam drivers (beam driven plasma wakefield acceleration) have been demonstrated and are active fields of research by, for example, the AWAKE group (proton) at CERN [11] and at FACET (electron) at the Stanford linear accelerator (SLAC), amongst others.

An accelerator that can take advantage of these large accelerating gradients has the potential to be built at a reduced size compared to conventional accelerators, with such structures being only millimetres in length. The current state of the art for laser plasma accelerators, as taken from many different experiments and for many different aspects are: electron energies of 4.2 GeV [9], electron bunch charge of 620 pC [12], small energy spreads of $\sim 5 - 10\%$ [9, 13], normalised transverse beam emittances of $0.15 \pm 0.06\pi$ mm mrad [14] and $0.21 \pm 0.01\pi$ mm mrad [15] and a plasma cell length of only 4 mm producing 1 GeV beams [10].

When the staging of many plasma cells is considered, it opens up the possibility of making a staged laser-wakefield-acceleration based linear-accelerator (linac). This would enable electron beams to reach the large energies desired by high energy physics experiments, while being relatively immune to the effects of synchrotron radiation. Given the 100 GeV/m gradients obtained with laser wakefield acceleration, it could be possible to reach TeV beam energies in a matter of tens of meters, rather than tens of kilometres. One proposal has

presented a concept 1 TeV linear collider, which aimed to utilise laser wakefield acceleration to allow a machine length of ∼100 m [16]. For all the perceived benefits of laser wakefield acceleration, there are some limitations to its possible applications. Firstly, this technology has only been demonstrated with electrons and positrons. Although hadron accelerators have been demonstrated by hitting a solid target with a laser pulse, hadron accelerators are yet to be demonstrated using wakefields. This is because the phase velocity of the wakefield using a single laser pulse is much higher than the initial hadron velocity. Proton acceleration has been proposed using a slow wakefield [17]. Until more work is invested into hadron acceleration, plasma wakefield accelerators are unlikely to compete with conventional accelerators in this field. Also, the charge yield is quite small compared to that seen in conventional accelerators, typically an order of magnitude lower. Hence, in industry, more conventional methods of producing high-intensity beams of electrons for material processing and fabrication remain preferable. However, to realise laser wakefield acceleration, the limiting factor becomes the laser technology used to drive a wakefield. Here a laser's peak power and repetition rate must be considered.

To drive a wakefield in a plasma, the density of the plasma-electrons must be tailored to suit the pulse duration of the laser. In order to generate a wakefield with the largest axial electric field, the plasma period, T_p should be $T_p = 2.7T_{FWHM}$ (where T_{FWHM} is the measured laser pulse duration, assuming it to be Gaussian [7]). Plasma periods different to this will still produce a wakefield capable of acceleration, but the axial electric field will not be as large as for the optimised case. The plasma period is related to the plasma frequency, which is given by [7]

$$\frac{2\pi}{T_p} = \omega_p = \sqrt{\frac{n_e e^2}{m_e \epsilon_0}}, \tag{1.3}$$

where n_e is the electron number density, e is the electric charge, m_e is the electron mass and ϵ_0 is the permittivity of free space. From Equation (1.3), it can be seen that e, m_e and ϵ_0 are constants, leaving n_e as the only free parameter. What Equation (1.3) tells us, is that for a given laser pulse duration, there is a corresponding 'optimal' plasma electron number density. Therefore, when choosing the density of the plasma to work at for a laser-wakefield acceleration experiment, the duration of the driving laser pulse is a key parameter.

Using Equation (1.3), the electron density required can be calculated by substituting in the condition for the largest axial electric field

$$n_e = \left(\frac{2\pi}{2.7T_{FWHM}}\right)^2 \frac{m_e \epsilon_0}{e^2}, \tag{1.4}$$

which for a typical 50 fs laser pulse produced by a Ti:Sapphire laser, gives an electron density of $6.8 \times 10^{17} cm^{-3}$. For comparison at room temperature and pressure, a gas would have a number density of $10^{21} cm^{-3}$, four orders of magnitude higher. The wakefield is established when the laser pulse propagates through the plasma and displaces the plasma-electrons, causing them to oscillate behind the pulse in a Langmuir wave. This displacement is driven by the ponderomotive force, given by [18]

$$F_p = -\frac{e^2}{4m_e\omega^2}\nabla|E|^2, \tag{1.5}$$

where e is the electric charge, m_e is the electron mass, ω is the oscillation frequency of the laser and E is the amplitude of the electric field. The ponderomotive force is the time-averaged force produced by a charged particle interacting with an inhomogeneous oscillating electric field. Quantitatively, Equation (1.5) is derived by Esarey et al. [7] and so won't be repeated here.

When a charged particle sees a homogeneous oscillating electric field of frequency ω, the charged particle will also begin to oscillate at ω. The charged particle will be displaced away from and then return to its original position in one full cycle. For an inhomogeneous oscillating electric field of frequency ω, during the first half of the plasma electron's oscillation cycle, it sees a stronger electric field amplitude in the direction of the weaker field amplitude. In the second half of the plasma electron's oscillation cycle, the plasma electron sees a weaker electric field amplitude pointing in the direction of the stronger electric field amplitude. The time-average over a single plasma electron cycle produces a net force pointing in the direction of the weaker field gradient. A simplified example is demonstrated in Figure 1.6.

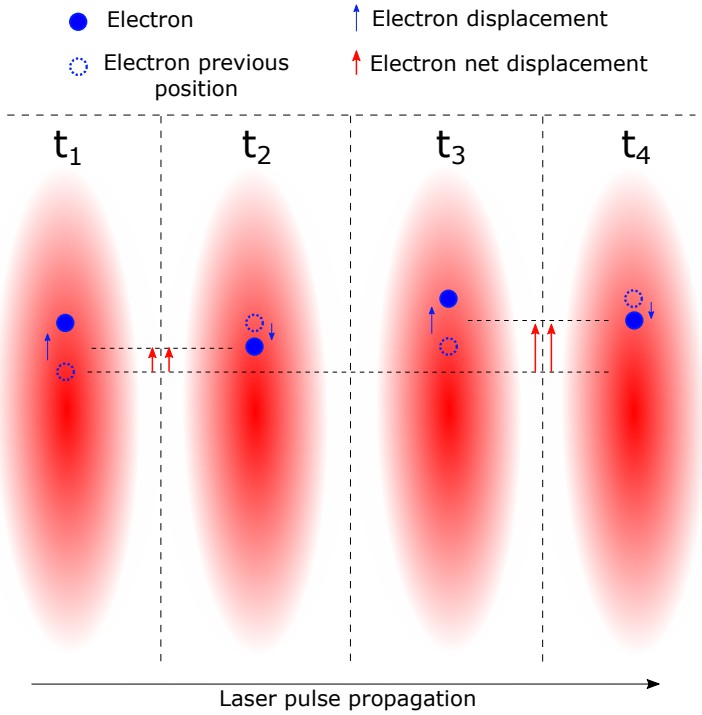


Figure 1.6: A simplified example to demonstrate the mechanism of the ponderomotive force. The shaded region represents the laser pulse and its spatial intensity distribution. In time step t_1 during the first half of the plasma-electron's oscillation cycle, it undergoes a displacement away from the laser pulse's axis of propagation proportional to the field amplitude. This displacement occurs in the direction of the negative gradient. In the second half of the plasma-electron's oscillation cycle, t_2, it sees a weak amplitude field directed towards the positive gradient, resulting in a smaller displacement back towards its original position. Another complete oscillation of the plasma-electron between time steps t_3 and t_4 leads to a displacement further away from the original propagation axis.

The laser pulse propagating through the plasma pushes the plasma-electrons out of the way with a force proportional to the gradient of the intensity, as per Equation (1.5). Steeper gradients will produce larger forces and so, to drive a wakefield, short, intense laser pulses should be used. This is achievable easily with current Ti:Sapphire lasers. Once the laser pulse has passed, the displaced plasma-electrons see a restoring force back towards the laser's axis of propagation. These plasma-electrons undergo oscillations at the plasma frequency, which establishes the Langmiur wave. Note that $\omega_p << \omega_l$, where ω_l is the frequency of the laser driver and ω_p is the plasma frequency. For linearly polarised beams, a useful parameter which describes the power of a laser pulse is a_0, the laser strength parameter. It is approximated as

$$a_0 \approx 0.85 \times 10^{-9} \lambda(\mu m) \sqrt{I(W/cm^{-2})}, \tag{1.6}$$

where $\lambda(\mu m)$ is the laser wavelength defined in μm and $I(W/cm^{-2})$ is the laser pulse peak intensity defined in cgs units. Cgs units are the standard way of presenting numbers in

the field of laser plasma wakefield acceleration and will appear frequently throughout the associated literature. A wakefield can be generated using a laser pulse with an $a_0 < 1$ and it is said to be in the linear regime. Higher a_0 values, typically $a_0 > 1$ [7], lead to non-linear wakefields and eventually the "bubble regime" or "blowout regime", where the laser pulse has sufficient intensity to expel all of the plasma-electrons from its path. In this highly non-linear regime, some of the plasma-electrons can become trapped in the longitudinal accelerating gradient of the wakefield and get accelerated as a bunch. Hence, a laser wakefield acceleration linac could be designed to seed itself rather than relying upon external injection from a conventional accelerator. External injection though feasible requires synchronisation between the laser system and the externally injected electron bunches. The advantage of self-injection is that the electron bunch created is synchronised with the laser pulse. Other injection schemes exist, such as colliding pulse injection [19, 20], ionisation injection [21, 22, 23, 24] and density gradient injection [25], but these will not be elaborated upon.

In summary, Equation (1.3) tells us that for a given laser pulse duration, there is a corresponding 'optimal' plasma frequency that enables the most efficient generation of a wakefield. To generate the largest on-axis electric field within the wakefield for a given intensity [7], laser pulses should be short ($T_p = 2.7T_{FWHM}$). For sufficient displacement of plasma-electrons to establish a wakefield useful for acceleration, the laser pulse should be intense to provide a large ponderomotive force. The intensity of a pulse can be described by the laser-strength parameter with non-linear wakefields being established for $a_0 \geq 1$. High-peak-intensity pulses can be produced given higher peak powers, assuming the focused beam area is the same.

1.3 Laser wakefield acceleration with terawatt class lasers

The lasers used to drive laser wakefield acceleration are typically terawatt class Ti:Sapphire ultrafast lasers, delivering energies in the range of milli-joules to hundreds of joules, in short pulse durations of ∼50 fs. An ultrafast laser is one where the laser energy arrives in discrete pulses, with pulse durations of picoseconds and below. The current state of the art for short pulse duration is 67 as [26], produced in the XUV. Ti:Sapphire systems have demonstrated pulses down to 6 fs [27]. To produce laser pulses with short pulse durations, the laser pulse

must be made up of a collection of laser frequencies, known as its spectrum. The limiting relationship between the pulse duration and its spectrum is described by the Fourier relationship between time and frequency. Assuming a Gaussian spectrum, the Fourier transform of a Gaussian is another Gaussian. Larger-width Gaussians in the spectral domain can transform into smaller-width Gaussians in the time domain. If all of the frequency components are travelling with a constant spectral phase with respect to the central frequency of the Gaussian spectrum, then the pulse can be said to have a "flat phase". This means that the group delay, the first derivative of the phase with respect to angular frequency, is zero. If a frequency-dependent spectral phase is introduced then the group delay is no longer zero, and the frequency components will arrive at different times. This results in a broadening of the pulse duration. A pulse with a flat spectral phase is the shortest pulse that can be generated from the given spectrum and is known as the "transform limit".

Ultrashort lasers with large peak powers and short pulse durations are suitable for providing the large gradients required for the ponderomotive force to displace electrons. The current state of the art for high-peak-power laser pulses using Ti:Sapphire technology is 4 PW [28], with other laser systems producing peak powers of 1 PW [9, 29]. Although these Ti:Sapphire laser systems have shown already that they are capable of driving a wakefield and accelerating electrons, they have low efficiencies and repetition rates. To reach these peak powers, many amplifiers are used, and each is pumped by either diode lasers or other pump lasers which themselves are pumped with flash lamps. Each of these stages are liable to heat build-up due to the quantum defect, where the energy difference between the pumping and lasing wavelengths causes heat build-up in the amplifier. This heat must be removed before the laser can be fired again, or thermal damage will occur. Ti:Sapphire lasers typically have a repetition rate in the range of 0.1-10 Hz with wall-plug efficiencies as high as 1%, but it is common for efficiencies to be half this or less [30]. Repetition rates can exceed 10 Hz if the system is made to produce smaller peak powers. Laser wakefield acceleration has been demonstrated with a Ti:Sapphire laser running at 1 kHz [31] but with pulse energies of approximately 1-10 mJ, orders of magnitude less than the 10-100 J used for petawatt-class lasers. At low repetition rates, the average power of these systems is also going to be low, typically in the 1-10 W region.

One reason for the interest in laser wakefield acceleration is for potential use in a 'tabletop' sized particle accelerator for medical or scientific applications. The hurdles which currently face Ti:Sapphire systems are their low repetition rates and low wall plug efficiencies. The low repetition rate prevents a medical system from integrating over many shots to build up a dose or image a patient. Meanwhile, the efficiency has a significant impact on the electricity bill required to run such a system and the heat built up in the amplifiers, which must be removed. Given the potential of laser wakefield acceleration to proliferate medical accelerators to hospitals everywhere, the challenges facing Ti:Sapphire lasers need to be overcome. For a laser wakefield accelerator system to be useful, it should have a repetition rate of $\sim$ kilohertz and wall plug efficiencies of tens of percent. One solution, is to make use of fibre lasers to bypass the limitations of Ti:Sapphire systems.

1.4 Fibre lasers

Fibre lasers have seen rapid gains in terms of performance and average power. Fibres are waveguides designed to channel infra-red and optical radiation. The most common type of fibre available is known as a step-index fibre. A step-index fibre comprises a core with a cladding material wrapped around it, each having different refractive indices. The difference that exists between the larger core-refractive-index and the slightly smaller cladding-refractive-index leads to total internal reflection of radiation propagating through the core. The transition angle between refraction and reflection is known as the critical angle and is taken with respect to the normal from the core wall. Angles smaller than the critical angle will lead to refraction of the radiation out of the core. Angles larger than the critical angle will cause radiation to undergo total internal reflection. Doping the core of the fibre with a rare-earth ion allows it to act both as a waveguide and a laser amplifier. Typically, the rare-earth ions in the fibre are excited to the metastable state using a pump laser, and the energy is liberated using a seed laser. In this thesis, when a 'pump' is referred to it implies a high average power CW laser. A 'seed' refers to the initial laser pulse prior to amplification and an 'amplified seed' refers to a laser pulse which has previously undergone amplification.

When classifying a fibre, a key characteristic is its core diameter. The core diameter indicates the cross-sectional area of the core and, subsequently, the volume of dopant available

when the fibre's length is known. The Mode Field Diameter (MFD) indicates the beam size that can be supported by a given core diameter. To convert from core diameter to MFD for step-index fibres, the following approximation is used,

$$\mathrm{MFD} = 2\mathrm{a}(0.65 + 1.619\mathrm{V}^{-3/2} + 2.879\mathrm{V}^{-6}), \tag{1.7}$$

where MFD is the Mode Field Diameter, a is the core radius and V is known as the V-parameter, which is defined as

$$\mathrm{V} = \frac{2\pi\mathrm{a}}{\lambda}\mathrm{NA}, \tag{1.8}$$

where a is again the core radius, λ is the wavelength of the laser and NA is the numerical aperture of the fibre's core, defined as $\mathrm{NA} = \sqrt{\mathrm{n}_{\mathrm{core}}^2 - \mathrm{n}_{\mathrm{cladding}}^2} = \sin(\theta)$ where $\mathrm{n}_{\mathrm{core}}$ is the refractive index of the core, $\mathrm{n}_{\mathrm{cladding}}$ is the refractive index of the cladding and θ is the maximum angle of incidence for a ray to be guided, which is taken with respect to the normal at the surface of the core. Larger-core diameters reduce the intensity of the light propagating through them, and reduce the risk of material damage, making them more suitable for applications requiring large peak powers. The modes that a fibre can support are also important to consider. Discussion on fibre modes is performed in more detail in Chapter 2; for now, it can be said that the fundamental mode is the preferred mode of operation. This mode looks like a J_0 Bessel function rotated 2π around the axis, as shown in Figure 1.7. Different designs of fibre allow larger core diameters to be realised, while still supporting only the fundamental mode. Fibres that support only one mode, such as the fundamental mode, are said to be single mode. Fibres which additionally support higher-order modes are known as multi-mode fibres. For step-index fibres, the refractive index difference between the core and cladding needs to be very small in order to ensure that the numerical aperture of the fibre is small. Smaller numerical apertures are necessary to ensure that the fibre only supports the fundamental mode.

Large-core-diameter step-index-fibres available for purchase typically have a numerical aperture of ~ 0.06 or greater. The inability to manufacture small differences between the refractive index of the core and cladding was once the largest obstacle preventing further

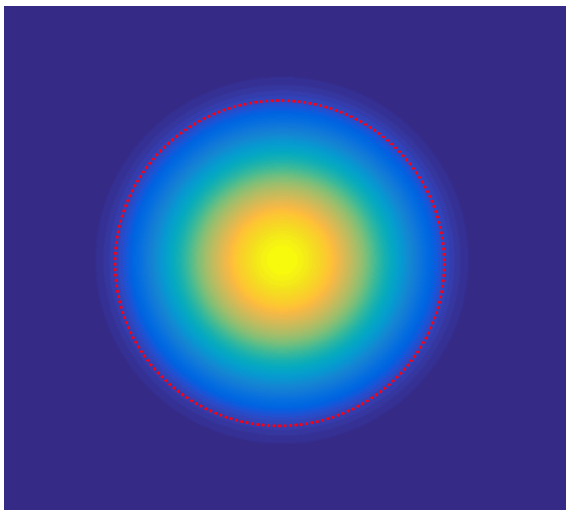


Figure 1.7: An example of the intensity distribution of the fundamental mode of a circular step index fibre. The beam within the fibre core has a Bessel profile, which meets the core wall and becomes an exponential tail which decays away. The red dotted line shows the edge of the core.

increases in core diameter using a step-index design. Now, with better technology to control the refractive indices of materials, the effects of losses such as those due to fibre bending become significant. The act of bending a fibre causes a large enough change in the refractive index between the core and cladding to prevent total internal reflection. This break down removes the advantage of a step-index fibre for coiling large quantities of fibre into a small space. For larger single mode fibres, less bending can be tolerated. A different design of fibre makes use of a photonic crystal structure to channel the radiation into the fibre core. Photonic crystal fibres (PCF) with core diameters an order of magnitude larger than step-index fibres have been demonstrated, with state of the art at the time of writing being 135 μm [32], operating at a wavelength of $\lambda = 1032$ nm. The analysis which is used to determine the MFD for a step-index fibre cannot be applied to PCFs due to their different structures. Calculating the MFD for PCFs is not trivial and is generally done using waveguide mode solver software. PCFs can be doped also with rare-earth ions to turn them into fibre amplifiers. The state of the art, amplified-seed pulse energy for a single PCF is 2.2 mJ, and this produces a compressed-pulse peak-power of 3.8 GW from a 480 fs pulse [33]. More discussion about PCFs will be presented in Chapter 2.

Fibre lasers can operate at MHz repetition rates and are capable of delivering high average powers. Their wall-plug efficiencies are also very high; with slope efficiencies (output-power increase for given pump-power increase) of the range of $60\% - 80\%$ [34, 35, 36]. When

pumped using diode lasers (now reported to possess wall plug efficiencies of over 70% [37]) the wall plug efficiency of a fibre laser system can be brought up to $\sim 40\%$ which is an order of magnitude above that of the most efficient Ti:Sapphire lasers. Higher wall plug efficiencies can lead to large savings in electricity costs. The one drawback that has prevented this architecture dominating over Ti:Sapphire lasers is the low peak power of the pulses. The state of the art pulse energy extracted from a single fibre laser is 2.2 mJ, producing a peak power of 3.8 GW from a pulse duration of 480 fs [33]. For the peak power, the state of the art is reported to be 7.8 GW for a single PCF [38] using a pulse duration of 200 fs, but a more modest pulse energy of 1.6 mJ. These reported results used a Yb^{3+}-doped-silica, large-pitch PCF, which had a large core diameter and a bandwidth theoretically capable of supporting pulses of 50 fs and below, assuming a spectral pulse FWHM bandwidth of 30 nm. Assuming 2.2 mJ of energy was put into a 50 fs Gaussian pulse and all of that energy was located under the central peak, the peak power generated would be 44 GW. The intensity required to generate a linear wakefield (for example, $a_0 = 0.75$) using a fibre laser operating at a wavelength of 1 μm and using Equation (1.6), would be $7.7 \times 10^{17}\,\mathrm{W/cm^2}$. If this pulse is focused to a beam diameter of 40μm (assuming a diffraction limited beam with an initial diameter of 30 mm for use with a standard 3" mirror), then the required peak power to give this intensity is $\sim 5\,\mathrm{TW}$.

For fibre lasers with 5 TW of peak power, the pulse energy required to make this peak power, given a pulse duration of 50 fs is $\sim 260\,\mathrm{mJ}$. The 50 fs pulse duration is a fundamental limit of the Yb^{3+} gain medium as it is not able to support larger bandwidths. This means that to make improvements upon the peak power figure using a Yb^{3+} gain medium, generally work is invested in increasing the pulse energy. Non-linear methods of pulse compression also exist, which take advantage of the intensity-dependent non-linear effects produced by pulses with high intensities in order to generate bandwidth. Increasing a pulse's bandwidth enables it to support shorter pulses, which can subsequently increase the pulse's peak power, despite the process being lossy. However, this is an active field of research, and more detail is provided in Chapter 10. Given that non-linear compression is not performed in this thesis, the goal of improving the peak power is focussed primarily on increasing the pulse energy, while maintaining the input pulse duration. Non-linear effects are detrimental to the pulse

duration and so should be minimised, or better, avoided. Non-linear effects are discussed in more detail shortly during the discussion of spectral phase.

For fibre laser technology, the current, highest pulse energy of 2.2 mJ from a single fibre amplifier will need amplification by two orders of magnitude before it reaches terawatt levels of peak power. Simulations [39] have shown that the maximum peak power a pulse could possess after amplification in a single fibre amplifier with a Mode Field Diameter (MFD) of 80 μm using a Chirped Pulse Amplification (CPA) scheme was $\sim$ 10 GW. CPA is a method of enabling greater amplification of a laser pulse by mitigating the effects of non-linearities. Chirping is the expression used to describe the temporal spreading of frequency components such that the frequency of radiation seen by an observer in the path of the pulse increases or decreases with time. Chirping a pulse with a finite bandwidth spreads the different frequency components of a laser pulse temporally. The observed effect of this (should an observer witness a positively chirped pulse with a spectrum centred in the visible) is that the laser pulse's colour would change colour from red, to green and then to blue. A pulse can be chirped by forcing the frequency components to travel different distances to one another. The temporal spreading of the frequency components reduces the pulse's peak intensity which in turn reduces the impact of non-linear effects. The temporally-spread laser pulse is said to be 'stretched' or 'chirped'. By acquiring lower quantities of unwanted non-linear effects, greater quantities of energy can be liberated from an amplifier before these non-linear effects begin to manifest themselves. Non-linear effects include self-focusing, self-steepening and intra-pulse Raman scattering. Once a pulse has been through the amplifier section of the laser system, the final part of CPA is compression. During compression, the spreading effect which reduced the peak intensity is undone by causing the frequency components to travel different distances. However, this time the distances travelled by the frequency components are inverted, such that those who travelled further distances before, now travel shorter distances and those that travelled shorter distances before are made to travel longer distances. After the laser pulse has been through the compressor, all of the frequency components should have travelled the same distance through the entire laser system. Providing non-linear effects were not experienced, the frequency components should have the same frequency-independent phase relationship after compression that they did before stretching. The net effect is that the pulse duration

of the stretched pulses reduces towards the transform limit, causing the peak intensity to see a rapid increase. A schematic diagram of a CPA system is shown in Figure 1.8.

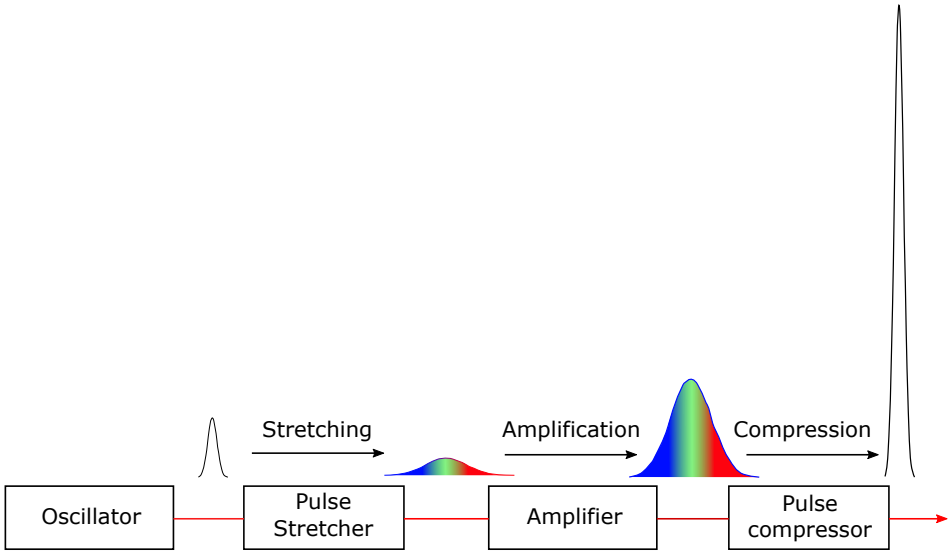


Figure 1.8: A schematic diagram of a chirped pulse amplification (CPA) system. An ultrafast laser pulse is stretched using a pulse stretcher, chirping the pulse. The chirped pulse is amplified before undergoing compression. The result is an ultrafast laser pulse with a similar pulse duration to the initial seed pulse but with a greatly increased peak power.

Although more energy can be extracted from the fibre by stretching the pulses by larger factors, it will reach a point where CPA will not provide further improvement to the peak power of a laser pulse. The main components used to make a stretcher and compressor (typically diffraction gratings are used for high power laser systems) cannot be made arbitrarily large as the cost of manufacture increases non-linearly with component size. Although coils of passive fibre can be used to achieve large stretch factors by making use of material dispersion, the aim of CPA is to match the stretcher and compressor, so that the effects of the stretcher are undone by the compressor. When compressing pulses, coils of passive fibre cannot be used as the high pulse intensities due to small core sizes could either cause damage to the fibre or will drive undesirable non-linear processes. This constrains compression to the use of a grating-pulse-compressor with size-limited components. The limit on component size provides a limit on the amount of compression that can be achieved, which ultimately influences the amount of stretching that can be applied to a laser pulse in the first place. The limit on the stretch-factor dictates the minimum intensity a given pulse can be reduced to before it

is amplified by a fibre amplifier. Given the upper peak power limit of $\sim$ 10 GW [39] for a fibre with an 80 μm MFD, fibre amplifiers cannot be made arbitrarily long to keep adding energy to improve the peak power. Additionally, non-linear effects such as self-focusing, can cause damage to the fibre's core-material, so generally, it is preferred to avoid large pulse intensities.

Another potential solution is to increase the transverse size of the amplifier to get more gain-medium, but the core diameters of the fibres cannot be made arbitrarily large due to the constraints of guiding the fundamental-mode alone, otherwise known as single mode operation. A workaround to this limitation is to make the fibre operate in 'effective-single-mode' operation. Effective single mode operation is where the design of Photonic Crystal Fibre (PCF) is capable of supporting higher order modes, but strategies to suppress higher order modes have been implemented. These strategies include forcing higher order modes to see greater loss than the fundamental, as they propagate through the fibre, and to position gain in the fibre to preferentially amplify the fundamental mode over the higher order modes. By designing fibres to operate in effective single mode operation, the core diameter can be increased further while the fibre still produces the fundamental mode. Effective single mode operation is an active field of research, which has led to the development of a different design of PCF called large pitch fibres. Large pitch fibres with MFDs of 100 μm [32] have been demonstrated and provide one route to surpassing the 10 GW peak power limit of 80 μm MFD PCFs. However, even these increases are not expected to lead to the development of millimetre-scale core diameters and so there remains a fundamental limit on the core diameter of a PCF. The contradiction of requiring larger core diameters for more energy, but maintaining the fundamental mode purity by keeping the core diameters small is a challenge for fibre based lasers systems.

One solution to the problem posed by requiring more gain-medium is to split the gain-medium across multiple fibres. The laser system architecture would be an array of fibre amplifiers fed by a laser pulse from a common source. The laser pulse is split and a copy of the pulse directed into each amplifier, where they are amplified in parallel. After amplification, the many amplified pulse replicas are recombined back into a single pulse, which

is then finally compressed using a pulse compressor. Splitting up a single pulse into many replicas and extracting energy from multiple fibres will lower the peak intensity of the pulse in each fibre and will result in a reduction of the non-linear effects experienced by each pulse. The most challenging aspect of this architecture is using fibre amplifiers. Generally, fibre lasers are preferred to operate in the fundamental mode, which makes working with the beam predictable in terms of its beam parameters. To excite the fundamental mode, the laser pulse needs to be correctly aligned to the core of the fibre and the seed laser pulse needs to have the correct fundamental-mode spatial-distribution. Aligning a beam to a fibre core requires micrometre level precision, and if the setup has not been connectorised, then precision alignment using mirrors and lenses must be performed. If the beam was to enter the core at a small angle to this axis, two things could happen: firstly, some of the beam could be refracted out of the core if it no longer satisfies the condition for total internal reflection; secondly, there will be excitation of the higher-order modes, should the fibre support them. A beam supporting multiple modes makes its intensity profile hard to predict. The probability that the beams from any two PCFs would share the same intensity distribution when exciting higher-order modes is remote. Matching the spatial-distributions becomes a problem when combining the beams as mismatched intensity profiles can result in a spatially varying polarisation state across the beam. So, it is preferable to use the fundamental mode when combining beams.

The generation of multiple pulse replicas is addressed by placing a series of beam splitters in the beam. Beam splitters are optical components designed to split an incident beam between a number of output ports. This can be done using the many diffracted orders from diffraction gratings, reflections from optical surfaces, or by using a laser's intrinsic polarisation. Unpolarised beam splitters are designed to split a beam by a specified ratio (10:90 or 30:70 for example) with the majority of the beam power being transmitted. Polarising beam splitters either: split an unpolarised beam into two orthogonal polarisations with a balance of 50:50 between the two output ports; or split a polarised beam between two ports with a variable ratio, depending on the incident beam's polarisation state with respect to the polarising beam splitter's transmission axis. Rotating an incident beam's polarisation state from P-polarised to S-polarised, using a Half Wave Plate (HWP) will see the output power

from the polarising beam splitter shift from one output port to the other. Adjusting the incident beam's polarisation state to 45 degrees from both S- and P-polarisation states will produce two beams of equal power, one emitted from each of the polarising beam splitter's output ports with S- and P-polarisation respectively. This equal splitting can be useful for creating a pair of pulse replicas with equal pulse energies. Further splitting can be performed by using a daisy-chain of beam splitters to create multiple pulse replicas. Before each splitting stage, the pulse must be first rotated to 45 degrees to achieve equal power splitting between the polarising beam splitter's output ports. This is demonstrated in Figure 1.9.

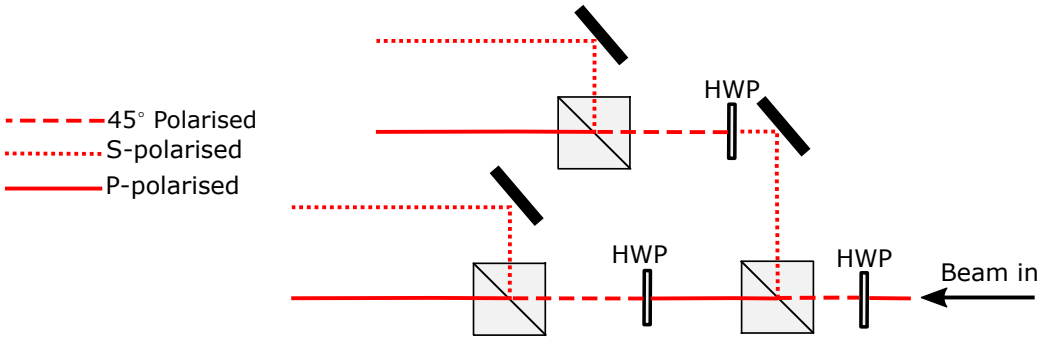


Figure 1.9: A P-polarised beam entering from the right can be rotated to 45° using a Half Wave Plate (HWP) and directed into a polarising beam splitter. A HWP is used to rotate the beams polarisation to 45° to ensure equal splitting at the next polarising beam splitter. The beam is then split into two orthogonal beams with S- and P-polarisation which can be split further.

The many pulse replicas can be directed towards separate fibre amplifiers for amplification. To recombine the pulses, the reverse process is performed. A pair of pulses with orthogonal S- and P-polarisations are sent towards a polarising beam splitter, but the two output ports are now input ports. Providing the orthogonally-polarised pulses have equal intensity and arrive at the same time, the pulses will combine to create a new polarisation state at 45 degrees to both the S- and P-polarisation states. This pulse can be rotated to either S- or P-polarisation state and sent into another polarising beam splitter for further stages of recombination. This is as per Figure 1.9 but done in reverse. The process of splitting a laser pulse at one optic and recombining the two pulse replicas at another is the function of an interferometer. Interferometers are designed to split a laser beam into two paths, known as arms, and then recombine them back, as shown in Figure 1.10.

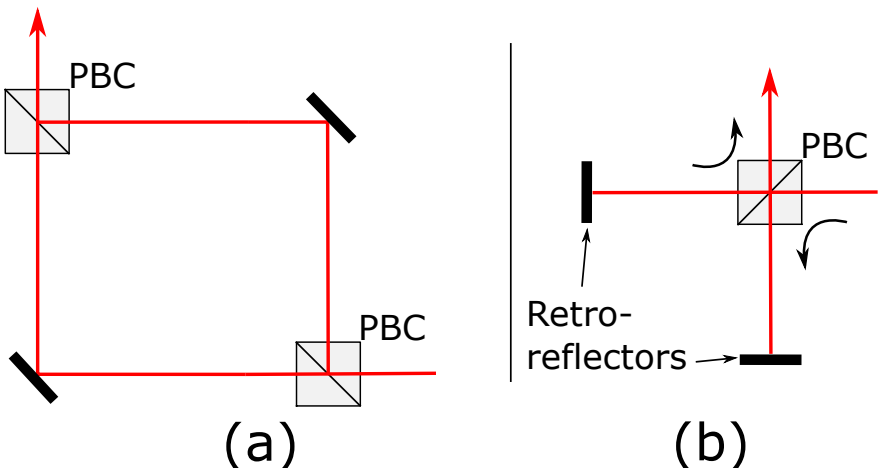


Figure 1.10: Interferometers split a laser beam (typically using polarisation splitting if efficient power recovery is desired) into two paths and then recombine them. The beams in each arm can then be manipulated separately before they are recombined. (a): A Mach-Zehnder interferometer: the beams are split at one polarising beam splitter and recombined at another. (b): A Michelson interferometer: the beams are split and recombined on the same polarising beam splitter. Not shown are the quarter wave plates needed to rotate the polarisation of each beam during retro-reflection. For an interferometer to combine two beams into one output, the beams must arrive at the combination optic in phase.

Interferometers find wide use in applications of distance measurement, with the most famous example being the detection of gravitational waves [40] in 2016. Also, interferometers have found application in combining the outputs of multiple fibres lasers into a single beam. For an interferometer to combine two ultrafast laser pulses correctly, the laser pulses must arrive at the interferometer's combination optic at exactly the same time. For example: for two laser pulses with pulse durations of 1 ps, a length mismatch of $\sim$ 50 μm between the two arms of the interferometer is sufficient to affect the combination of the two pulses. A length mismatch greater than 300 μm is enough to prevent any combination at all. This is to say that the two arms of the interferometer must have the length of their arms matched to micrometre accuracy to ensure combination. At this point, it is convenient to introduce the concept of coherent combination. Coherent combination is the process of ensuring the phase relationship between the combining beams is fixed, so that there is a stable time delay between the arrival of the two pulses at the combination optic. By controlling the phase of the two pulses, assuming they are temporally overlapped, there exists a fixed polarisation state, whether it be circular, elliptical or linear. In order to achieve coherent combination, methods of stabilising the two arms of the interferometer need to be employed and this, along with several methods of coherent combination, will be discussed further in Chapter 3.

Projecting forward the possibility of making a large, facility-scale laser system using coherently combined fibre lasers, design studies have been performed by the International Center for Zetta-Exawatt Science and Technology (IZEST) looking at the potential construction of the so-called International Coherent Amplification Network (ICAN) [41]. This system is designed to coherently combine the outputs of hundreds of fibres and is proposed to operate at a repetition rate of 15 kHz with output pulses of 32 J each, giving an average output power of ∼500 kW. Assuming pulse durations of 200 fs, based upon recent experiments [38], would give the system an approximate peak power of 160 TW, making it suitable for driving laser-wakefield acceleration. Advances in generating shorter pulse durations, potentially below 100 fs, could see the peak power climb even higher. Cost-estimates of this system place the cost/W (average power) at a level approximately three orders of magnitude less than that of a petawatt class Ti:Sapphire laser system [41].

1.5 Research in the field of fibre lasers

Despite the challenges faced by a fibre laser architecture driving a laser wakefield, the most significant being their low peak powers, different groups have pursued research into increasing a fibre laser's peak power. The Friedrich-Schiller-University (FSU) in Jena has demonstrated the coherent combination of four PCF lasers to produce pulses containing 5.7 mJ of energy with 22 GW of peak power [38]. A further experiment using eight fibres was demonstrated also where a combined pulse energy of 12 mJ was achieved, giving a peak power of 35 GW [42]. In addition, the same authors extracted a record 2.2 mJ pulse energy from a single fibre amplifier, giving a peak power of 3.8 GW at 480 fs [33]. Work on non-linear compression has been reported also, using a krypton gas-filled hollow-core PCF to induce spectral broadening on a pulse [43]. This allowed the compression of a pulse from 340 fs down to sub 30 fs. Solid core PCFs can be used for spectral broadening also, and this has been demonstrated by a few groups [44, 45]. However, solid core PCFs are not suitable for use at multi-gigawatt / terawatt peak powers due to material damage.

The scalability of a coherently combined fibre laser system has been demonstrated also, combining 64 CW fibre lasers into a single beam using a tiled aperture setup [46] (discussed in Chapter 3, Section 3.1). This experiment showed excellent potential for the combination

of a large number of amplifiers. The difference between a tiled aperture setup and a filled aperture setup, the latter used for the combination of four PCFs [38], is in the way the beams are combined. For tiled aperture setups, the beams remain temporally locked but spatially separated until they are combined in the far field. Filled aperture setups combine the beams both spatially and temporally in the near field. Tilled aperture and filled aperture setups are covered in more detail in Chapter 3.

Fibres lasers used in broadband oscillators have also seen improvements to their pulse duration, pulse energy and average power. State of the art, compressed pulse durations down to 28 fs have been reported [47] from a ring cavity using Yb^{3+}-doped fibre amplifiers as part of ring cavity. Also, pulse durations of 42 fs have been reported [48] with an all-fibre-integrated (all-fibre ring-cavity) Yb^{3+}-doped fibre laser. Other groups have reported 100 fs pulse durations, but with pulse energies of 534 nJ, giving peak powers of 3.2 MW from a fibre oscillator [49].

Pulse synthesis is another avenue for creating short pulses and involves splitting a broad spectrum of 50 nm into two pieces using a dichroic mirror, and then amplifying them separately [50]. These amplified pulses are coherently combined into a single pulse with a larger bandwidth of 19 nm, compared to an equivalent single pulse with a bandwidth of 12 nm, after gain narrowing. Pulse synthesis has been used with fibre lasers to generate pulse durations of 130 fs. Pulse synthesis has been used also to demonstrate the recovery of 40 nm of bandwidth from the initial bandwidth of 50 nm. The 40 nm of bandwidth was achieved by splitting the 50 nm spectrum of the input pulse into 10 small, overlapping slices in the spectral domain (imagine 10 Gaussians in the spectral domain side by side and overlapping near their edges). Each small bandwidth is a pulse in its own right, with a central frequency and bandwidth. Each pulse is amplified separately before the pulses are recombined. The 40 nm recovered bandwidth gives pulse durations of 100 fs without any phase compensation; pulse durations of 70 fs are predicted given proper phase compensation [51].

All the work highlighted above aimed to improve the peak power of a fibre laser. This has resulted in the improvements to the pulse energy, pulse duration and bandwidth while mitigating the effects of non-linearities and higher order mode amplification in larger core

fibres. Many of the fundamental principles required to create a large system like ICAN have been demonstrated successfully under ideal conditions and on a small scale. One aspect that still requires work is the bandwidth of a pulse post-amplification, to enable pulse durations of sub 100 fs in larger systems. In addition to this, other issues might arise when combining the beams from a large number of fibre lasers. In particular, coherently-combining the outputs of fibres which have originated from different batches, which could potentially have different lengths, core diameters and internal structures. Can manufacturers accurately control the core diameter and fibre length within acceptable tolerances when thousands are being mass produced? Statistically, what could be the maximum variation between the very largest fibre core diameter and the smallest? Would companies be given a PCF design to build or would they meet the criteria for core diameter with their own design? What effect would coherently combining two fibres with different structures have on the final beam?

In order to answer some of these questions, some assumptions have been made about the fibres that could be used for such a project. Firstly, it is assumed that the manufacturers are likely to design their own PCFs using the equipment, intellectual property and licences they have to meet the design requirements on core diameter and dopant concentration. Secondly, the core diameters of all the fibres will be a truncated Gaussian distribution around the target core diameter, to model the uncertainty in manufacture. Moreover, the lengths of the fibres will not be exactly the same, with the fibre lengths also being a truncated Gaussian distribution around a target length. Fibres produced slowly under controlled conditions are going to have less uncertainty in their core diameter and length, compared to fibres produced on a large scale where volume is the main driver. In mass production, checking every fibre to ensure it is the correct core diameter is not practicable. In this case, mass-produced fibre products will have tolerances on core diameter and length. For experiments that have demonstrated the coherent combination of fibre lasers, it is expected that the fibres/PCFs would be chosen to be as closely matched as possible. For some of the experiments reported, the near identical matching of the fibre lengths was essential [46]. However, for a larger scale laser system utilising hundreds/thousands of fibre lasers, like ICAN, it may not be economical to produce near-identical PCFs.

Therefore, using the previously stated assumptions, it was considered that, from a large distribution of PCFs, many of the PCFs would have different core diameters. However, manufacturers today can ensure the core diameter is within ∼10% of the target core diameter. Assuming the target core diameter is 80 µm, there could be PCFs with core diameters ranging from 72 µm up to 88 µm being coherently combined. One of the underlying causes of this core diameter variation is the challenge in controlling the size of the holes used for the photonic structures. Because of this difficulty, PCFs are not made at the same rate as standard step-index fibres, therefore are produced in smaller quantities and at much slower speeds. The obvious solution to the core diameter variation problem is to inspect every PCF and reject any which fall outside a set tolerance band. However, this is time-consuming and could generate a significant amount of waste, which will drive up the cost. Another solution is to see whether two PCFs with different core diameters may be coherently combined and, if so, how much the combination will be degraded due to the core diameter mismatch.

The length of the PCFs is considered to be variable also. During manufacture, the length of the PCF may need to be shortened below the specified figure. This could be to remove an end-section of a PCF that has sustained damage or because it contains manufacturing imperfections. If the tolerance band for the PCF length is too stringent, the wastage may lead to higher manufacturing costs per PCF. As with the core diameter, the slightly shorter PCFs may still be usable in a coherent combination setup.

To investigate coherently combining PCFs with very different properties, two PCFs were used which were found in storage in the Oxford physics department fibre lab. These PCFs were produced by the same company, NKT Photonics, but were purchased four years apart. The two PCFs had vastly different lengths: the shorter one was 520 mm long while the longer one was 820 mm long. The core diameters were very different also: the shorter PCF had a core diameter of 70 ±5 µm, while the longer PCF had a core diameter of 85 ±5 µm. Finally, the structures of the PCFs were very different: the 70 µm PCF had stress rods for polarisation maintenance while the 85 µm PCF had not. These three properties made the PCFs different in many of the ways which could potentially affect the combination in a large experiment where 1000's of PCFs were used.

Because of the differences between these two PCFs, it was seen as a good analogy to the scenario where a large, facility-scale laser system procured PCFs from multiple vendors and different batches. Therefore, I wanted to investigate whether or not PCFs with exaggerated differences could be coherently combined. Using PCFs with exaggerated differences enabled easy identification of problems that may prevent their coherent combination. If problems arose, solutions could be developed to mitigate them.

1.6 Quantifying coherent combination

In order to assess whether the coherent combination of two fibre amplifiers is deemed to be 'good', the efficiency of the system can be analysed using the combination efficiency and the system efficiency [52]. The combination efficiency is a measure of how well a number of beams have combined together, as measured at the output of the system. The effects of losses due to polarisation and absorption in optical components are neglected [52]. It is defined as

$$\eta_{\mathrm{comb}} = \frac{\mathrm{P}_{\mathrm{comb}}}{\sum_{\mathrm{j}} \mathrm{P}_{\mathrm{j}}}, \tag{1.9}$$

where η_{comb} is the combination efficiency, $\mathrm{P}_{\mathrm{comb}}$ is the combined output power when the system is stabilised and P_{j} is the individual channel output power measured at the output. P_{j} assumes that all of the power has reached the system output and has not undergone power losses at polarisation elements, due to an incorrect polarisation state. As mentioned in [52], the combining efficiency does not take into account the losses of the beams through the system; this can lead to overly optimistic efficiencies. Therefore, a better metric of system performance is the system efficiency. The system efficiency is defined in the same way as Equation (1.9), but the powers of the individual channels are measured prior to any combination elements. The system efficiency takes into account the losses experienced by the beams at the combining polarising beam splitter and any losses due to absorption. It determines how much power from the individual channels is present in the final beam, despite the cumulative losses the beams experience through the system. For this thesis, the combination efficiency was omitted, and only the system efficiency was calculated.

There is a further metric, which takes the ratio of the linearly polarised beam power transmitted through a polarising beam splitter and the power incident to that polarising beam splitter. This gives a metric called the fraction of available beam power, and it is defined as

$$\mathrm{FABP} = \frac{\mathrm{P_L}}{\mathrm{P_t}}, \tag{1.10}$$

where $\mathrm{P_t}$ is the total combined beam power incident upon the polarising beam splitter and $\mathrm{P_L}$ is the total linearly polarised beam power transmitted through the polarising beam splitter. This metric defines the amount of useful polarised laser power that is available to an end user. For the coherent combination of fibres to be deemed 'good', the fraction of available beam power should be +90%. However, any beam power fraction over 50% is deemed to be gaining, due to the coherent combination of fibre lasers. This is the important metric used to characterise the quality of combination, and it relates to the useful beam power which will contribute to further stages of coherent combination between N lasers in a larger system.

1.7 Outline of the remainder of thesis

This thesis is broken down into nine further chapters. Chapter 2 gives a more detailed description of fibre lasers and fibre modes. PCFs and their differences with step-index fibres are discussed, followed by introducing PCFs in general and the PCFs used in this thesis. Chapter 3 discusses the coherent combination of two or more fibre lasers, the pulse-phase stabilisation techniques necessary to allow this to happen and the challenges to be overcome for laser systems designed to combine large numbers of fibre lasers. Chapter 4 looks at pulse propagation theory and introduces the mathematics and models required for the coherent combination of mismatched PCFs.

Chapter 5 discusses the experimental setup I constructed, looking at the equipment and diagnostics used to ascertain the peak power, beam profile and efficiency of combination. Chapter 6 presents the work I performed on characterising the seed laser to understand the laser pulse characteristics prior to the coherent combination experiment. Chapter 7 investigates some of the wavefront and thermal effects that needed to be overcome in order to make the coherent combination experiment more reliable and repeatable. Chapter 8

presents the results of the coherent combination of two mismatched PCFs which I recorded in low and high power regimes. Chapter 9 builds upon the results seen in Chapter 8 and presents analysis and further modelling work. Finally, Chapter 10 discusses the conclusions of the experiment and identifies the further work that is needed as a result of this research.

Chapter 2

Fibre lasers

2.1 Solid state ultrafast lasers

Laser systems exist which are able to drive a wakefield to accelerate electrons. At the time of writing this thesis, the state of the art for laser wakefield acceleration was a self-injected electron bunch accelerated to 4.2 GeV using the 1 petawatt laser at the Berkeley Laser Lab Accelerator (BELLA) [9]. This is a solid-state laser making use of Ti:Sapphire crystals as the gain medium to generate pulses with 40 fs pulse durations, pulse energies of 42.2 J and at a repetition rate of 1 Hz. However, the wall plug efficiency of the BELLA laser system is not high, with one estimate placing it at 10^{-4} [53]. Solid state lasers are laser systems which employ either semiconductors, doped-glass or -crystals to act as the gain medium and are typically macroscopic bulk amplifiers, being millimetres to centimetres in transverse scale. The doped crystal chosen will depend on the wavelengths in which it can absorb or emit radiation and also on the bandwidth it can support. Crystals which can amplify larger bandwidths will be able to support shorter pulses which make them suitable for laser wakefield acceleration. Ti:Sapphire is one such crystal that has seen widespread use in facilities across the world. It was used to achieve the current record for laser wakefield acceleration generated electron energies [9]. This is due to its ability to support laser pulses with very short pulse durations, typically 50 fs or less and to produce pulse energies between 100 mJ - 60 J [54, 55, 56]. These short pulse durations and high pulse energies lend themselves to providing high peak powers, with the state of the art being over 1 petawatt [9, 29]. The world record is held by a 4 petawatt laser in South Korea [28]. Future facilities, like the Extreme Light Infrastructure (ELI) aim to reach 10 petawatt peak powers [57].

A schematic layout of a terawatt class Ti:Sapphire laser is shown in Figure 2.1. Because of the peak powers involved, an ultrafast (ps to fs) transform limited pulse travelling around the system would experience non-linear effects at almost every lens encountered. In addition, damage to optics would be highly likely. To mitigate these effects, these systems employ chirped pulse amplification to spread the energy out temporally, and large beam sizes to spread the energy out spatially. The spatial and temporal compression will occur as close to the point of delivery as possible.

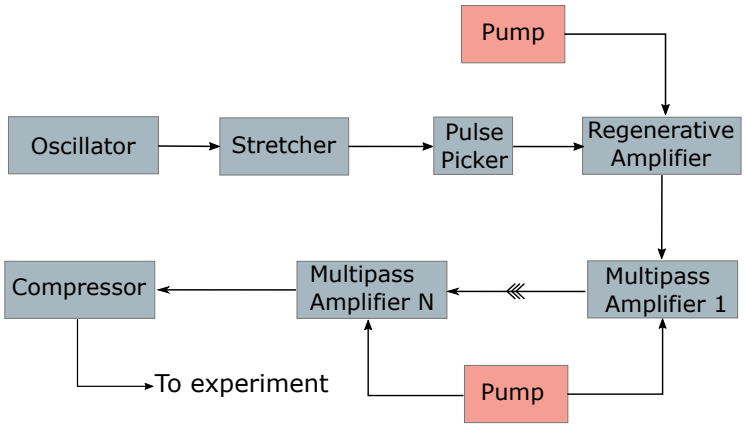


Figure 2.1: A schematic layout of a typical terawatt laser design. Although the oscillator is able to output pulses at high repetition rates, in order to attain high pulse energies, the repetition rate of the system is reduced. This reduction is performed firstly to let the gain in the amplifier build up between pulses, and secondly to prevent excessive heating of the components, which can lead to damage. The repetition rate reduction occurs after the stretcher as the peak power of the pulses emanating from the oscillator can be sufficient to cause damage to the pulse picker used. Typical pulse pickers are acoustic optical modulators or Pockels cells. Terawatt solid-state laser systems tend to output at repetition rates of sub-Hz to Hz to avoid significant heat build up in the amplifiers. Multipass amplifiers generally use Ti:Sapphire crystal with the pulse taking several trips through the crystal to build up sufficient energy. After amplification, the pulse is compressed and directed towards the target area.

2.1.1 Drawbacks of solid state lasers for laser wakefield acceleration

The effectiveness of spreading the energy out spatially is limited by the size of the optical components available. For precision work, where the wavefront of the laser pulse is important, a surface flatness of $\lambda/10$ and below is required. The difficulty in creating optics that are optically flat to within a fraction of a wavelength over a desired area increases non-linearly given the area and the flatness required. Therefore, larger components cost more to buy and replace than smaller components. The presence of a stretcher and compressor will

add additional design challenges in terms of extra system components and their size. A consequence of these larger optics is that the physical space taken up by a 100 terawatt - 4 petawatt system can be large, about the size of a basketball court in some cases, with ELI [57] expected to be the largest yet. Ti:Sapphire systems are not known for their compact designs, usually spanning several rooms of a facility. An additional consideration is the rapid rate at which these systems succumb to non-linear effects, with even a single lens able to add enough non-linear phase to reduce the peak power. This influences the design of such a system and often results in reliance on reflective optics rather than transmissive ones.

Other limitations of these systems are their wall plug efficiencies and their repetition rates. Given the optical pump power delivered to the amplifiers, the amount which typically gets transferred to the seed pulse is approximately 30-50% [55, 58, 29] and is described as the slope efficiency. The remainder of the energy used to excite the lasing ion is lost as heat in the crystal, with the amount being proportional to the quantum defect: the energy difference between the absorption and emission wavelengths. This heat build-up is one of the largest barriers to higher repetition rates. If this heat cannot be removed, then thermal effects will begin to distort the seed pulse and eventually, damage will occur to the amplifier medium. The pump laser used will depend on the chosen seed wavelength and the pump absorption spectrum of the crystal. Longer pump wavelengths can take advantage of the NIR-IR range of diode lasers, giving pump wall plug efficiencies of $\sim 50\%$. At the time of writing, the state of the art for high wall plug efficiency diode lasers belongs to nLight, who demonstrated a 975 nm pump laser with an output of > 100 W at a wall plug efficiency of 73% [37]. Pump lasers operating in the visible range (pumping of Ti:Sapphire near 490 nm, for example) required flash lamps previously to provide the necessary wavelengths, and these systems would typically have wall plug efficiencies of just a few percent. Today, pump lasers can use frequency-doubled, solid-state lasers, which can benefit from the efficiency of diode lasers. Frequency conversion creates another source of loss, but with reported intracavity frequency upconversion efficiencies of $> 36\%$ [59], this could give a final upconverted-pump wall plug efficiency of $\sim 25\%$ using the best-case diode wall plug efficiency of 73%. Taking the $\sim 30\%$ conversion efficiency from pump to seed, Ti:Sapphire lasers might climb to wall plug efficiencies of

$$\eta_{\mathrm{w}} = \eta_{\mathrm{PD}} \times \eta_{\mathrm{FUC}} \times \eta_{\mathrm{PSC}} = (73\%) \times (36\%) \times (30\%) = 7\%, \tag{2.1}$$

where η_{w} is the total Ti:Sapphire wall plug efficiency, η_{PD} is the pump diode wall plug efficiency, η_{FUC} is the efficiency of frequency up conversion and η_{PSC} is the efficiency of pump to seed conversion. Even with the best possible numbers used for the efficiency calculation of a solid state laser, the wall plug efficiency of such a system is still below 10%. If a laser wakefield accelerator is going to be produced and deployed into a hospital, then this efficiency must be improved upon. As mentioned in Chapter 1, fibre lasers are seen as a solution to overcome the drawbacks of solid states lasers when used for laser wakefield acceleration. This chapter aims to introduce the reader to the key concepts of a fibre laser so they can appreciate the challenges involved in their construction and use. Photonic crystal fibres will subsequently be introduced and their advantages over step-index fibres discussed. Finally the photonic crystal fibres used in this thesis will be presented.

2.2 Fibres

2.2.1 Fibre modes

The intensity profile of a fibre's beam can be represented as a linear combination of fibre modes. For two fibres to combine, the intensity profile of the output beams must be as closely matched as possible. Close matching ensures that the combined beam's polarisation state is uniform and not spatially varying, an important point that will be discussed further in Chapter 7. Therefore there is a requirement to understand what modes are produced by a given fibre design. Fibre modes are light waves, whose spatial field distributions do not change as they propagate down the fibre, due to total internal reflection. For a wave to exist and propagate in a fibre, it must satisfy both the wave equation and the applicable boundary conditions, Nonlinear fiber optics, 2013, Agrawal, GP [60]. For step-index fibres, their cylindrical symmetry lends itself to re-writing the wave equation in cylindrical coordinates, as per Equation 2.2

$$\frac{\delta^2 \mathrm{E}(\rho,\mathrm{z},\mathrm{t})}{\delta\rho^2} + \frac{1}{\rho}\frac{\delta \mathrm{E}(\rho,\mathrm{z},\mathrm{t})}{\delta\rho} + \frac{1}{\rho^2}\frac{\delta^2 \mathrm{E}(\rho,\mathrm{z},\mathrm{t})}{\delta\phi^2} + \frac{\delta \mathrm{E}(\rho,\mathrm{z},\mathrm{t})}{\delta \mathrm{z}^2} + \mathrm{n}^2\mathrm{k}_0^2\mathrm{E}(\rho,\mathrm{z},\mathrm{t}) = 0, \tag{2.2}$$

where $k_0 = 2\pi/\lambda$, n is the refractive index and $E(\rho, z, t) = E(\rho)e^{i\beta z - \omega t}$ is the electric field. From the three different field components E_ρ, E_ϕ, E_z, only one of them is independent, and this is also true for the magnetic fields H_ρ, H_ϕ, H_z. The reason for the presence of the z components is due to the finite nature of the waveguide in which they travel. If a finite boundary exists, so too must the z components of the electric and magnetic fields. The different types of modes that can travel down a fibre are transverse electric (TE), transverse magnetic (TM) and hybrid modes (HE/EH, the order depending on which component is larger). TE modes are defined as having $E_z = 0, H_z \neq 0$, TM modes are defined as having $E_z \neq 0, H_z = 0$, and hybrid modes are defined as having $E_z \neq 0, H_z \neq 0$. TE and TM modes propagate via total internal reflection and pass through the centre of the fibre core as they do so. For TE modes, as the wave propagates, the electric field points in the same direction, hence, always remains transverse to the direction of travel. For hybrid modes, the wave vector spirals around the centre of the fibre axis, never crossing through it. To propagate in this fashion, both the electric and magnetic field must have a component facing in the z-direction. Given that all modes will have either E_z or H_z components, it is convenient to represent the wave equation in terms of the z components only. Once both E_z and H_z are known, then all other field components can be determined, provided the waveguide is longitudinally homogeneous [60]. The field components for E_ρ, E_ϕ, E_z and H_ρ, H_ϕ, H_z, in terms of the other field components, can be found from Maxwell's equations in cylindrical coordinates

$$\frac{1}{\rho}\partial_\phi E_z - i\beta E_\phi = i\omega\mu H_\rho, \tag{2.3}$$

$$i\beta E_\rho - \partial_\rho E_z = i\omega\mu H_\phi, \tag{2.4}$$

$$\partial_\rho E_\phi - \frac{1}{\rho}\partial_\phi E_\rho = i\omega\mu H_z, \tag{2.5}$$

$$\frac{1}{\rho}\partial_\phi H_z - i\beta H_\phi = i\omega\epsilon E_\rho, \tag{2.6}$$

$$i\beta H_\rho - \partial_\rho H_z = i\omega\epsilon E_\phi, \tag{2.7}$$

$$\partial_\rho H_\phi - \frac{1}{\rho}\partial_\phi H_\rho = i\omega\epsilon E_z, \tag{2.8}$$

where ρ is the radius, ϕ is the rotation angle and z is the longitudinal distance. ∂ is short hand for a derivative with the subscript indicating the parameter which is being differentiated, ω

is the angular frequency and ϵ and μ respectively are the permeability and permittivity of free space. All of the transverse components ($\mathrm{E}_\rho, \mathrm{E}_\phi$ and $\mathrm{H}_\rho, \mathrm{H}_\phi$) can then be expressed in terms of $\mathrm{E_z}$ and $\mathrm{H_z}$

$$(\mathrm{k}_0^2 - \beta^2)\mathrm{H}_\rho = \mathrm{i}\beta\partial_\rho \mathrm{H_z} - \frac{1}{\rho}\mathrm{i}\omega\epsilon\partial_\phi \mathrm{E_z}, \tag{2.9}$$

$$(\mathrm{k}_0^2 - \beta^2)\mathrm{H}_\phi = \mathrm{i}\omega\epsilon\partial_\rho \mathrm{E_z} - \frac{1}{\rho}\mathrm{i}\beta\partial_\phi \mathrm{H_z}, \tag{2.10}$$

$$(\mathrm{k}_0^2 - \beta^2)\mathrm{E}_\rho = \frac{1}{\rho}\mathrm{i}\omega\mu\partial_\phi \mathrm{H_z} - \mathrm{i}\beta\partial_\rho \mathrm{E_z}, \tag{2.11}$$

$$(\mathrm{k}_0^2 - \beta^2)\mathrm{E}_\phi = \frac{1}{\rho}\mathrm{i}\beta\partial_\phi \mathrm{E_z} - \mathrm{i}\omega\epsilon\partial_\rho \mathrm{H_z}, \tag{2.12}$$

where $\mathrm{k}_0^2 = \omega^2\mu\epsilon$. As mentioned, this shows that once $\mathrm{E_z}$ and $\mathrm{H_z}$ are known, then all other components can be deduced. From Equation (2.2), k_0 is the wave vector of the incident light wave travelling down the fibre. However, this can be decomposed into a longitudinal and transverse component, which is represented by the wave vector triangle as seen in Figure 2.2.

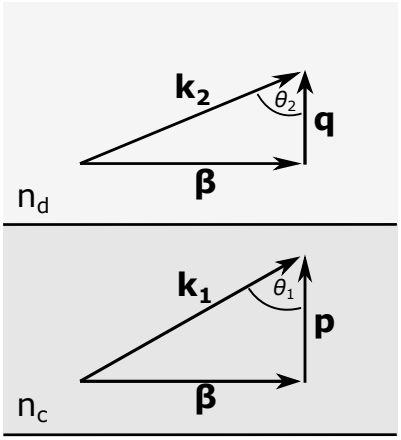


Figure 2.2: The wave vector triangle depicted inside the core region (material with refractive index $\mathrm{n_c}$) and the cladding region (material with refractive index $\mathrm{n_d}$). $\mathrm{k_1}$ is the wavevector of the wave propagating down the core of the fibre, $\mathrm{k_2}$ is the wavevector of the wave propagating in the inner cladding, β is the longitudinal wave vector or the "propagation constant" and p and q are the transverse wave vectors in the core and cladding materials respectively. θ_1 and θ_2 are the angles subtended between the normal of the core cladding interface and $\mathrm{k_1}$ and $\mathrm{k_2}$ respectively.

From Figure 2.2, β is the longitudinal wave vector or the "propagation constant", p is the transverse wave vector in the region of $\mathrm{n_c}$ and q is the transverse wave vector in the region of $\mathrm{n_d}$. The condition for a mode to travel down the fibre is that the transverse component, p, should reflect back and forth between the walls of the fibre and undergo constructive inter-

ference with itself, known as the transverse resonance condition. This resonance condition is defined as

$$2\mathrm{pd}\cos(\theta) + 2\Phi(\theta) = 2\mathrm{m}\pi, \tag{2.13}$$

where d is the diameter of the fibre, m is an integer, Φ is the phase shift produced at the core/cladding interface and θ is the angle subtended between k_0 and the normal vector from the core wall. The field must make a complete round trip in order to interfere with itself, hence the factor of two in front of the phase term. Because m is an integer, θ can only assume certain discrete values in order to satisfy the transverse resonance condition. This subsequently defines discrete values for β using the wave vector triangle as seen in Figure 2.2.

For the cases where $\mathrm{k}_0 > \beta$, the transverse component is real and describes the propagation of a wave down the fibre. For cases where $\mathrm{k}_0 < \beta$, the transverse component is imaginary and, therefore, the wave sees an exponential decay of its field in the transverse direction, known as an evanescent wave. For a guided mode, the wave must be real inside the core, whilst being an evanescent wave outside of the core. These two conditions can be defined as

$$\mathrm{p} = (\mathrm{n_c^2 k_0^2} - \beta^2)^{1/2} \quad (\rho \leq \mathrm{a}), \tag{2.14}$$

$$\mathrm{q} = (\beta^2 - \mathrm{n_d^2 k_0^2})^{1/2} \quad (\rho > \mathrm{a}), \tag{2.15}$$

where $\mathrm{n_c}$ is the refractive index of the core, $\mathrm{n_d}$ is the refractive index of the cladding, p is the transverse wave vector inside the core region, q is the transverse wave vector outside the core region, a is the core radius and ρ is the radial distance from the axis of the fibre. From Equations (2.14) and (2.15), for a mode to exist, beta must therefore satisfy the condition

$$\mathrm{n_d k_0} < \beta < \mathrm{n_c k_0}. \tag{2.16}$$

For the case where $\beta < \mathrm{n_c k_0}$, Equation (2.2) is solved using separation of variables, leading to the general solution

$$\mathrm{E_z}(\rho, \omega) = \mathrm{A}(\omega)\mathrm{F}(\rho)\mathrm{e}^{\mathrm{im}\phi}\mathrm{e}^{\mathrm{i}\beta\mathrm{z}}, \tag{2.17}$$

where $A(\omega)$ is the spectral amplitude, m is an integer, $F(\rho)e^{im\phi}$ is the transverse spatial distribution and β is the propagation constant. $F(\rho)$ is a solution of

$$\frac{d^2F}{d\rho^2} + \frac{1}{\rho}\frac{dF}{d\rho} + \left(n^2k_0^2 - \beta^2 - \frac{m^2}{\rho^2}\right)F = 0. \tag{2.18}$$

Solutions to Equation (2.18) are Bessel functions in the case where $\beta < n_c k_0$

$$F(\rho) = C_1 J_m(p\rho) + C_2 N_m(p\rho), \tag{2.19}$$

where J_m and N_m are Bessel functions of the first and second kind respectively. p is defined as $p = (n_c^2k_0^2 - \beta^2)^{1/2}$ where n_c is the refractive index for $\rho \leq a$. Solutions for N_m have a singularity at $\rho = 0$ making them unphysical, therefore, $C_2 = 0$. This leaves the solution

$$F(\rho) = J_m(p\rho) \qquad \rho \leq a, \tag{2.20}$$

where the constant C_1 has been absorbed into the spectral amplitude $A(\omega)$ and 'a' is the core radius. These mark solutions inside the fibre core; outside the fibre core, the solutions should exhibit exponential decay. This is represented by the modified Bessel function

$$F(\rho) = C_3 I_m(q\rho) + C_4 K_m(q\rho) \qquad \rho \geq a, \tag{2.21}$$

where $q = (\beta^2 - n_d^2k_0^2)^{1/2}$ with n_d being the refractive index for $\rho > a$ and; I_m and K_m are the modified Bessel functions of the first and second kind respectively. The solutions to I_m tend to infinity at large distances which are unphysical in the cladding region, therefore, $C_3 = 0$. Considering boundary conditions, the field components tangential to the boundary $\rho = a$ and the field gradients must be continuous across the boundary. This is to say that

$$J_m(pa) = K_m(qa) \qquad \text{at} \qquad \rho = a, \tag{2.22}$$

$$J_m'(pa) = K_m'(qa) \qquad \text{at} \qquad \rho = a \tag{2.23}$$

This leads to an eigenvalue equation which allows the propagation constant β to be calculated for various fibre modes. The derivation is discussed in other textbooks [61], the final result is quoted here

$$\left[\frac{\mathrm{J'_m(pa)}}{\mathrm{pJ_m(pa)}}+\frac{\mathrm{K'_m(qa)}}{\mathrm{qK_m(qa)}}\right]\left[\frac{\mathrm{J'_m(pa)}}{\mathrm{pJ_m(pa)}}+\frac{\mathrm{n_d^2}}{\mathrm{n_c^2}}\frac{\mathrm{K'_m(qa)}}{\mathrm{qK_m(qa)}}\right]=\left(\frac{\mathrm{m}\beta\mathrm{k_0(n_c^2-n_d^2)}}{\mathrm{an_c p^2 q^2}}\right). \tag{2.24}$$

Equation (2.24) provides solutions for β for a given value of 'm'. The different solutions can be labelled with the subscript 's' accordingly allowing individual fibre modes to be described by β_{ms}. A root solving algorithm is used to calculate which modes can propagate down the core as Equation (2.24) is a transcendental equation. As transcendental equations do not have analytical solutions, numerical methods must therefore be employed. Some of the modes that are supported by a cylindrical waveguide are illustrated as examples in Figure 2.3. Fibres which are said to be multi-mode are able to support a number of different modes simultaneously within their cores. This becomes a problem when trying to predict the beam's intensity profile while it propagates down the fibre. Furthermore, the non-linear effects experienced by a laser pulse vary dramatically between different modes and modal amplitudes. Hence, much effort has been invested into creating fibres which guide the fundamental mode ($\mathrm{LP_{01}}$) only. All other modes are classified as higher order modes and are deemed to be undesirable. For the condition where only the fundamental mode is desired, the cutoff frequency for the first higher order mode can be determined from Equation (2.24) and is found to give a V parameter of $\mathrm{V_c} = 2.4048$, where $\mathrm{V_c}$ is the V parameter at the cutoff frequency.

The normalised frequency / V-parameter cutoff value can be used to determine the maximum core diameter for single mode operation given specific core and cladding materials and operating wavelength. The normalised frequency can be derived by enforcing the condition that at the boundary between the core and the cladding there should be phase matching between the real and evanescent waves. This is to say that β is equal across the boundary. Using Equations (2.14) and (2.15), β can be eliminated by squaring and adding the two equations together

$$\mathrm{(p^2+q^2)=k_0^2(n_c^2-n_d^2)}=\frac{(2\pi)^2}{\lambda^2}\mathrm{(n_c^2-n_d^2)}. \tag{2.25}$$

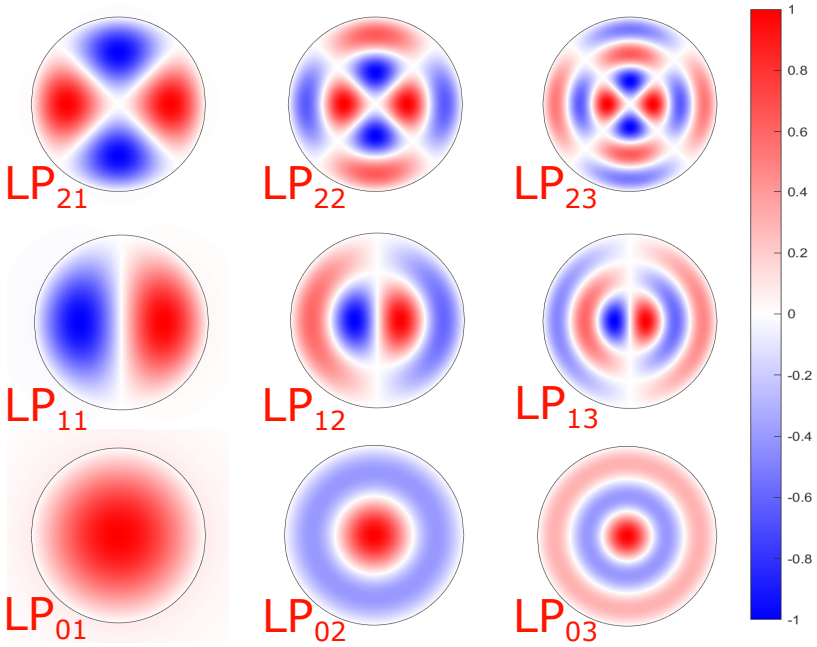


Figure 2.3: The electric fields of fundamental mode (LP_{01}) and a few higher order modes supported by a cylindrical waveguide whose boundary is marked by a black line. The relative, normalised field strength is indicated by the colour bar shown on the right. LP stands for linearly polarised with the subscripts identifying the number of zeros in the azimuthal and radial direction respectively. This is to say that LP_{12} has a single zero crossing in the azimuthal direction (as taken across the entire beam diameter) and two zero crossings in the radial direction. The diagram does not show the degeneracy which can increase the number of modes present. The degeneracy is dependent upon the polarisation state of the propagating light wave. For example, the fundamental mode can have double degeneracy with each mode having an orthogonal polarisation state. Higher order modes with larger numbers of lobes will have greater amounts of degeneracy.

For a guided mode, cutoff occurs when $\beta = n_d k_0$ (from Figure 2.2, when $\theta_1 = \theta_c$, the critical angle) therefore, q=0, as per Equation (2.15). By applying a square root to both sides of the equation and multiplying by a, the core radius, we arrive at

$$pa = \frac{2\pi a}{\lambda}(n_c^2 - n_d^2)^{1/2} = \frac{2\pi a}{\lambda}NA, \tag{2.26}$$

which is the normalised frequency / V parameter as was discussed before in Chapter 1. From Equation (2.26), the core radius and numerical aperture required to give exclusively single mode operation with the fundamental mode (LP_{01}) depends on the cutoff value for the first higher order mode, LP_{11}. Once single mode operation has been established, the peak intensity in the fibre needs to be reduced, which can be done by using large core diameters. In the interest of maximising the core diameter, it is desirable to optimise the fibre parameters to put the V-parameter as close to the LP_{11} cutoff frequency as possible. Finely tuning the

core and cladding refractive index will allow the creation of a small difference between the two. The smaller the refractive index difference, the larger the core diameters can become. Manufacturing materials with precise refractive indices was, at one point, the largest barrier to step index fibres getting larger core diameters. At the time of writing, the largest core diameter commercially available for a single mode step index fibre is ~ 25 μm. However, today, losses due to bending sensitivity limit the practical use of larger core diameters. Photonic crystal fibres (PCFs) are able to achieve larger core diameters than step index fibres whilst remaining in single mode operation. However, the eigenmode analysis performed with the step-index fibres breaks down for PCFs as the core/cladding boundaries are more complex. In order to describe mode propagation in a PCF, modifications are required to Equation (2.26) and also other methods, which will be discussed in Section 2.3.

2.3 Photonic Crystal Fibres

Photonic crystal fibres (PCFs) are the next iteration of fibre technology. PCFs aim to increase the core diameter of a fibre, whilst still maintaining single mode operation, to produce lower intensity beams for a given peak power. This directly impacts on the amount of non-linearity acquired by a pulse as it propagates down the fibre. PCFs make use of an array of holes around a central core region which creates a photonic structure. This photonic structure creates a guiding effect despite the refractive index of the core not being greater than that of the cladding. In this case, it is essential for the refractive index to be matched between the core and cladding so that the guiding is performed by the holes and not the index step between the core and cladding. A schematic of a PCF is shown in Figure 2.4.

The current standard for endlessly single-mode PCFs is a core diameter of 25 μm, similar to that of step-index fibres. Endlessly single-mode is the term used to describe fibres that are single mode for all wavelengths shorter than the structure's theoretical maximum single mode wavelength. The small core diameters of these PCFs allow them to be bent around into coils to allow large amounts of fibre to be used in a small space. For larger guiding structures, the larger core and fibre diameters become fragile due to mechanical stress. Additionally, PCFs are not immune from bending losses, so as larger core diameters are created, it is necessary to keep the fibres straight. Rod fibres are PCFs with fibre diameters of a few millimetres. To

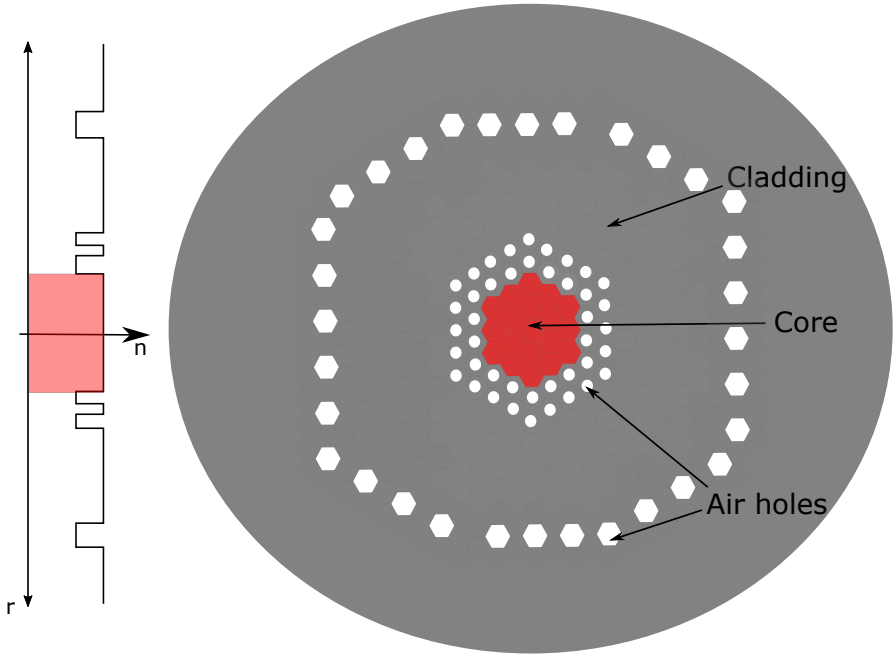


Figure 2.4: A schematic diagram of the internal structure of a PCF. The structure is made by building up a large, preform block using rods fitted together in tessellating fashion. This preform is then heated and drawn out into a fibre. The core region is highlighted in red and around the core are the rods with air holes. During fabrication, the air hole size can be altered by changing the drawing tension on the fibre at a given temperature [62]. This allows the effective refractive index of the cladding region containing air holes to be fine-tuned. The graph on the left gives an indication on the refractive index profile across the diameter of the fibre. It should be noted that the refractive index across the fibre should be the same in both the core and cladding, despite any doping that may be present in the core region. The guiding is performed by the holes, with the seed confined to the core region and pump laser confined to the cladding region.

keep rod fibres straight, the fibres are supported along their entire length by metal structures with V-slot grooves cut into them. Using rod fibres supported in V-grooves, state of the art core diameters of up to 135 μm have been reported [32]. The largest core diameter PCF used in this thesis was $\sim$ 85 μm, with a mode field diameter of 65 μm. Although not state of the art, this is the largest core diameter PCF commercially available.

However, rod fibres are not exclusively single-mode. These fibres operate in a regime known as effective single mode operation [63]. Effective single mode operation is where the design of the fibre would normally allow higher order modes to propagate, making the fibre multi-mode. Strategies are then implemented to filter out these higher order modes to allow the fibre to operate like a single mode fibre. These strategies include bending and differential attenuation. Bending is an approach commonly used with smaller step-index fibres and bendy PCFs and works because higher order modes see high losses when a bend is encountered [64]. Differential attenuation looks at the design of the fibre such that higher order modes see

higher losses as they propagate. The simplest solution is the careful positioning of the gain in a doped fibre. By placing the gain nearer the centre of the core, higher order modes which have lobes of intensity beyond this central region, will see less gain. The fundamental mode is therefore preferentially amplified by the fibre. The alternative choice to effective single mode operation is to change the structure of the PCF to prevent higher order modes propagating at all. The structure of the PCF can be changed by adjusting the holes. The free parameters are the hole size and hole pitch, as shown in Figure 2.5.

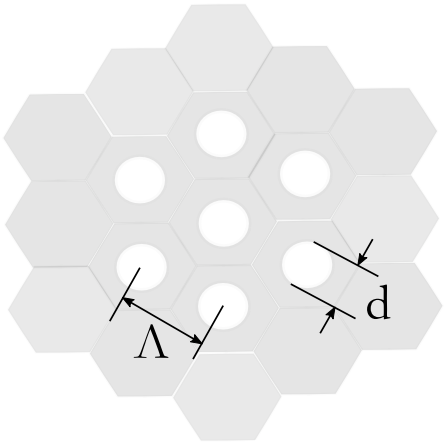


Figure 2.5: A schematic diagram of how the hole size and pitch are defined in a PCF; d is the hole size and Λ is the hole pitch.

The simple argument used in older literature, is that presence of these air holes averages out the refractive index of the cladding material and allows for very accurate fine tuning of the refractive index difference between core and cladding. Assuming a simple total internal reflection model, this allows for large core diameters, which can remain in single mode operation. Accurate tuning is possible as the air hole size, pitch and quantity can be directly controlled and fine adjustment can be made while drawing the fibre [62] (see Figure 2.5). A more rigorous argument, says that the relationship between the holes and their spacing determines whether the fibre will be single mode or support higher order modes [65]. The definition of hole size and pitch is shown in Figure 2.5. Plotting different values of the normalised fill factor, d/Λ, and for different normalised wavelengths, λ/Λ, Figure 2.6 shows the point at which a fibre will transition from single mode operation to one which supports higher order modes [65].

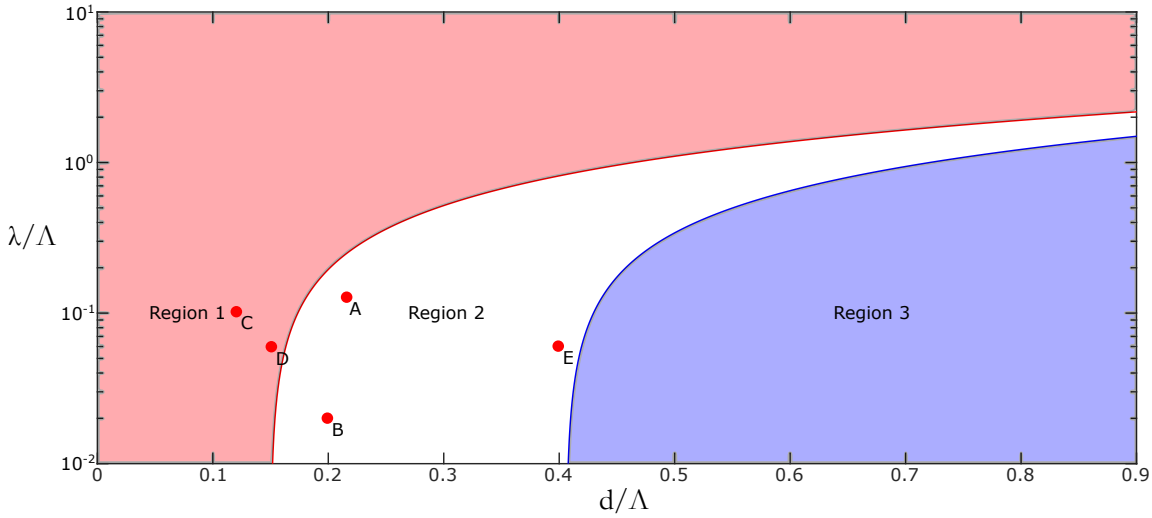


Figure 2.6: A plot of the normalised fill factor (d/Λ) against normalised wavelength (λ/Λ). The red curve and the blue curve separate the diagram into 3 regions. Region 1 is where the single mode penetrates into the cladding and is weakly guided. Region 2 is where the core is single mode. Region 3 is where the core will support higher order modes. [65]. Plotted are a variety of fibre parameters from the literature for reference. A[33], B[32], C[66], D[67], E[68]. Points C and D are earlier designs of PCFs, A,B and E are more recent.

Figure 2.6 shows that if d/Λ and λ/Λ are chosen appropriately, a PCF can be made to be single mode. For d/Λ values which allow λ/Λ to go to very small values and not cross into Region 3, the PCF is endlessly single mode. Using the d/Λ parameter, if a PCF can be constructed to stay under d/Λ ≈ 0.4, then core diameters can tend to infinity [69]. This has led research into Large Pitch Fibres (LPF) which aims to minimise λ/Λ, whilst keeping d/Λ under 0.4 [32]. State of the art LPFs have core diameters as large as 135 μm, whilst maintaining effective single mode operation [32]. Further research, done by other groups, aims to take this a step further and create band-gap structures, which can further enhance the higher order mode filtering. The arrangement of air holes is known as a distributed mode filter [70]. As the core diameters of PCFs get larger and their structures more complex, the fibre designs themselves become fixed to that of a rod, as the usual bending strategies used to suppress higher order modes for smaller core fibres, now affect the single mode. The PCF's susceptibility to bending losses mean they need to be kept straight and be supported along their entire lengths.

2.3.1 End caps

Because of the high pulse energies that can be generated in PCFs, fibres can have an endcap sealed to the end of the fibre as shown in Figure 2.7. The endcap is a macroscopic piece of

fused silica making it suitable only for PCFs rigid enough to support its weight. Because the rod-fibres are prevented from bending to avoid higher order modes, this rigidity also enables them to support the weight of an endcap.

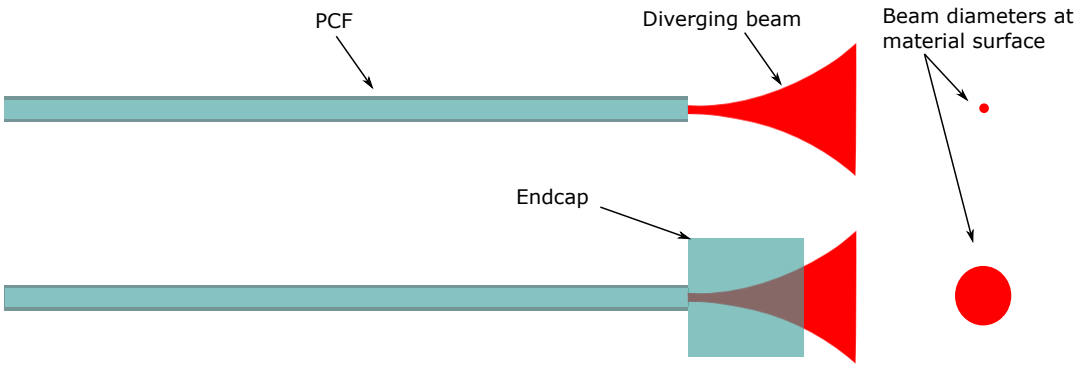


Figure 2.7: A comparison of a fibre with and without an endcap. The fibre without the endcap will see a higher electric field at the air-fibre interface than that with the endcap. If there are any surface imperfections, this could lead to damage occurring due to local electric field enhancement. For high pulse energies, the endcaps become essential.

This endcap's function is to allow the beam to propagate and diverge from its small beam diameter before leaving the fibre material. Because the air-fibre interface is more susceptible to damage than the bulk material, the endcap allows the beam to diverge within the bulk to reduce the intensity before it finally encounters the air-fibre interface.

2.3.2 Advantages and disadvantages of photonic crystal fibres

Given the high average powers possible, cooling of the Photonic Crystal Fibres (PCFs) will need to be considered, as there is a smaller volume to surface area ratio with larger core PCFs. A heat build-up could lead to damage or disturbance of the refractive index difference between the core and cladding, leading to a breakdown of the single mode operation. These 'mode instabilities' are an active field of research and work reported so far [71, 72, 73, 74] has identified a threshold-like behaviour for this breakdown linked to the average power level in the core.

PCFs are able to have large core diameters, whilst maintaining effective single-mode operation. Larger core diameters, which support higher order modes, can be used if strategies are implemented to preferentially-select and amplify the single mode, whilst making the higher order modes see greater losses. Despite the larger core diameter, the PCFs are thin enough to provide excellent heat dissipation properties compared to bulk amplifiers. Choosing

the pump and seed wavelengths to have a small quantum defect will allow these systems to be cooled passively at average powers, which are orders of magnitude greater than Ti:Sapphire bulk amplifiers.

2.4 Photonic Crystal Fibres used in this thesis

The amplifiers that were used for the experiments in this thesis were Yb-doped rod-type PCFs made from fused silica. Yb^{3+} ions doped in silica have a gain spectrum as seen in Figure 2.8 [75].

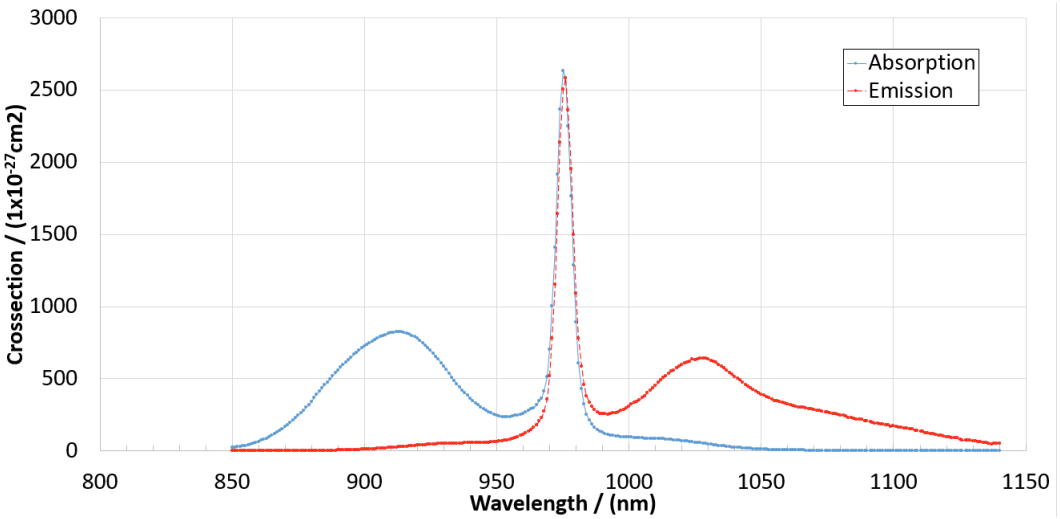


Figure 2.8: The absorption and emission gain spectra for Yb^{3+} ions in fused silica with a resolution of one point per nanometre. Used with permission from Amplitude Systême.

The absorption cross-section seen in Figure 2.8 drives the choice of the pump laser. There are two obvious choices from Figure 2.8: the first is 900-910 nm to take advantage of the broad bump around 910 nm, the second is 976 nm, to take advantage of the high absorption cross-section. The pump lasers used in this thesis fall into the latter category. Given this choice of pump wavelength, a seed laser was chosen to operate at a wavelength $\sim$ 1030 nm. This choice takes advantage of the large hump seen around 1030 nm in the emission spectra from Figure 2.8, but also allowed the use of dichroic mirrors to separate the 976 nm and 1030 nm wavelengths prior to and after amplification. This is necessary as the final turning mirror used to reflect and couple the seed beam to the core of the fibre, must also be transparent to 976 nm, in order to allow the pump laser to couple to the fibre.

For the experiments in this thesis, two PCFs were used. Each fibre was positioned in one of the arms of an interferometer, as discussed in Chapter 1, Section 1.4 to amplify one of the split pulse copies. The PCFs had different core diameters: one had a core diameter of 70 μm; while the other had a core diameter of 85 μm. For the remainder of this thesis, the PCFs will be identified by their core diameter, e.g. the 70 μm PCF. In addition to their core diameters, the PCFs had other parameters and structures which differed between them. This section aims to highlight these differences.

2.4.1 The 70 μm Photonic Crystal Fibre

The 70 μm PCF had a length of 520 mm, a core diameter of 70 μm (mode field diameter of 55 ± 5 μm), mode field area of 2200 ± 200 μm^2 and an inner cladding diameter of 200 ± 5 μm. The numerical aperture for the cladding was significantly larger than that of the core (~ 0.5 compared to ~ 0.02) to enable large powers from multimode pump diodes to be guided. The core was designed to operate in effective single mode operation, which required the numerical aperture to be very small. The detailed information can be found on the Datasheets at Appendix A. A diagram of the internal structure of the PCF is shown in Figure 2.9.

This PCF is from NKT Photonics [76] and is one of their older designs. As seen in Figure 2.9, the PCF contained stress rods which were used to induce stress on the core to make it polarisation maintaining [77] for certain linear polarisations. If the input seed had a linear polarisation state at a different angle to the transmission axis of the core, then it would be guided less strongly and would leak into the inner cladding. A half wave plate was positioned before the fibre to rotate the linearly polarised light to the orientation accepted by the PCF's core. To determine the appropriate lens to couple to the PCF's core, the simple Gaussian beam-focusing equation was used [78]

$$f^2 = \left(\frac{\pi\omega_1\omega_2}{\lambda}\right)^2 + \mathrm{XY}, \tag{2.27}$$

where ω_1 is the incident beam radius, ω_2 is the focused beam radius, λ is the wavelength and X and Y are the distances of ω_1 and ω_2 from the focal points, f, on each side of the lens respectively. To find the waist size ω_2 expected at the focus, we set Y=0 as seen in Figure 2.10.

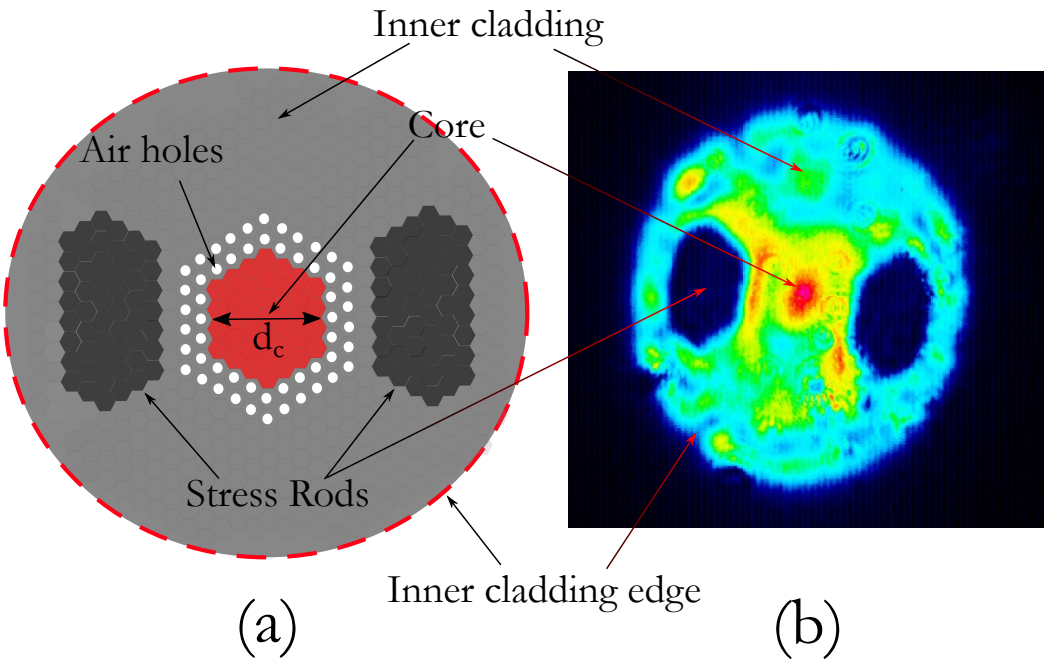


Figure 2.9: (a): A schematic slice of the 70 µm PCF (not to scale), excluding the outer cladding. The red area is the core region and is doped with Yb^{3+}. The core has a mode field diameter of 55 µm, whilst the core diameter is defined by d_c. In the simpler total internal reflection (TIR) model, the air holes surrounding the core serve to average out the refractive index and create a small index step, as previously seen in figure 2.4. The darker grey areas are stress rods, which induce birefringence into the core to make the fibre polarisation maintaining. The red dashed ring indicates the inner cladding diameter. (b): The image of the end of the 70 µm PCF was taken with a beam profiling camera to show the stress rods. The beam was misaligned from the core to allow the fibre's structure to be observed.

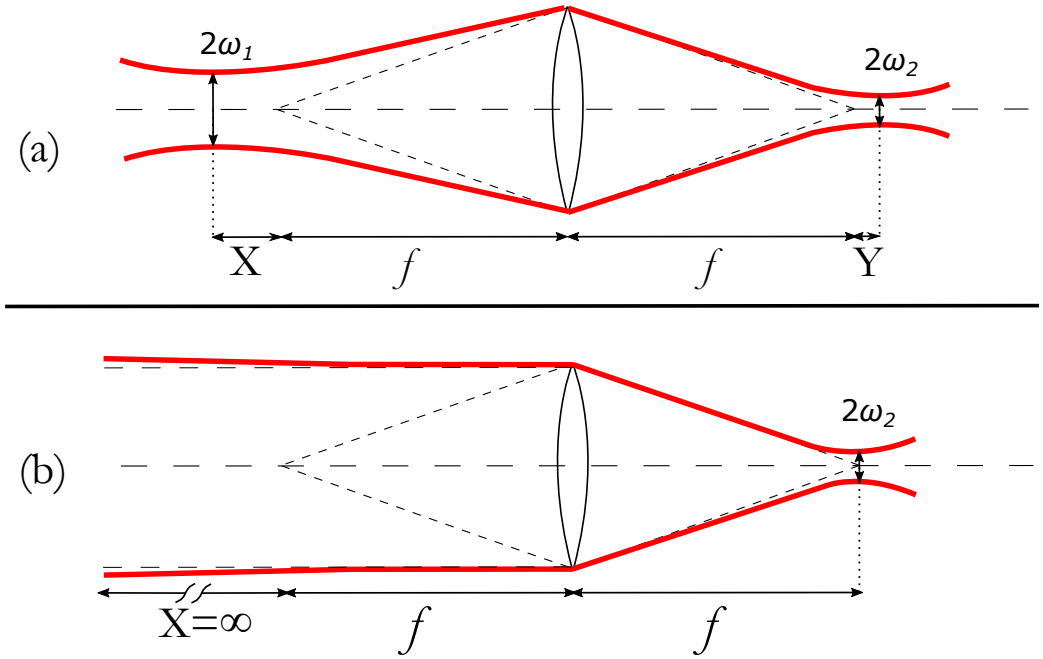


Figure 2.10: (a) Equation (2.27) shown schematically. Dashed lines converge at the lens's focii, dotted lines indicate the beam waist. (b) The result of setting Y=0.

The calculated focal length, was 62.8 mm, which assumed a collimated input beam with a beam diameter of $2\omega_2 \approx 1.5$ mm. From the lenses that were readily available, a focal length lens of $f = 60$ mm was found to be the closest match. However, using this lens to couple the beam to the core of the fibre resulted in some leakage into the inner cladding. This suggested that the diameter of the beam was slightly too large. Given that the above calculation was a simple approximation it was understandable that there would be a discrepancy. To correct this discrepancy, a smaller focal length lens was chosen, with the next available lens being an $f = 50$ mm lens. A beam profiling camera was then used to observe the focused beam spot to confirm its beam diameter. Normally aspheric lenses would be used to couple the seed to a fibre to prevent spherical aberrations. However, no $f = 50$ mm, aspheric lenses were available for purchase from Thorlabs. Instead, Thorlabs sold a series of 'best-form' lenses which were designed to reduce aberrations from spherical lenses. 'Best-form' lenses offered a good compromise between spherical lenses and aspheric lenses, and allowed focal lengths greater than those available using aspheric lenses. The lens which was used to image the opposite end of the PCF and collimate the amplified seed beam as it exited the core, (Figure 2.11(a)), was also the lens used to couple the pump beam to the inner cladding. This lens had a dual functionality, so optimising one beam at the expense of the other would either result in reduced amplification or a poorly collimated amplified seed beam. Both the pump and the amplified seed beams had to be considered when choosing a lens and positioning it relative to the end of the PCF. Figure 2.11(b) shows the effect on the seed beam when the pump coupling lens is moved to improve the pump coupling.

2.4.2 The 85 μm Photonic Crystal Fibre

The 85 μm PCF had a length of 820 mm, a core diameter of 85 μm (mode field diameter of $65 \pm 5\,\mu$m, corresponding to an approximate mode field area of $\sim 3300\,\mu\text{m}^2$) and an inner cladding diameter of $200 \pm 5\,\mu$m. The numerical aperture of the cladding and core was 0.5 and 0.02 respectively, as for the 70 μm PCF. The detailed information is found on the Datasheets at Appendix A. This PCF was a newer design and the Company had moved away from using stress rods to create a polarisation maintaining effect. It was found that

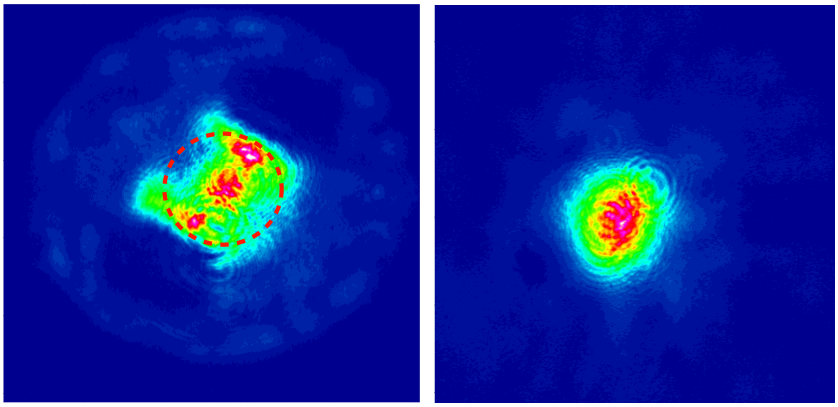


Figure 2.11: (a): The coupling of the seed beam to the 70 µm PCF using the 50 mm best-form lens, with the pump coupling lens used to image the end of the fibre. Seen here is the seed beam exiting the PCF from the pump side of the PCF. There is still some leakage of light into the inner cladding, but the contrast compared to the core (indicated with a red dashed ring) is very good and is similar to what is seen by groups using a similar fibre [77]. Figure 2.9(b) showed the beam slightly misaligned to the core. Here, the majority of the beam power had leaked into the inner cladding. (b): The beam profile in (a) improved when the pump coupling lens was moved closer to better optimise the pump beam coupling. This was a result of focusing the amplified seed beam emitted from further inside the core, which improved the amplified-seed beam-profile.

the rod-type PCFs were able to maintain a linear polarisation state, provided they were kept straight. However, small bends in the fibre could lead to a beam's linear polarisation state deteriorating [79]. This could be avoided by placing the PCFs in V-grooves, which supported the PCF along its entire length. End caps were fitted to prevent breakdown at the ends of the fibre due to surface imperfections. This meant that the fibre could support larger powers [63]. The schematic structure of the fibre is shown in Figure 2.12.

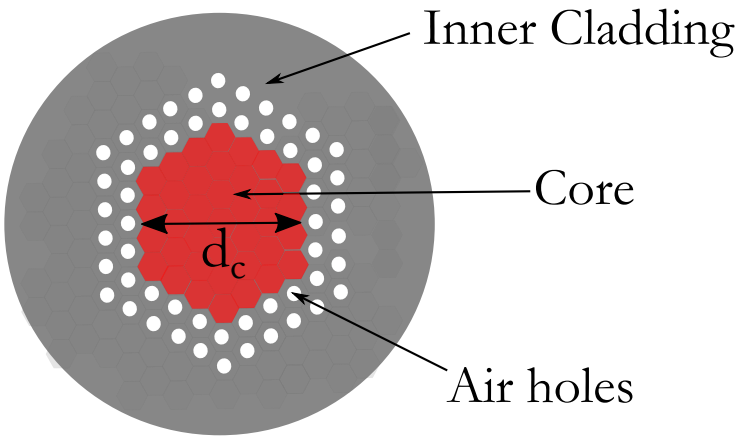
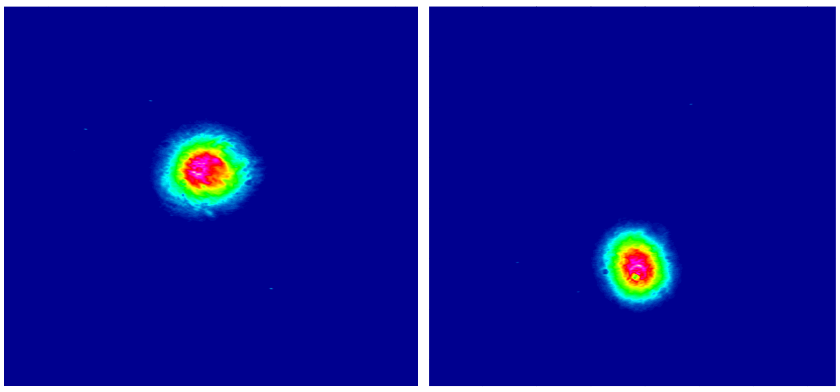


Figure 2.12: A schematic diagram of the structure of the 85 µm PCF. Shown is the core (red) surrounded by the air holes and the inner cladding (grey). The mode field diameter is 65 µm.



(a) (b)

Figure 2.13: (a): The image of the end of the PCF produced with the pump coupling lens, when the seed was coupled to the core using the 60 mm best-form lens at the other end of the PCF. (b) The image of the end of the PCF with the pump coupling lens moved closer to the PCF, to better optimise the pump coupling. The small difference between the images of the beam profile seen in (a) and (b), unlike that seen in Figure 2.11, is due to the newer design of the PCF. The contrast between the core and inner cladding is also superior to that of the 70 μm PCF (Figure 2.11).

The seed coupling lens chosen had a beam of diameter $2\omega_2 \approx 1.5$ mm incident upon it and, using Equation (2.27), gave a calculated focal length of 74.2 mm. This suggested the use of a $f = 75$ mm lens. This lens caused a small amount of leakage into the inner cladding which indicated that the beam diameter was too large. Instead, a $f = 60$ mm best form lens was chosen as there were no aspheric lenses large enough to provide this focal length. Using this best-form lens to couple the seed to the fibre, the image of the end of the fibre produced using the pump coupling lens showed good coupling to the core, as shown in Figure 2.13. The seed beam profile emitted from this PCF, and seen through the pump coupling lens, can be approximated very well to a Gaussian distribution as shown in Figure 2.14.

2.4.3 Pump coupling

For this experiment, the pumps were coupled to the PCFs in a counter-propagating geometry. This means that a pump was coupled to the PCF at one end, and the seed was coupled at the other. Each PCF had its own pump diode laser, with the 70 μm and 85 μm PCF having a 100 W and 500 W pump respectively. Both pump lasers operated at 976 nm and both had pump beam delivery fibres with core diameters of 400 μm and numerical apertures of 0.22,

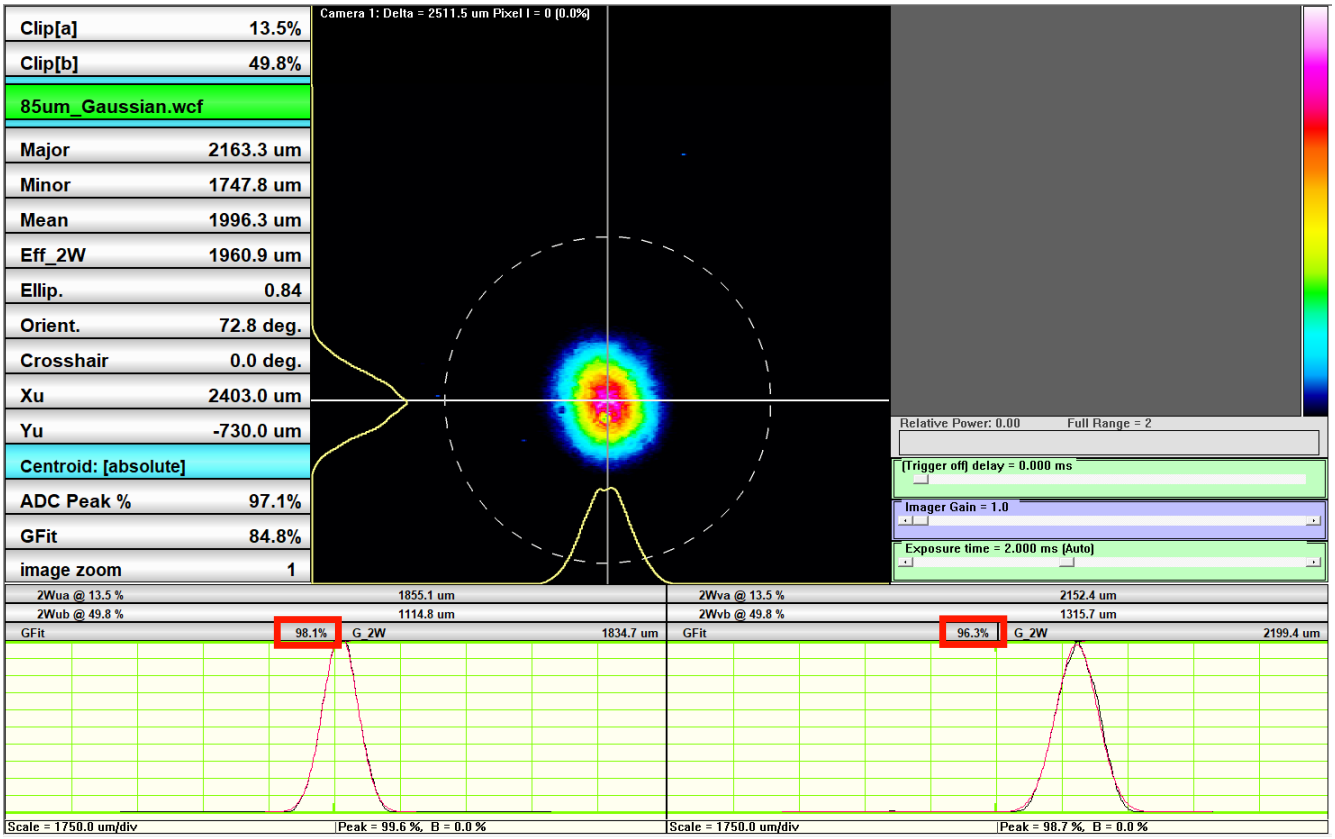


Figure 2.14: The beam from the 85 µm PCF as viewed with a beam profiling camera through the pump beam coupling lens. The projection of the seed beam onto the vertical and horizontal axis are shown in the graphs below with a Gaussian fitted. The GFit parameter gives an indication of the goodness of fit of the fitted Gaussian with the beam projection. Here, both axes gave a GFit of over 96%.

making their output beams highly multimode. The 500 W pump was an older, bespoke diode-laser produced when diode technology for 976 nm was still developing. It had a broader pump spectrum ($\sim$ 3 nm FWHM) compared to the newer 100 W pump laser ($\sim$ 0.3 nm FWHM), which was better designed to target the 976 nm absorption peak (Figure 2.15). The wall plug efficiencies of the pump diodes are $\sim$ 12.5% and 35% for the 100 W and 500 W pumps respectively.

To couple the pump beams into each fibre and to collimate the seed beam leaving each fibre, the position of the pump coupling lens needed to be adjusted to optimise for both pump beam coupling to the inner cladding and seed beam decoupling to produce a Gaussian-like beam. A poorly collimated seed beam would make it difficult to work with whilst a poorly coupled pump beam would mean little power would be transferred to the seed beam. In addition, the 85 µm PCF, unlike the 70 µm PCF, had endcaps that required the pump

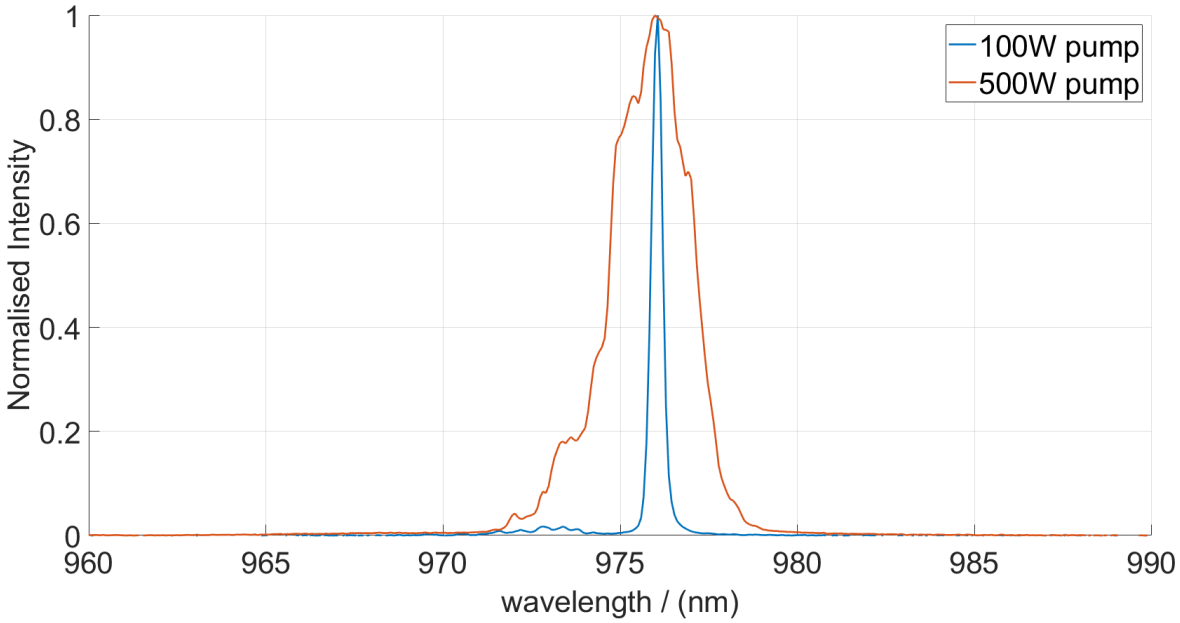


Figure 2.15: The spectra of the two pumps used in the experiment, measured with an ocean optics USB2000 spectrometer supporting a resolution of 0.1 nm per pixel. The older 500 W pump had a significantly broader spectrum than the 100 W pump diode.

coupling aspheric to sit at a different distance to the end of the PCF, compared to the 70 µm PCF, in order to couple the seed beam into the core. This meant that each fibre needed to be treated separately. To simplify the problem, it was decided to put the pump coupling aspherics in the best position to optimise the seed beam, and then add and adjust optics in the pump beamlines to get the best pump beam coupling possible. It was accepted that the optimum slope efficiencies of the PCFs may not be achievable due to non-ideal pump coupling. However, this small detail could be addressed by later experiments. To give extra degrees of freedom, while optimising the pump beam coupling, a telescope was added to allow adjustment of the pump beam's diameter and divergence, as seen in figure 2.16.

The pump-coupling-aspherics were positioned closer than their focal lengths to the ends of the PCFs to allow collimation of the seed beam emitted from further inside the core. This produced a seed beam with better spatial quality (in the case of the 70 µm PCF) and improved the amount of pump beam coupled to the inner cladding. Once suitable seed beam profiles were obtained from each fibre using a collimating ashperic lens, a beam propagation model was constructed in MATLAB, which used the measured beam diameters and divergences of each pump beam. Using simple ABCD transfer matrices to describe the pump beam line, several lens combinations were trialled for the telescope, using what was available in the laboratory. The MATLAB script estimated where the lenses should go with respect to each

PCF to size and collimate the beam correctly. The inner-cladding mode-field-diameter and numerical aperture were used to estimate what beam diameter and divergence were needed upon reaching the PCFs to realise good pump coupling.

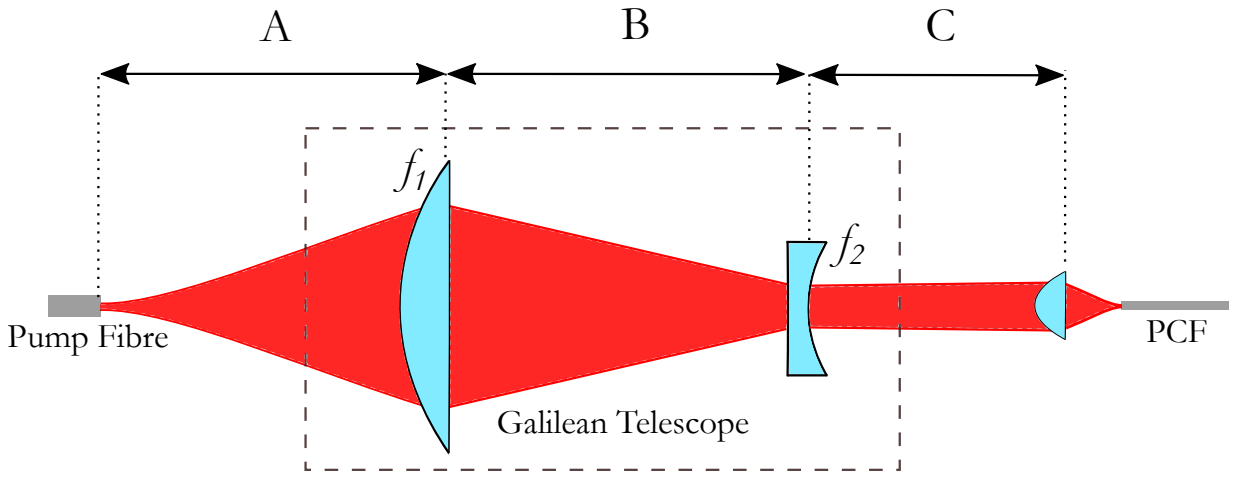


Figure 2.16: The pump beam telescope. 'A' is the distance from the pump beam delivery fibre to the first lens of the telescope. B is the length of the telescope, 'C' is the distance from the end of the telescope to the pump coupling aspheric lens. f_1 and f_2 are the focal lengths of the two telescope lenses.

After several iterations, it was decided that a 26 mm aspheric pump coupling lens, along with a Galilean telescope, would provide sufficient pump coupling and seed beam collimation. The telescope used for each PCF was slightly different due to the extra path length travelled by one of the pump beams and the relative position of the dichroic mirrors with their finite 25 mm apertures. Dichroic mirrors are optics, which are coated to allow for a high reflectivity at one wavelength, whilst allowing another wavelength to see high transmission. They were used to transmit the 976 nm pump beam, while reflecting the 1032 nm seed beam.

2.5 Summary

In this chapter, the challenges facing solid state lasers and their use in laser wakefield acceleration were discussed. Solid state lasers are able to deliver high peak power pulses to drive a laser wakefield, but are unable to do so at high repetition rates or wall plug efficiencies. Fibre lasers are seen as a solution to address these drawbacks. The factors that influence the design of a fibre were discussed and emphasis was placed on achieving single mode operation. Fibres that are made to produce the fundamental mode only are useful for applications where multiple fibre outputs can be combined together. For combination, the intensity distributions

of each combining beam must be matched to prevent a spatially varying polarisation state from being produced. Therefore, it was important to understand what modes a fibre is able to produce and how to filter out the desired mode. Photonic crystal fibres (PCFs) were subsequently introduced as the next iteration of fibre laser technology and the challenges faced by them in achieving larger core diameters were discussed. The arrangement and size of the air holes in a PCFs structure can be used to design PCFs with very large core diameters. Larger core diameters will guide greater numbers of higher order modes however, so one of the challenges faced by designers is how to produce a PCF with a large core diameter that is exclusively single mode. The solution to this problem is to operate in effective single mode operation. State of the art for PCF core diameters which are effectively single mode, currently sit at 135 μm [32]. Large pitch fibres aim to surpass this core diameter whilst maintaining effective single mode operation but, it is unlikely that these will results in millimetre scale core diameters. Hence, the coherent combination of multiple fibres is seen as the next logical step to increasing the amount of gain medium whilst maintaining single mode operation.

Finally, the two Yb-doped PCFs used for the coherent combination experiments in this thesis, the 70 μm and 85 μm PCF, were presented. The optics used to decouple the seed beam from each PCF also acted as the pump beam coupling lens. The dual functionality of this lens meant that there was a requirement to optimise for either the seed beam collimation or the pump beam coupling. Because the aim of the work in this thesis was to coherently combine the outputs of two PCFs, the seed beam collimation was chosen to be optimised at the expense of the pump beam coupling. The pump lasers used for the 70 μm and 85 μm PCFs were 100 W and 500 W laser diodes respectively. Each pump laser had a telescope in the beam line to assist in matching the pump beam to the fixed position aspheric lens, previously optimised for the seed beam collimation. The telescope was identified as the solution to overcome the poor pump beam coupling, whilst retaining the seed beam collimation.

Chapter 3

Coherent Combination of Fibre Lasers

As mentioned previously, in Chapter 1, the one limiting factor which has prevented fibre laser technology dominating is its peak power performance. Simulations [39] have shown that a single fibre amplifier with a core diameter of 80 μm will have a maximum theoretical output peak power of $\sim$ 10 GW. This limits the usefulness of single fibre and chained amplifier systems for reaching terawatt peak powers. However, one solution to this problem is to combine the outputs of many fibres together using coherent combination. This makes it possible to overcome the 10 GW peak power limit [39] of a single fibre.

Coherent combination is the process of carefully controlling the phase relationship between the combining beams such that they are mutually coherent. By carefully controlling the path length each beam travels from the point of splitting to the point of combination, the beams can be stabilised with a phase difference of zero, giving perfect recombination. Coherent combination has been demonstrated in both passive and active stabilisation architectures, both of which act to keep the combining beams locked with a phase difference of zero. By contrast, incoherent beam combining uses beams where their coherence to one another is not controlled. So even if an S- and P-polarised beam, for the sake of example, were incident into a polarising beam splitter (PBS) via different ports, the brightness of the combined beam will have increased, but the beam will have an undefined polarisation state, preventing further stages of combination. This is a demonstration of the second law of thermodynamics, where the entropy of a system cannot be decreased over time. The two incoherent beams are in two different states, so cannot be collapsed down into a single, well-defined state. Beams

which are in the same state can be coherently combined repeatedly in a cascade-like fashion, enabling the power to be scaled beyond that of an incoherent beam combining architecture. In this Chapter, a summary of the different methods of coherent combination found in the literature are presented. Tiled aperture setups and filled aperture setups are introduced, stating their advantages and drawbacks alike. The methods used to stabilise the setups will also be presented, culminating with the stabilisation method employed in this thesis.

3.1 Tiled aperture setups

Tiled aperture setups split a beam into multiple amplifiers, and the outputs are combined in the far field only. The output of each amplifier is aligned to be propagating in parallel to every other beam, creating an array of beams [46, 80]. This array is treated as a single beam when turning mirrors are encountered, directing the array of beams towards the target area for combination, see Figure 3.1.

The coherent combination happens in the far field and is performed using a focusing optic. Additionally, tiled aperture coherent combination has the benefit of spreading out the energy of the beams across an optic, meaning the damage threshold of the optic would not be surpassed. A drawback of tiled aperture coherent combination is the low energy contained in the combined beam. A square 8×8 array of tightly packed Gaussian beams was demonstrated to contribute only 30-40% of the total beam energy into a combined beam [46]. This was because the focusing optic acted as a Fourier transformer, with the Fourier transform of an array of Gaussian beams leading to a central beam with several lobes around it, as illustrated in Figure 3.2. Because of the amount of energy outside the central lobe, more work is needed to improve the energy in the central beam. If Gaussian beams are instead used in a hexagonal arrangement of beams, then the amount of energy that can be recovered in the central lobe increases towards $\sim 50\%$. This figure can be improved further by changing the individual beam profiles from those of a Gaussian to those of a top hat [63].

The optimal scenario would be to improve the packing of the beams and adjust their intensity such that their combined intensity distribution approximated that of Gaussian beam, which would Fourier transform to give another Gaussian carrying the majority of the beam

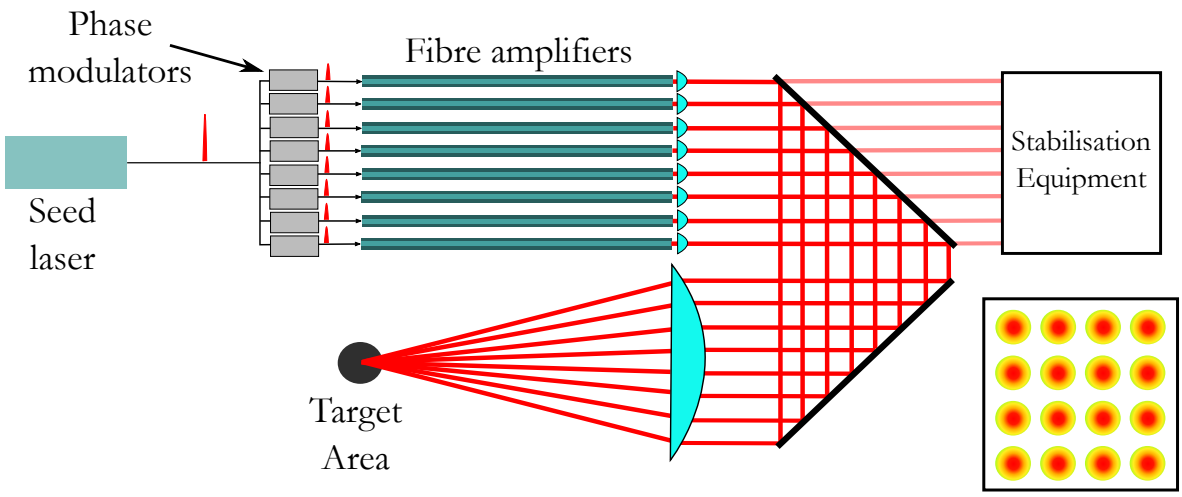


Figure 3.1: A schematic diagram of a tiled aperture array. An oscillator pulse is split into N arms, each of which passes through a phase modulator and their own amplifier, making this an active stabilisation setup. After amplification, the output beams are individually collimated using lenses and made to propagate in parallel with one another. The array of beams is then treated as a single beam, which can be reflected using mirrors. Inset is an example of the beam profile that would be seen on a camera after the fibre amplifiers. The path length for each pulse starts from the point of splitting and ends at the point of combination. When the beams are due to be combined, a focusing optic is used to bring all of the beams together at the focus. A lens has been shown, but a focusing mirror could be used also. Phase modulators are used to allow the phase difference to be controlled with respect to the other fibres. The stabilisation equipment contains the CCD required to observe the beams and the image processing equipment needed to extract the phase difference between the beams. More details follow about how the phase difference can be determined.

power. Given that this level of packing is not feasible with the current state of the art, the side lobes will always be present in tiled aperture coherent combination. This presents tiled aperture setups with an additional design problem, as the beam power of the side lobes cannot be disregarded. As much as 70% of the total energy is carried by the side lobes in the example of the 8×8 array. An additional consideration will be the size of the beam transport optics, as the more fibres that are used, the larger the mirrors and the focusing lens will need to be to ensure all the beams are captured and not spatially clipped. For ultrafast lasers, pulse compression will need to be considered also, which will add complexity to the setup. If the setup was built to support sub-100 fs pulses, the combination lens (and any other post-compression lenses) could be replaced with a parabolic mirror to avoid adding material dispersion and non-linear phase to the compressed pulses.

A small 1% reflection/transmission is taken from the array of beams for stabilisation purposes, leaving the rest to propagate to the target area. There are two distinct methods of phase stabilising a tiled aperture system: the first is quadriwave lateral shearing interfer-

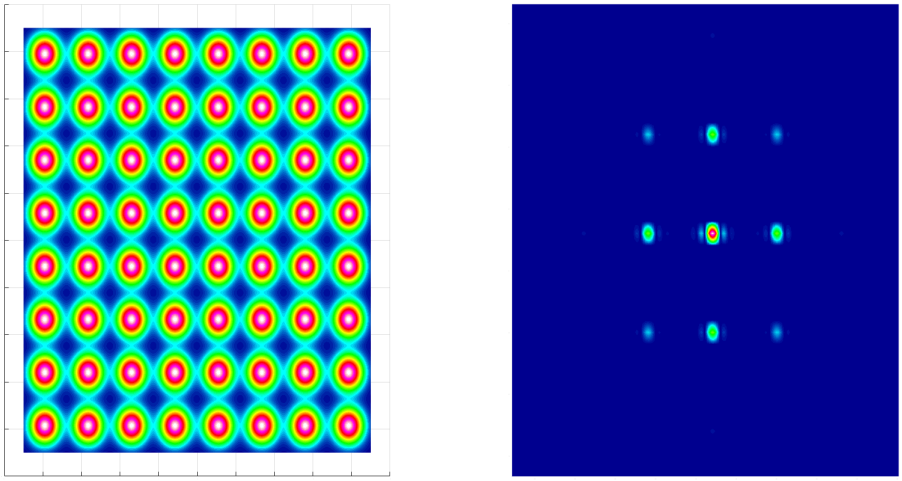


Figure 3.2: (a) An 8×8 array of simulated Gaussian beams. (b) The Fourier transform of an array of simulated Gaussian beams. Given that a lens acts as a Fourier transformer, the coherently combined beam profile for the tiled aperture setup, assuming zero phase difference between the combining beams, will resemble that of Figure (b). In this particular example, only $\sim$ 30% of the total energy in (b) ended up in the central lobe, with the rest being distributed into the side lobes.

ometry [81]; the second is self-adaptive digital holography [82]. Both of these methods offer efficient use of locking mechanisms. Unlike the filled aperture setups seen in Section 3.2, they allow a large number of fibres to be locked together using only a single CCD detector, which is placed to observe the 1% reflection/transmission. The individual beams in these setups are made to either interfere with each other, in the case of quadriwave lateral shearing interferometry, or interfere with a single reference beam in the case of self-adaptive digital holography. The interference between the beams produces interference fringes on the surface of the CCD, which can be analysed in software using image processing. By identifying the fringes and the relative position of their maxima and minima compared to all of the other beams, the phase difference between the beams can be determined and minimised. The complex calculations for these two stabilisation methods are performed by software, so even though there is only a single detector and computer present, each beam will need a separate feedback loop in software. Technically, the individual pixels on a CCD can be treated as separate detectors, but the proliferation of CCD technology in the mass market photography industry has seen these high-density detectors become larger and just as affordable as single photodiodes. These two methods of stabilisation, along with tiled aperture coherent combination setups are alternative methods of addressing the problem of coherently combining

a large number of fibre amplifiers. However, as tiled aperture setups and their stabilisation methods do not feature in the work presented in this thesis, they have been discussed here as an outline of alternative methods of coherent combination and will not be discussed any further.

3.2 Filled aperture setups

For optimum energy recovery, the polarisation state of the beam is exploited to perform the pulse splitting and done so with a polarising beam splitter (PBS). Filled aperture setups differ from tiled aperture setups in that the temporal and spatial overlap happens in the near field and continues into the far field. Filled aperture coherent combination setups have been demonstrated many times [38, 83, 84, 85], with the two varieties of filled aperture setup demonstrated schematically in Figure 3.3.

The first method shown in Figure 3.3(a), makes use of a diffractive optical element such as a diffraction grating, to disperse the beam into different orders. Each beam is amplified separately and is recombined in the opposite way. CW operation has been demonstrated successfully with 5 and 15 fibre amplifiers, giving 1.9 kW and 600 W output beams respectively [86, 87]. One of the benefits of using a DOE is that many beams can be split and combined with a single optic. This can reduce the number of components required for a combination setup. However, there are two drawbacks to using a DOE. The first is the additional output beams from the higher orders. The power lost to these higher order beams is small (5-10%) [88] compared to the zeroth order, but it is a significant amount of power, especially when working at kilowatt level average powers. At these power levels, safe dumping of the higher order beams must be considered. The second drawback is when ultrafast lasers are used, the beams of the higher orders (other than the zeroth order) acquire an angular chirp, which gets larger as the orders get higher. To use these beams, the angular chirps would need to be compensated. Using a similar principle to the stretcher, as shown in Figure 4.2, each beam would need a 4-f lens setup and an additional DOE to create a grating-image separation of D=0 [63], as shown in Figure 3.4. The total number of DOEs used, assuming N beams,

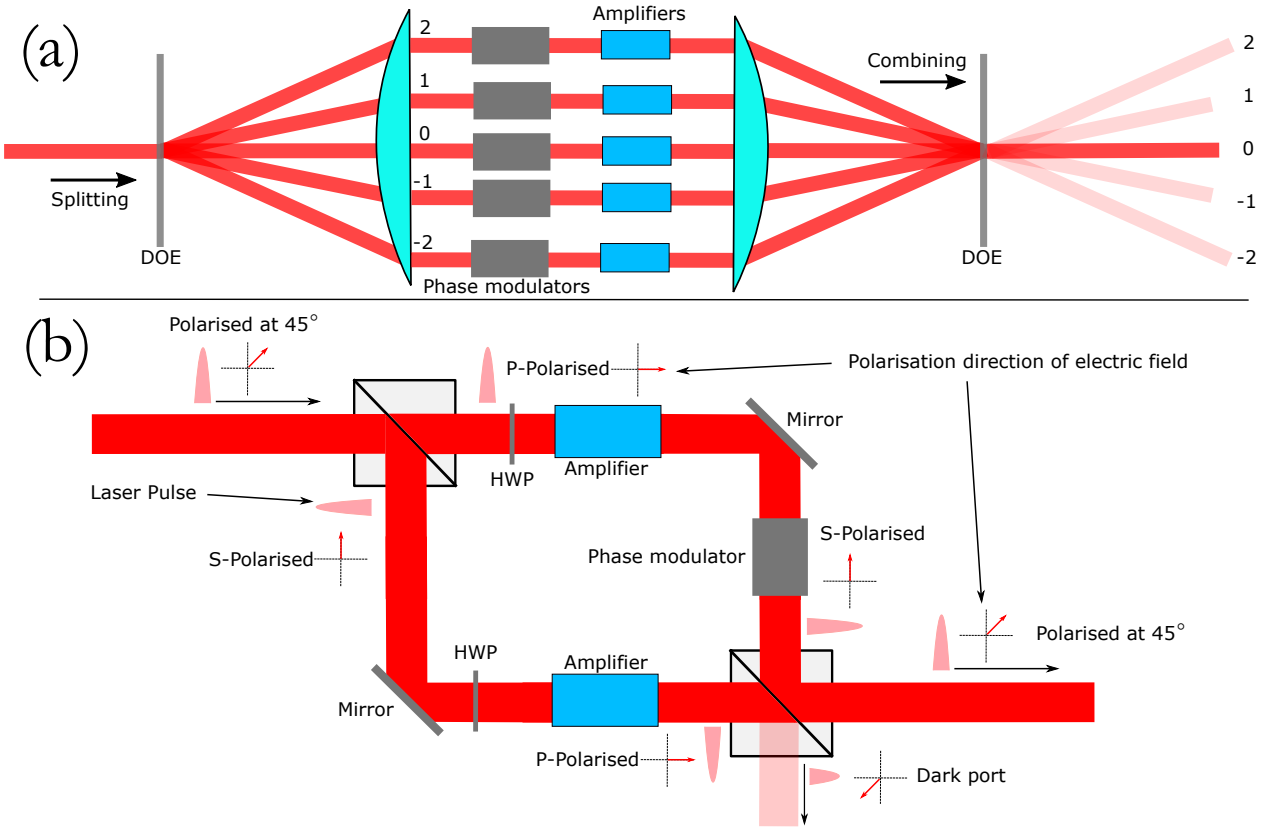


Figure 3.3: A schematic diagram showing two distinct methods of filled aperture coherent combination. (a) A filled aperture setup using the diffracted orders of a Diffractive Optical Element (DOE). All of the diffracted beams are collimated using a lens, with the majority of the beam power being sent into the zeroth order beam. Each beam is sent into a separate amplifier, before the reverse process is used to recombine the beams. Beam combination using a DOE does come with the drawback that, after combination, there will still be some diffracted beams leaving the DOE, as shown in Figure (a). The advantage over tiled aperture setups is that beam energy lost to the diffracted orders (5%) is an order of magnitude less than that lost to side lobes (70%) as seen in Figure 3.2. (b) A filled aperture setup utilising PBSs in a Mach-Zehnder interferometer. These types of setups can be used for both passive and active coherent combination setups, due to their simplicity over setups such as that shown in (a). By carefully controlling the polarisation of the beams using half-wave plates (HWP) as they propagate through the interferometer in free space, the split beams can be recombined into a single beam with negligible loss. Also shown in these setups are the phase modulators used by the stabilisation mechanism to zero the phase difference between the two beams. The phase modulators ensure each beam has travelled the same path length through the different arms of the interferometer.

would be 2N. This number does not include the additional lenses required for the 4-f setups. Splitting using a 1:N beam splitter would require N beam splitters and half wave plates, with an additional N beam splitters required to reverse the process for N:1 beam combination, giving 2N.

The second method shown in Figure 3.3(b) takes advantage of the polarisation of the beam. A beam with its polarisation oriented at 45° to the transmission axis of a Polarising

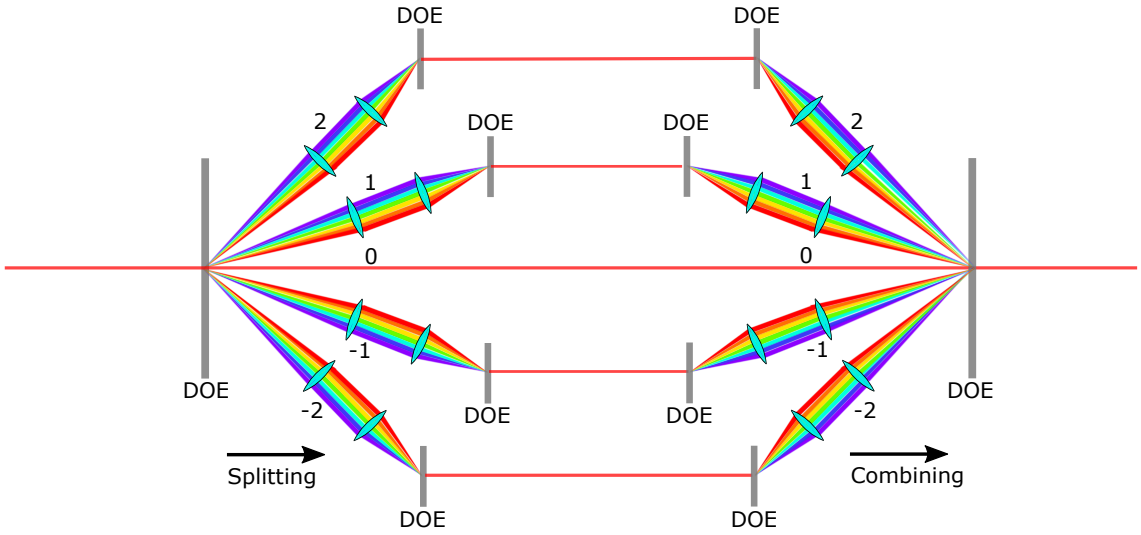


Figure 3.4: A schematic diagram illustrating how a 4-f lens setup can be used to compensate for the angular chirp, when using a DOE as a beam splitter for ultrafast laser pulses. Each beam has been labelled with its diffracted order. Each beamline, except for the zeroth order beam, makes up a two grating stretcher, with a stretcher length of D=0.

Beam Splitter (PBS) will be split into two beams with S- and P-polarisation states, as shown in Figure 3.3(b). For linearly polarised, ultrafast laser pulses with S- and P-polarisations, the two pulses can also be combined using a PBS. If the pulses enter the polarising beam splitter possessing a phase difference of zero, then they will combine to form a pulse with a polarisation state oriented 45° between the S- and P-polarisation states. To save repetition, when a beam is said to be at 45°, this is the polarisation with respect to the transmission axis of the pre or proceeding PBS, which always will be horizontal or vertical. To perform a cascade of pulse splitting to seed N amplifiers, a Half Wave Plate (HWP) can be positioned into both beams after the splitting PBS. The HWP can be used to rotate the beam's polarisation state to 45° in relation to the proceeding PBS's transmission axis, as shown in Figure 1.9. For combination, the reverse process is performed, provided that a pulse incident on a PBS is oriented either in S- or P-polarisation states to be properly reflected or transmitted at the PBS. This is demonstrated in Figure 3.5, using a Mach-Zehnder interferometer layout.

At each combination stage, there are two output beams, one is the desired combined beam, and the other is the rejected beam. The rejected beam has orthogonal polarisation to the desired output beam and exists due to the non-perfect combination of the two beams. This could be due to the phase difference between the frequency components of the pulses

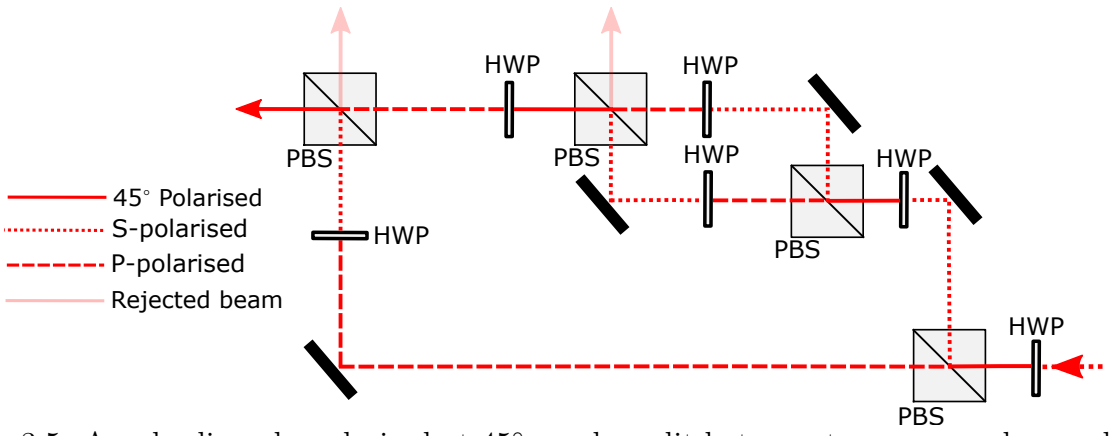


Figure 3.5: A pulse linearly polarised at 45° can be split between two arms and recombined at a PBS to produce the same 45° polarisation state. This splitting and combination can be cascaded as shown in the upper arm. At each combination stage, there are two output beams, one is the desired beam, and the other is the rejected beam. The rejected beam has orthogonal polarisation to the desired output beam. For the beams to be optimally combined, there must be a phase difference of zero between the combining beams, and this must be true for both interferometers.

not being identical, as discussed later in Chapter 4, or the combining beams not having pure linear polarisation states. Polarisation splitting and combining rely upon the beams having a well-defined linear polarisation state to ensure that the majority of the energy is recovered. Any amplifiers used will also have to preserve the beam's polarisation state and should be chosen for their ability to do so. Demonstrations of coherent combination of four PCFs [38], followed by eight PCFs [42] have been shown in the literature. These two coherent combination experiments demonstrated that such systems are capable of producing quasi-continuous wave beams with pulse energies four times greater than the damage threshold of a single channel [38, 42]. Additionally, increasing the number of fibre amplifiers in a coherent combination system leads to that system's efficiency flattening out to a fixed value [52]. It has been shown that spatial multiplexing of fibre amplifiers can be scaled linearly to very large numbers, leading the way for further experiments using a large number of fibre amplifiers.

A polarisation splitting setup is straightforward to assemble and align, but it needs additional optics in order to build large setups. Furthermore, as the number of fibres to be combined increases, so will the peak power of the combined pulse. To prevent damage to the optics, the size of the optics will need to increase also. However, the optic size can be increased gradually until, after the final stage of combination, the largest optics are required to transport the combined beam. Tapering the increase in optic size as beams are combined will help to decrease the cost. Unfortunately, larger optics are an unavoidable consequence

of transporting high energy pulses and is the same problem experienced by Ti:Sapphire systems. To combine an array of fibre amplifiers, a PBS would be required for every beam split into two, and every two beams combined into one. So for N beams combining in a cascade of pairs, N-1 combination PBSs are required [63], with this number being doubled again to account for the splitting PBSs, giving 2N-2 PBSs in total. Taking the example of 16 beams, after the first round of combination, 8 PBSs would be required. Four would be required for the next stage, followed by two and finally one more, giving a sum of $8 + 4 + 2 + 1 = 15$ PBSs. Regardless of whether a DOE or polarisation splitting architecture is chosen; for all filled aperture setups there are two distinct stabilisation regimes in which to choose from. These are passive and active stabilisation. Both of these regimes are discussed in turn in sections 3.3 and 3.4.

3.3 Passive stabilisation of beams

Passive stabilisation takes advantage of the layout of the components to ensure that the split pulses travel the same path length. The split pulses consequently arrive back at the splitting Polarising Beam Splitter (PBS) with a phase difference of zero. Typical solutions either use retro-reflections to send pulses back along the paths in which they travelled or by using a ring based system known as a Sagnac interferometer.

3.3.1 Sagnac interferometer

A Sagnac interferometer is a closed ring interferometer. The separate pulses can be sent in opposite directions around the ring [89, 90] as shown in Figure 3.6. In Figure 3.6, the pulses travel in opposite directions around the ring and will experience the same path length and encounter the same optical components. Because the pulses have been treated the same, they arrive back at the splitting PBS, with a phase difference of zero. However, the 10 GW peak power limit for pulses with pulse durations < 1 ns applies to these amplifiers, so a chain of amplifiers cannot be used in the ring to get to large peak powers beyond 10 GW.

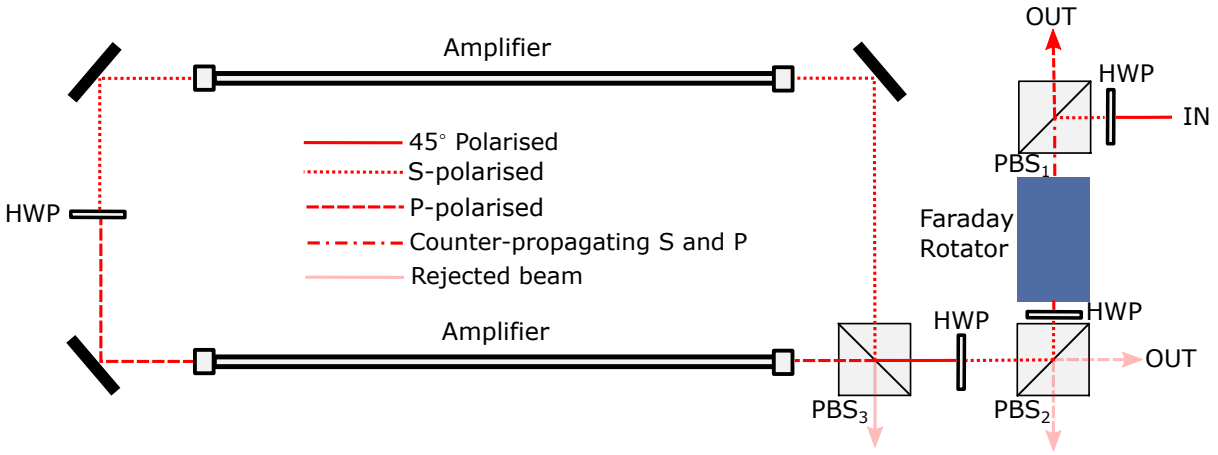


Figure 3.6: A schematic diagram of a passively stabilised coherent combination setup. Here we see a Sagnac interferometer with two amplifiers in the ring. A pulse is reflected at PBS_1 and travels through a Faraday rotator before it is split equally at PBS_3. Each copy of the pulse travels through both amplifiers in turn before returning to PBS_3. As each pulse travelled the same path length, they will have a phase difference of zero when recombined at the PBS_3. The combined pulse travels back through the Faraday rotator, whose purpose is to rotate the polarisation of the pulse by 90° with respect to the incident polarisation after two passes through it. This rotation allows the amplified pulse to be extracted via transmission through PBS_1.

3.3.2 Divided Pulse Amplification

For completeness, divided pulse amplification is included here. This is another method of splitting a pulse into many copies, which are amplified separately. This splitting occurs temporally, creating a train of pulses, which propagate along a single beam line and witness the same amplifiers. The splitting can be achieved by using small Mach-Zehnder [91] or Michelson [91] interferometers or by using birefringent crystals [92, 93, 94] as shown in Figure 3.7.

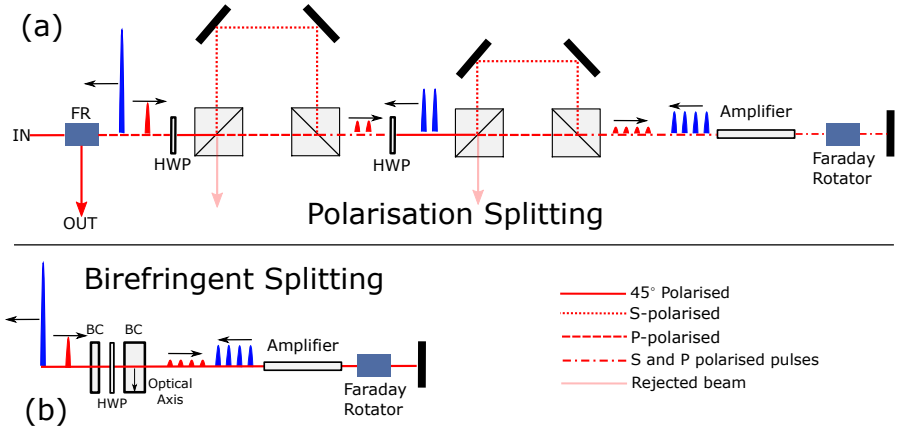


Figure 3.7: (a) The incident pulse is split many times using small Mach-Zehnder interferometers with arms of different path lengths. The number of pulses is controlled by the number of interferometers used. At every stage, the number of pulses is multiplied by two. (b) The same effect can be achieved using Birefringent Crystals (BC) with the fast and slow axis aligned 45° to the beam's polarisation state.

Once the pulse train has been amplified, it can be sent back through the arrangement of interferometers or birefringent crystals to be recombined into a single pulse. This setup works well for a beam, which takes a double pass through an amplifier using a Faraday rotator and retro-reflector as shown in Figure 3.7(a). A beam returning from the amplifier which is counter-propagating to the incident beam should already be aligned to the network of interferometers / birefringent crystals. It is sensible to use the same interferometers for splitting and recombining the pulses, as all of the path lengths will be exactly right to give zero phase difference between the combining pulses. Creating a second set of interferometers to combine the beams rather than relying on a double pass through the first set, as seen in Figure 3.7, is not impossible but is challenging. The second set of interferometers would need to have their path lengths matched to less than $n\lambda_0 \pm \lambda/20$, with n < 15, to recover $> 95\%$ of the energy into a linearly polarised pulse [95]. Without the ability to scan the length of the interferometers, it will be difficult to match the path lengths of the second set exactly. Thus some beam power would be lost into a rejected beam. Divided-pulse amplification has been demonstrated both actively and passively, with pulse splitting [92, 93, 94] using either polarisers or birefringent crystals, as shown in Figure 3.7B and 3.7C. Divided-pulse amplification, has been demonstrated with a pulse energy of 1.1 mJ from just two pulse copies [96]. Active stabilisation of a system has improved upon this pulse energy and has demonstrated a pulse train of four split copies, being split between 8 channels of a coherent combination fibre experiment. A total output pulse energy of 3.7 mJ was improved to 12 mJ by the use of actively stabilised, divided pulse amplification [42].

3.4 Active stabilisation of beams

Active stabilisation uses equipment to monitor the phase difference between two pulses and to take corrective action to minimise the phase difference, bringing the pulses back into phase. The monitoring and corrective action combined forms a closed feedback loop. As mentioned previously, the phase difference between the two pulses is linked to the path length each has travelled. If the pulses travel along different paths to get the point of combination, then environmental conditions will individually effect each path causing them to vary in length

independently. Therefore active stabilisation is required to ensure the pulses remain locked in phase by changing the path length using a movable mirror [38, 83] or by applying a voltage to a phase modulator [46, 80]. A passively controlled system could be reconfigured to use an active stabilisation system if desired.

To stabilise the combined beams, a different approach is required from that for tiled aperture coherent combination. One approach is to use Locking of Optical Coherence by Single detector Electronic frequency Tagging technology (LOCSET), and another approach is Hänsch-Couillaud polarisation locking. LOCSET requires that each beam to be combined, has its path length to the combination optic dithered by a piezo-electric mounted mirror or phase modulator. Each beam is dithered at a unique fixed frequency, which tags the beam. The combined output beam is then observed with a photodiode, whose output can be analysed using computer software. Dithering a channel at a known frequency allows a computer to demodulate it by heterodyning the combined signal with the modulation frequency. This isolates that particular channel's impact on the output power. Feedback is used to optimise the output power by controlling the phase difference between a particular channel and every other channel, producing a coherent beam [97]. This method is powerful because of the ability to lock a large number of beams together, using only a single detector. LOCSET works well for both tiled aperture [98] and filled aperture setups [85, 99], and particularly for the coherent combination of beams using DOEs, as shown in Figures 3.3 and 3.4.

However, the locking electronics used to perform the frequency tagging phase stabilisation for LOCSET are very expensive. Furthermore, as LOCSET was not used to coherently combine pulses in this thesis, it will not be discussed further. Hänsch-Couillaud polarisation locking, on the other hand, is a simpler stabilisation method, as the electronics make extensive use of simple Proportional Integral Differential (PID) controllers. However, the drawback of Hänsch-Couillaud polarisation locking is that it can only combine two beams at a time. For multiple beams, pairs of beams are combined in a cascade-like fashion, with a detector and feedback loop required for each pair of combining beams. Using the previous example of 16 beams, 15 detectors and feedback loops would be required to combine 16 beams into one.

However, given that the experiments in this thesis only considered the coherent combination of two PCFs, Hänsch-Couillaud polarisation locking was satisfactory.

3.4.1 Hänsch-Couillaud polarisation locking

Hänsch-Couillaud polarisation locking is used in filled aperture setups, which use polarisation based splitting and it is used to lock the phase of two beams together. Figure 3.8 shows an example of two pulses being phase locked together. For a system to phase lock, the detector needs to produce an error signal, based upon the phase delay between the two pulses. To generate this phase information, the polarisation state of the combined beam is analysed. This analysis is done using a Hänsch-Couillaud detector [100], shown in Figure 3.9, which is a pair of photodiodes observing the two output ports of a PBS.

The intensity of light seen by either photodiode can be used to determine the polarisation state entering the Hänsch-Couillaud detector. A Quarter Wave Plate (QWP) is added before the Hänsch-Couillaud detector to determine the direction in which the phase difference between two pulses occurs. The QWP converts the elliptical polarisation state of the combined beam into a linear polarisation state with an angle rotated away from 45°. An alternative way of viewing it, is that the QWP converts both of the orthogonal linearly polarised beams into circularly polarised beams with opposite chirality. Providing the two beams have equal amplitude, these beams interfere with each other and produce a linear polarisation state. The angle of the linear polarisation state is dependent upon the phase difference between the two beams. This process maps the phase difference of the beams onto the angle of the linearly polarised beam, which can be detected using the Hänsch-Couillaud detector. The operation of the Hänsch-Couillaud polarisation locking system is demonstrated schematically in Figure 3.10.

The Hänsch-Couillaud polarisation locking system is suitable for locking two pulses together, because the magnitude and sign of the intensity difference between the two photodiodes provide information about the direction and the magnitude of the phase difference between the two pulses. Therefore, the magnitude and sign of the intensity difference between the photodiodes can be used as an error signal for a PID controller. The PID controller uses

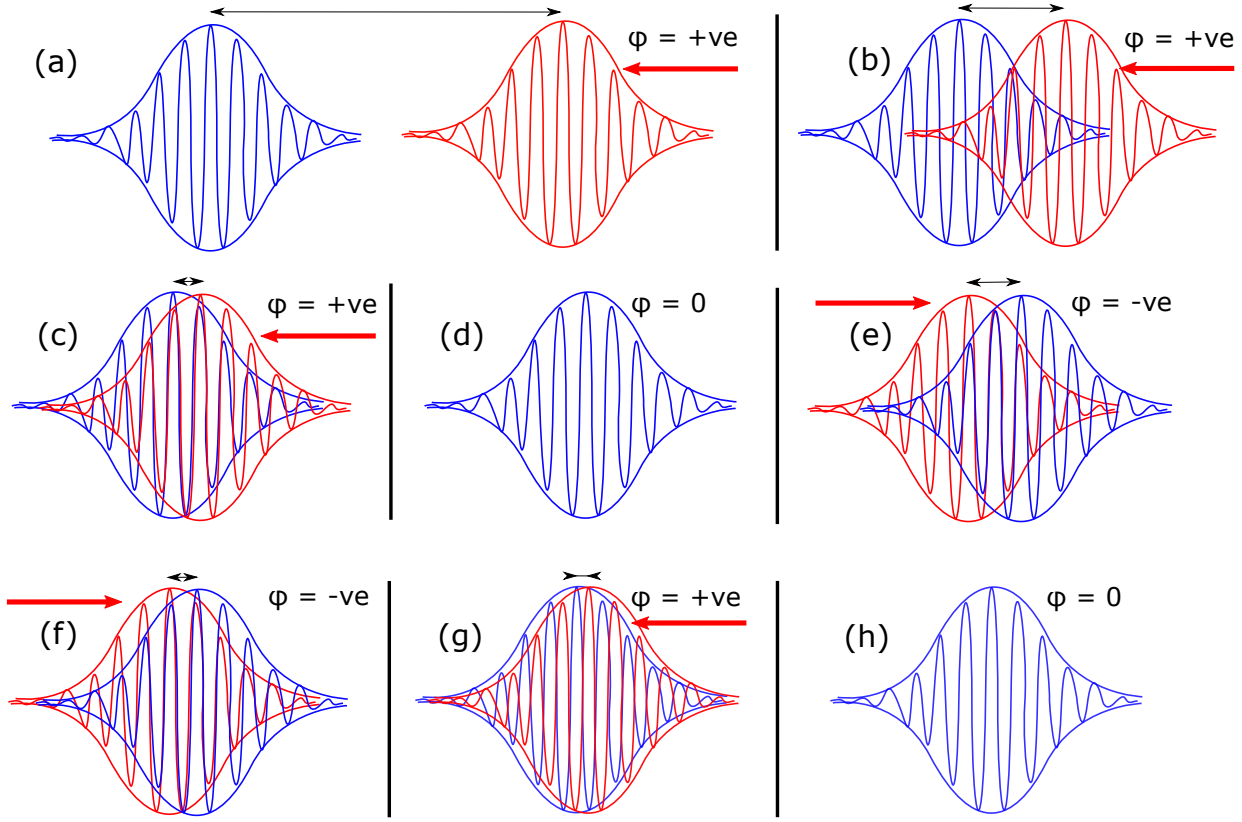


Figure 3.8: Phase locking aims to minimise the phase difference between two laser pulses, distinguished by their colours. The red arrow indicates the direction of travel of the red pulse, required to reduce the phase difference. ϕ indicates the relative phase difference between the two pulses and the black arrow indicates the relative size of the phase difference. Initially, the pulses need to be overlapped to allow them to interfere. This is demonstrated in (a)-(f) with the pulses being scanned through each other to detect and find the point of maximum overlap, (d). Maximum overlap is observed when the visibility of the interfering beams viewed on the photodiode through a PBS is maximised. Once as close to state (d) as can be achieved manually, the interference between the combining pulses represents a mapping of the phase delay to the polarisation state. A Hänsch-Couillaud detector [100] is used to observe the polarisation state and infer the phase delay between the two pulses. The locking system can only stabilise the pulses from within $\lambda/4$ of optimal overlap. The phase locking equipment uses a PID controller to change the path length of the red pulse to reduce the phase difference to zero as shown in **g** to **h**.

this error signal to generate a control signal which can drive a phase shifter to adjust the path length of one beam and bring the two pulses back towards a phase difference of zero. The phase shifter is usually a movable mirror mounted on a piezo-electric ceramic. This is because a mirror can operate at the high combined average powers seen in coherent combination setups. Also, the piezo-electric mounted mirror can be mounted on translation stages to compensate for larger differences in path length variation. Electro-optic modulators can be used to compensate for phase delay, but are usually limited to phase compensation of up to 2π, which can prevent stable locking in a system affected by noise and thermal variations.

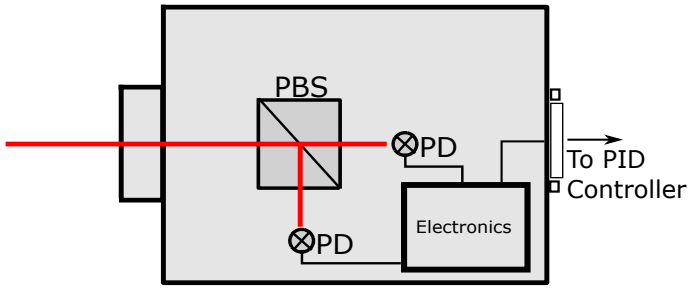


Figure 3.9: The layout of a Hänsch-Couillaud detector. Red is the path taken by the incident beam, 'PD' represents photodiodes and PBS is a polarising beam splitter.

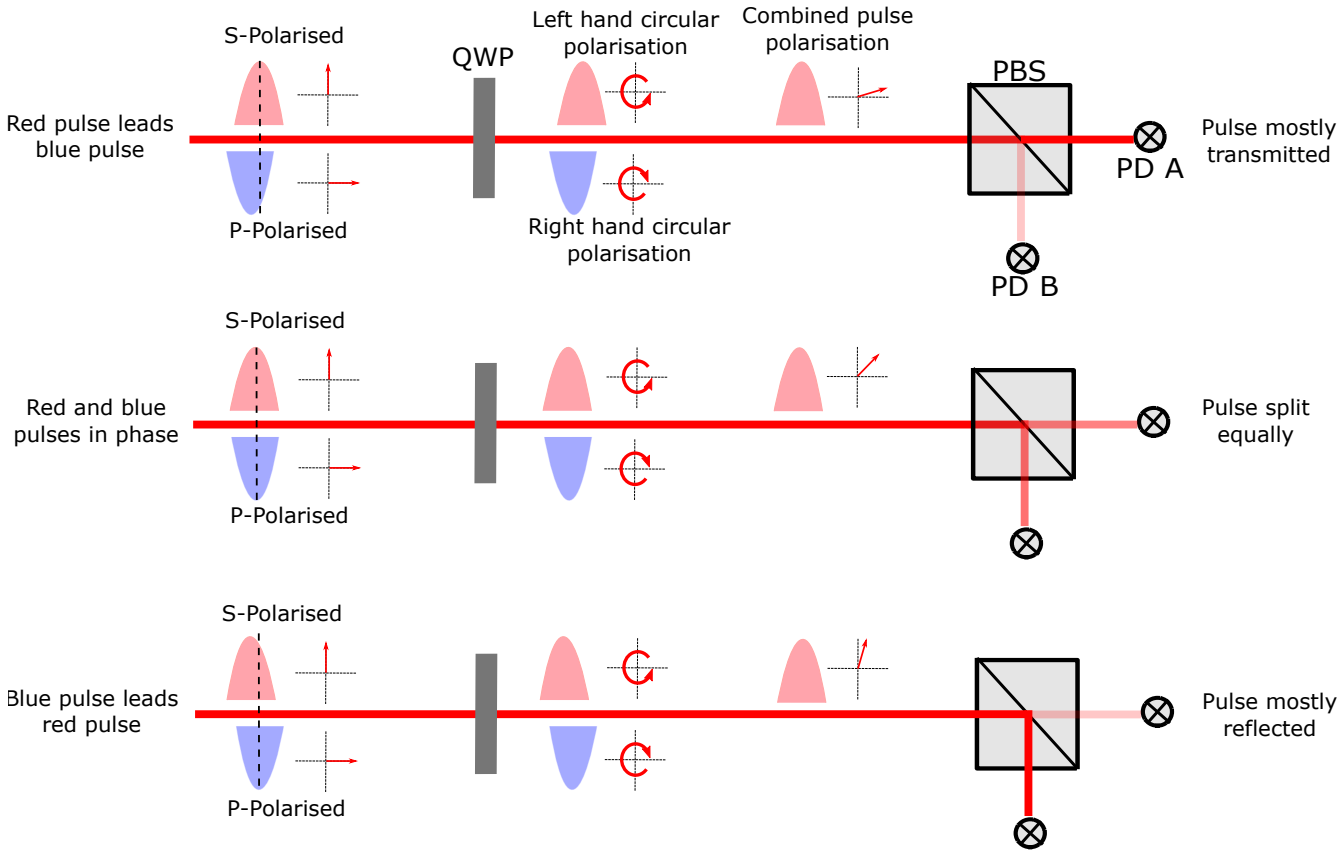


Figure 3.10: A schematic example of how the phase delay between the linearly polarised pulses can be mapped onto the angle of linear polarisation of the combined pulse. QWP is a quarter wave plate, PD A is photodiode A, PD B is photodiode B, and PBS is a polarising beam splitter. The polarisation of each pulse is represented by the small cross with a red arrow, with the direction of the pulse propagation (with respect to the polarisation state) coming out of the page. The two pulses of equal amplitude are sent through a QWP giving them both a circular polarisation state of opposite chirality. These interfere with one another to produce a linear polarisation state with an angle dependent upon the phase difference between the two pulses. When the red pulse leads the blue pulse, this causes the power to be mainly transmitted through the PBS. For a blue pulse leading, the combined pulse is mainly reflected at the PBS. For pulses with zero phase difference, the power seen by photodiodes A and B is equal.

For this thesis, a Hänsch-Couillaud polarisation locking system was the preferred choice to lock two photonic crystal fibres together due to its simplicity. The stabilisation system used a PID controller which monitored the error signal produced by the Hänsch-Couillaud

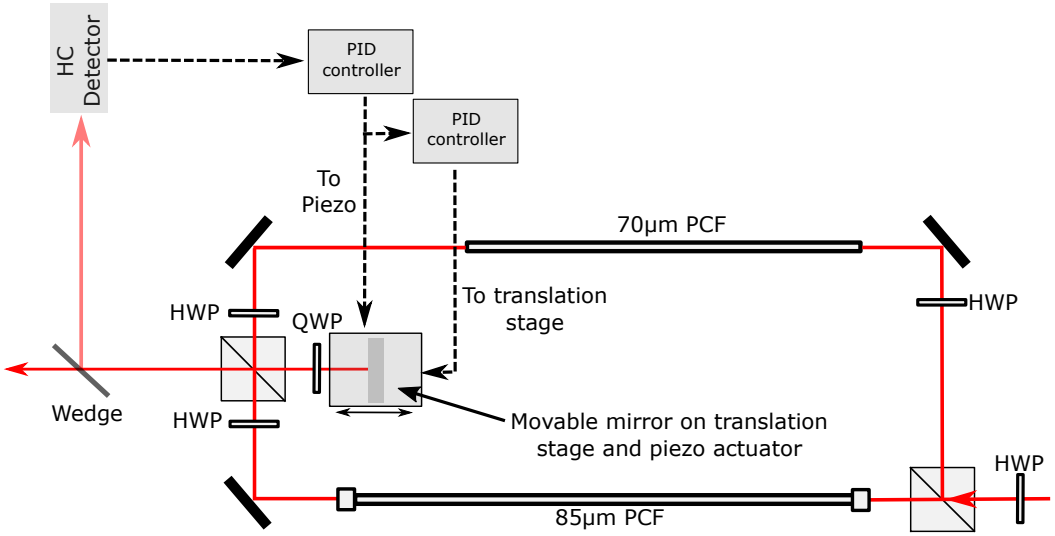


Figure 3.11: A schematic diagram of the Mach-Zehnder interferometer for the two PCFs and the Hänsch-Couillaud polarisation locking system used to stabilise it. An additional feedback system is present which treats the output from the first PID controller as an error signal. The second PID controller then performs two functions; the first is to keep the piezo-electric mounted mirror near the centre of its range of travel to improve its stability. The second is to enable the system to stabilise over very large path length variations and therefore improve the robustness of the system.

detector to generate a control signal. The PID controller used the control signal to drive a piezo-electric mounted mirror, positioned at the combination element of the interferometer, as shown in Figure 5.3. The piezo-electric mirror adjusts the path length of one of the arms of the interferometer, therefore affecting the signal observed by the Hänsch-Couillaud detector.

3.5 Summary

Coherent combination architectures fall into one of two categories: tiled aperture or filled aperture setups. Tiled aperture setups allow a large number of fibres to be coherently combined, while only using a single detector. This makes them very efficient in the use of optical components and electronic hardware. The main drawback is that $\sim 70\%$ of the total beam power is lost into the surrounding lobes of the combined beam, as demonstrated in Figure 3.2. The packing arrangement of the fibres could improve this figure, but even with the most efficient hexagonal packing of Gaussian beams, the loss still would be $\sim 50\%$ of the total beam power. Filled aperture setups allow all of the beam power to be captured in the final combined beam, but require more optical components for a similar number of amplifiers.

For the combination of two mismatched PCFs as covered in this thesis, the filled aperture architecture was chosen due to its simple construction. The locking mechanism chosen was

Hänsch-Couillaud polarisation locking, as it was a simpler and cheaper system compared to Locking of Optical Coherence by Single detector Electronic frequency Tagging technology (LOCSET). The drawback of using Hänsch-Couillaud polarisation locking is that a detector and feedback loop is required for every pair of combining beams. LOCSET is more complicated, but has the benefit of requiring only one detector to observe the final combined beam. However, the LOCSET system is limited by the number of different modulation frequencies available to tag each beam. Also, the cost of a LOCSET system is far higher that of a standard Hänsch-Couillaud polarisation locking system.

Chapter 4

Ultrafast laser theory

4.1 Outline

Ultrafast pulses are generated by oscillators, which have very low pulse energies, typically of the order of nanojoules. These seed pulses are sent through laser amplifiers to reach the joule-level pulse energies required for laser-wakefield acceleration. Amplifiers are materials which contain a lasing ion capable of being excited by a pump source and retain that energy for a short period of time. The seed laser pulse enters the medium and liberates energy through stimulated emission. The more amplifiers the pulse propagates through, the more energy it will acquire, but this process does not happen indefinitely.

The relationships between the frequency components of the laser pulse are described by their phases, normally defined in relation to the central frequency of the pulse's bandwidth. As the pulse propagates through a material, each frequency component will behave differently according to the refractive index, $n(\omega)$, of the material and the spectral intensity of the pulse. Both the refractive index and spectral intensity contribute to linear and non-linear effects on the pulse which are undesirable. The non-linear effects are intensity dependent so become more prominent with high peak powers. Given that the peak power depends upon the pulse energy and duration, changing the phase relationship between the frequency components of a transform-limited pulse will result in an increase in the pulse duration and a reduction in the peak power. To generate the highest peak powers, the pulse should be kept as short as possible. This means that control of a pulse's phase is required and methods to mitigate

the effect of non-linearities need to be considered. In addition, the large electric fields of the pulse can lead to damage in the amplifiers, so care must be taken to ensure the pulse's fluence does not exceed a material's damage threshold. In terms of the non-linear contributions, the most dominant is the non-linear phase and is defined as

$$\Phi_{\mathrm{NL}} = \frac{2\pi}{\lambda} \int_0^{\mathrm{L}} \mathrm{n}_2(\omega)\mathrm{I(z)dz}, \tag{4.1}$$

where I(z) is the intensity, L is the length of the medium and $\mathrm{n}_2(\omega)$ is the the non-linear refractive index of the medium (typically with values of 2.7×10^{-27} $\mathrm{W}^{-1}\,\mathrm{m}^2$ for fused silica) and is related to the intensity dependent refractive index of a material by

$$\mathrm{n}(\omega) = \mathrm{n}_0 + \mathrm{In}_2(\omega), \tag{4.2}$$

where n_0 is the refractive index of the material defined in the usual way. From Equation (4.1), the non-linear phase acts to spread out temporally the frequency components of the laser pulse in such a way that prevents compression back towards the transform limit. When seeking to achieve high peak powers, the non-linear phase can cause a drop in the maximum peak power achieved for a given pulse energy. To mitigate the effects of non-linear phase, other non-linearities and to prevent material damage, the intensity of the pulse needs to be reduced. Increasing a pulse's mode area reduces the pulse intensity and is one tactic used to mitigate non-linear phase and material damage. However, there is a restriction on how far this size increase can be taken due to physical limits on the size of the optical components. These limits stem from manufacturing limitations, tolerances and cost. Stretching the pulse in time can reduce the intensity considerably, and stretching factors of several thousand have been demonstrated [101]. Amplifying a pulse after it has been stretched is known as Chirped Pulse Amplification (CPA). A complete CPA system uses two components known as a pulse stretcher and pulse compressor, as was seen in Figure 1.8. Material dispersion also contributes to the chirp of a pulse and its impact should be considered when designing a CPA system. Any phase differences which cannot be removed by the pulse compressor will impact on the achievable pulse duration and subsequent peak power. This highlights the importance of having a pulse compressor, which is able to compensate for the initial chirp

from the stretcher and any other dispersion experienced by the pulse as it travels through an amplifier chain. To produce a compressor that is able to do this, as is done in Section 5.5, it is important to model an ultrafast laser pulse to understand what phase the pulse has acquired from stretching and amplification in order to compensate for them. This will be discussed in more detail in Section 4.2.1.

4.2 Ultrafast lasers

4.2.1 Pulse theory

In order to be able to predict what a known laser pulse will look like after propagating through a system, it is necessary to model a pulse and apply the effects of spectral dispersion, gain and non-linearities [60, 102]. Typically, pulse models begin by defining the electric field of the pulse

$$\mathbf{E(t)} = \mathrm{E_0 e}^{-\frac{4\ln(2)\mathrm{t}^2}{2\tau_{1/2}^2}} \mathrm{e}^{-\mathrm{i}\phi(\mathrm{t})_{\mathrm{init}}} \mathrm{e}^{\mathrm{i}(\mathbf{k}\cdot\mathbf{r}-\omega\mathrm{t})}, \tag{4.3}$$

where it is assumed to be a Gaussian where $\mathrm{E_0}$ defines the amplitude of the temporal pulse, $\tau_{1/2}$ is the pulse FWHM, $\phi(\mathrm{t})_{\mathrm{init}}$ is the initial phase distribution in time and $\mathrm{e}^{\mathrm{i}(\mathrm{kr}-\omega\mathrm{t})}$ is the pulse carrier. If any additional phase is added to this pulse, be it due to material dispersion or non-linear phase, the electric field is multiplied by an additional exponential to give

$$\mathbf{E_1(t)} = \mathbf{E(t)}\mathrm{e}^{-\mathrm{i}\phi_1(\mathrm{t})}, \tag{4.4}$$

where $\phi(t)_1$ is additional phase added. In the spectral domain this extra phase is expressed as

$$\mathbf{E_1}(\omega) = \mathbf{E}(\omega)\mathrm{e}^{-\mathrm{i}\phi_1(\omega)}, \tag{4.5}$$

where $\mathrm{E_1}$ indicates that the electric field is different due to the added phase. Because the bandwidth of the pulse is finite, an approximation of the additional phase can be made by Taylor expanding about the central frequency into a sum of terms

$$\phi_1(\omega) = \phi_0(\omega_0) + \phi^{(1)}(\omega)(\omega-\omega_0) + \frac{\phi^{(2)}(\omega)}{2}(\omega-\omega_0)^2 + \frac{\phi^{(3)}(\omega)}{6}(\omega-\omega_0)^3 + \frac{\phi^{(4)}(\omega)}{24}(\omega-\omega_0)^4 + \dots . \tag{4.6}$$

The first two terms are the total accumulated phase (ϕ_0) at ω_0 and the group delay ($\phi^{(1)}$). Both of these terms contribute to the phase difference between the carrier wave and the envelope of the pulse, known as the Carrier-Envelope Offset (CEO) [103]. The terms of interest here are those which act to disperse the frequency components away from each other, and these are second order and above. This dispersion results in a temporal stretching of the pulse (assuming the pulse was transform limited to begin with). The larger the bandwidth of the pulse, the more important the higher order terms will become. For the bandwidth in this thesis, terms above fourth order do not provide significant improvement in modelling the phase function, so can be dropped. If the second order term and above from Equation (4.6) are substituted into Equation (4.5), we arrive at

$$\mathbf{E_1}(\omega) = \mathbf{E}(\omega)\mathrm{e}^{-\mathrm{i}\frac{\phi^{(2)}(\omega)}{2}(\omega-\omega_0)^2}\mathrm{e}^{-\mathrm{i}\frac{\phi^{(3)}(\omega)}{6}(\omega-\omega_0)^3}\mathrm{e}^{-\mathrm{i}\frac{\phi^{(4)}(\omega)}{24}(\omega-\omega_0)^4}. \tag{4.7}$$

Equation (4.7) allows the addition of second, third and fourth order phases originating from either material dispersion, pulse stretchers or pulse compressors. There are other things which can add phase to the system, but these are the most relevant contributions to this experiment. The material phase function [104] is defined as

$$\phi(\omega) = \frac{\omega \mathrm{n}(\omega)\mathrm{L_m}}{\mathrm{c}}, \tag{4.8}$$

where ω is the frequency of interest, $\mathrm{n}(\omega)$ is the refractive index, $\mathrm{L_m}$ is the length of dispersive material, and c is the speed of light. The refractive index differentials that result from differentiating Equation (4.8) with respect to ω, can be derived from the Sellmeier Equation ([105]), which approximates empirically how the refractive index varies with wavelength in an optically transparent medium

$$\mathrm{n}(\lambda)^2 - 1 = \sum_{\mathrm{i}} \frac{\mathrm{B_i}\lambda^2}{\lambda^2 - \mathrm{C_i}}, \tag{4.9}$$

where λ is the wavelength, $\mathrm{B_i}$ and $\mathrm{C_i}$ are empirically derived constants [106]. The second, third and fourth order phase contributions from material dispersion are defined by substituting

Equation (4.9) into the second, third and fourth order differentials of Equation (4.8)

$$\phi^{(2)}(\omega) = \frac{\delta^2\phi}{\delta\omega^2} = \frac{\lambda^3 \mathrm{L_m}}{2\pi \mathrm{c}^2}\frac{\delta^2 \mathrm{n}(\lambda)}{\delta\lambda^2}, \tag{4.10}$$

$$\phi^{(3)}(\omega) = \frac{\delta^3\phi}{\delta\omega^3} = -\frac{\lambda^4 \mathrm{L_m}}{2\pi^2 \mathrm{c}^3}\left(3\frac{\delta^2 \mathrm{n}(\lambda)}{\delta\lambda^2} + \lambda\frac{\delta^3 \mathrm{n}(\lambda)}{\delta\lambda^3}\right), \tag{4.11}$$

and

$$\phi^{(4)}(\omega) = \frac{\delta^4\phi}{\delta\omega^4} = \frac{\lambda^5 \mathrm{L_m}}{2\pi^3 \mathrm{c}^4}\left(12\frac{\delta^2 \mathrm{n}(\lambda)}{\delta\lambda^2} + 8\lambda\frac{\delta^3 \mathrm{n}(\lambda)}{\delta\lambda^3} + \lambda^2\frac{\delta^4 \mathrm{n}(\lambda)}{\delta\lambda^4}\right). \tag{4.12}$$

As Equation (4.9) is defined in terms of λ, the refractive index differentials have been substituted from their $\mathrm{n}(\omega)$ form to the $\mathrm{n}(\lambda)$ form. Similarly to the above, second, third and fourth order phase equations for a grating compressor can be calculated from the grating compressor phase function [104]

$$\phi_{\mathrm{grating}}(\omega) = \frac{\omega}{\mathrm{c}}\frac{\mathrm{D}}{\cos\gamma}\cos(\gamma - \alpha(\omega)), \tag{4.13}$$

where D is the grating separation, γ is the angle of incidence to the first grating and $\alpha(\omega)$ is the angle of diffraction from the first grating, as illustrated in Figure 4.1. Here, the assumption has been made that the grating compressor is made of two gratings, in a double pass configuration, both subtending the same angles between the incident and diffracted beams. For a single grating compressor, D is the distance to the image of the grating created using lenses and a roof retro-reflector. In this case, D becomes twice the distance from the grating to the roof retro-reflector. Taking derivatives up to the fourth order

$$\phi^{(2)}_{\mathrm{grating}}(\omega) = \frac{\delta^2\phi}{\delta\omega^2} = -\frac{\lambda^3 \mathrm{D}}{2\pi \mathrm{c}^2 \mathrm{d}^2 \cos^3\alpha}, \tag{4.14}$$

$$\phi^{(3)}_{\mathrm{grating}}(\omega) = \frac{\delta^3\phi}{\delta\omega^3} = \frac{3\lambda^4 \mathrm{D}}{4\pi^2 \mathrm{c}^3 \mathrm{d}^2 \cos^3\alpha}\left(1 + \frac{\lambda \sin\alpha}{\mathrm{d}\cos^2\alpha}\right), \tag{4.15}$$

and

$$\phi^{(4)}_{\mathrm{grating}}(\omega) = \frac{\delta^4\phi}{\delta\omega^4} = -\frac{\lambda^5 \mathrm{D}}{8\pi^3 \mathrm{c}^4 \mathrm{d}^2 \cos^3\alpha}\left(1 + \frac{8\lambda \sin\alpha}{\mathrm{d}\cos^2\alpha} + \frac{\lambda^2}{\mathrm{d}^2}\left(1 + \tan^2\alpha\left(6 + 5\tan^2\alpha\right)\right)\right). \tag{4.16}$$

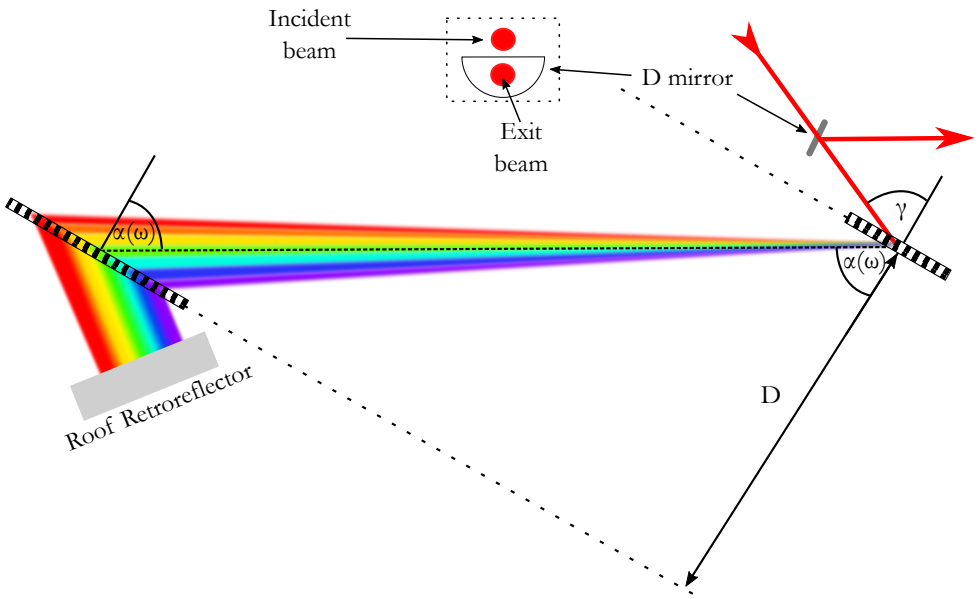


Figure 4.1: An illustration showing the quantities D, $\alpha(\omega)$ and γ.

To modify Equations (4.14-4.16) to be used for a stretcher, the parallel grating separation, D, should be made negative. Negative D can be achieved by using a 1:1 telescope added between the two gratings. Assuming a fixed $2f$ lens separation, the imaging system creates an image of the first grating, either in front of or behind the second grating, depending upon the distance of the telescope from the first grating. D then represents the distance between the plane of the gratings G_2 and G_2' [101]. However, changing D directly by moving grating G_2 would have affected α and γ of the two gratings, so the parameter D_c has introduced to decouple D from α and γ. This is illustrated in Figure 4.2.

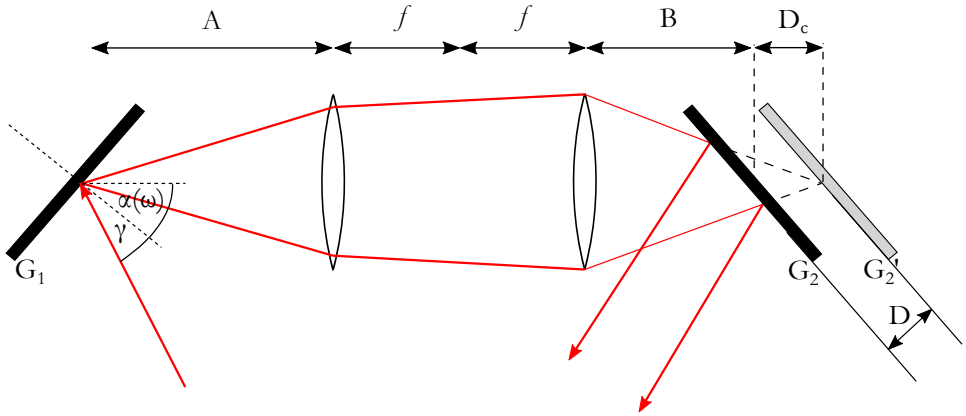


Figure 4.2: A 1:1 telescope placed between the gratings of a compressor can be used to create an image of grating G_1, behind grating G_2. The grating distance, D_c is defined as the distance between the centre of G_2 and the centre of the image of the first grating, G_2'. If G_2' appears before grating G_2 then D_c will be positive. If G_2' appears behind grating G_2 as shown above, then D_c will be negative. D_c needs to be converted to the parallel grating separation, D, using trigonometry before it can be used in Equations (4.14-4.16).

The phase equations for the two-grating compressor and stretcher are identical, except for the sign of D. This means that Equations (4.14-4.16) could be used to create a matched stretcher and compressor system, ideal for chirped pulse amplification (CPA). The importance of the phase acquired by a pulse is evident when the pulse is compressed. Given that the goal was to build a compressor, which could undo the dispersion added by the pulse stretcher and the materials in the system, building a matched stretcher and compressor would not provide optimum compression. If a compressor was unable to compensate for all the phase acquired by the pulse, then there will be some residual phase left over. This could have detrimental effects on the compressed pulse, preventing the pulse duration from reaching the transform limit that could be supported by the bandwidth and, therefore, reducing the achievable peak power.

4.2.2 Non-linear effects

In an amplifier, as a pulse liberates energy the pulse intensity is going to rise. This can lead to several non-linear effects, the largest contributors being non-linear phase (as defined in Equation (4.1), self-focusing, self-steepening and Raman scattering. The non-linear phase, self-focusing and self-steepening are related to the Kerr effect, which is the change in refractive index of a material in response to an applied electric field, as per Equation (4.2). Self-focusing is caused by high pulse intensities locally increasing the refractive index. Assuming a pulse with a Gaussian intensity spatial profile, there will be a refractive index gradient radiating out from the centre where the intensity is highest, creating a focusing effect for positive n_2. This focusing leads to an increase in the intensity at the centre of the pulse's spatial distribution. These high intensities contribute to further self-focusing and increase the non-linear phase acquired by a pulse. Self-steepening occurs when the Kerr effect causes the peak of the pulse to propagate slower than the less intense parts of the pulse. The result of the more intense regions propagating slower than the less intense regions, is that the peak shifts towards the back of the pulse, leading to a steepening of the trailing end.

Of these non-linear effects, the non-linear phase is the most dominant. Non-linear phase causes an impact due to the different phase values acquired by different frequency components in a non-flat spectrum. In a Gaussian spectrum, the centre will acquire the highest proportion

of non-linear phase with the wings of the pulse acquiring the least. Should each frequency component have the same amount of non-linear phase, then a pulse should undergo optimum compression. It is variations in the non-linear phase across the spectrum of the pulse, which leads to its detrimental effects. However, small amounts of non-linear phase of approximately one radian (between the peak and the edges of the spectrum), can be tolerated and are small enough for a system to achieve good compression. Beyond one radian, the pulse duration begins to broaden [107]. As previously mentioned, large amounts of non-linear phase can act as a barrier to optimum pulse compression which, if not compensated, can prevent a pulse being compressed back to the transform limit. However, there is an exception to this rule, where acquiring non-linear phase is desirable and can be used to shorten the pulse duration beyond its transform limit. It is known as non-linear compression and involves coupling optimally compressed pulses to the core of a fibre.

As a pulse propagates down the core of a fibre, it acquires non-linear phase as per Equation (4.1), which results in spectral broadening [108]. This spectral broadening leaves a residual chirp on the pulse, while leaving the temporal profile of the pulse unchanged. Following the spectral broadening stage, a compressor is used to remove this chirp and compress the pulse towards the transform limit of the broadened spectrum. For pulses with very high intensities, it is necessary to send them down hollow-core PCFs, filled with a noble gas, to avoid material damage [43, 44, 45]. Non-linear pulse compression is an active field of research with the aim of achieving high peak powers from fibre based laser systems. None of this research is found in this thesis, so it will not be discussed. For this thesis, non-linear phase is a quantity to be minimised, compensated for, or avoided where possible. To mitigate the intensity related non-linear effects, it is necessary to reduce the intensity of the pulse. Of course, achieving this by reducing the pulse energy is counter to the aims of laser wakefield acceleration. Two possible solutions are: to increase the area of the beam, and to stretch the pulse out in time. Both serve to reduce the pulse intensity, without affecting the pulse energy. Using a pulse stretcher, a spectral phase difference between the frequency components can be created as per Equation (4.6) and Equations (4.14-4.16).

Using Equations (4.14-4.16) and (4.10-4.12), the calculated phases due to the stretcher and material dispersion can be applied to the electric field in Equation (4.7). Once the chirped pulse has been amplified at this lower intensity, the phase added from the stretcher and material dispersion can be undone by a compressor. Chirped pulse amplification (CPA) is used to either mitigate non-linear effects for a given output peak power or to increase the compressed peak power obtainable for a given amount of non-linearity. For the energy levels needed for laser wakefield acceleration, it is expected that, even with CPA, the pulse will acquire considerable non-linear phase if it propagates down several amplifiers in series. To reduce the onset of non-linear effects, the intensity needs to be reduced further. Additional stretching is possible, but this is limited by the size of the components, the space to put them and the damage threshold of the compressor gratings. A practical limit will be reached where the cost of production outweighs the benefit. A better solution is to spread the pulse out spatially, and this is done by making the beam size and amplifiers larger. Terawatt class solid-state lasers use this technique along with CPA to reach high pulse energies, but this is not without drawbacks, as will be covered in Section 4.3. Similar to the cost and manufacturing limitations of the stretcher, the size of the optics and amplifiers are limited by these constraints also. As covered in Chapter 3, to further reduce the intensity, one pulse may be split into many which are then sent through their own separate amplifiers before being recombined back into one. Section 4.3 covers the mathematics of modelling coherent combination.

4.3 Modelling coherent combination

4.3.1 Model of combination

The coherent combination of two pulses can be mathematically modelled to allow an estimate to be made of the combined pulse's degree of linear polarisation and the fraction of available beam power [95]. An electric field in the temporal domain can be extracted from the spectral electric field

$$\mathbf{E(t)} = \frac{1}{\sqrt{2\pi}} \int_0^\infty \mathbf{E}(\omega) \mathrm{e}^{-\mathrm{i}\omega \mathrm{t}} \mathrm{d}\omega. \tag{4.17}$$

For two pulses which have acquired spectral phase from stretching and material dispersion, Equation (4.17) can be written as

$$\mathbf{E_1}(\mathbf{t}) = \frac{1}{\sqrt{2\pi}} \int_0^\infty \mathbf{E_1}(\omega) \mathrm{e}^{\mathrm{i}\phi(\omega)} \mathrm{e}^{-\mathrm{i}\omega \mathrm{t}} \mathrm{d}\omega, \tag{4.18}$$

and

$$\mathbf{E_2}(\mathbf{t}) = \frac{1}{\sqrt{2\pi}} \int_0^\infty \mathbf{E_2}(\omega) \mathrm{e}^{\mathrm{i}\phi(\omega)} \mathrm{e}^{-\mathrm{i}\omega \mathrm{t}} \mathrm{d}\omega, \tag{4.19}$$

where $\mathbf{E_1} \cdot \mathbf{E_2} = 0$. A phase delay between two pulses can be defined as [95]

$$\Delta\phi_{\mathrm{D}}(\omega) = \frac{2\pi\Delta \mathrm{L}}{\lambda} = \frac{\Delta \mathrm{L}}{\mathrm{c}}\omega, \tag{4.20}$$

where ΔL is the difference in optical path length travelled by two pulses. The phase delay is added as an additional phase to one of the spectral electric fields, to give

$$\mathbf{E_1}(\mathbf{t}) = \frac{1}{\sqrt{2\pi}} \int_0^\infty \mathbf{E}(\omega) \mathrm{e}^{\mathrm{i}\phi(\omega)} \mathrm{e}^{-\mathrm{i}\omega \mathrm{t}} \mathrm{d}\omega, \tag{4.21}$$

and

$$\mathbf{E_2}(\mathbf{t}) = \frac{1}{\sqrt{2\pi}} \int_0^\infty \mathbf{E}(\omega) \mathrm{e}^{\mathrm{i}\phi(\omega)} \mathrm{e}^{-\mathrm{i}\Delta\phi_{\mathrm{D}}(\omega)} \mathrm{e}^{-\mathrm{i}\omega \mathrm{t}} \mathrm{d}\omega. \tag{4.22}$$

The resulting electric field from the pulses' combination is found by summing

$$\mathbf{E_t}(\mathbf{t}) = \left(\mathbf{E_1}(\mathbf{t}) + \mathbf{E_2}(\mathbf{t})\right), \tag{4.23}$$

$$|\mathbf{E_t}(\mathbf{t})| = \frac{1}{\sqrt{2}} \left(|\mathbf{E_1}(\mathbf{t})| + |\mathbf{E_2}(\mathbf{t})|\right), \tag{4.24}$$

where the $\frac{1}{\sqrt{2}}$ factor is used because the pulses have perpendicular polarisation states and combine at 45°. The intensity is the time averaged Poynting vector. The Poynting vector is defined as

$$\mathbf{S} = \frac{1}{2} \langle \mathbf{E} \times \mathbf{H} \rangle, \tag{4.25}$$

where $\mathbf{E}$ and $\mathbf{H}$ are the electric and magnetic fields. For a plane wave in a linear and homogeneous medium, the amplitude of $\mathbf{E}$ and $\mathbf{H}$ are equal, therefore, when Equation (4.25)

is time averaged and a substitution is made for the magnetic field, the intensity can be approximated as

$$\mathrm{I} = \frac{\mathrm{n}\epsilon_0\mathrm{c}}{2}|\mathbf{E}|^2, \tag{4.26}$$

where the $\mathrm{n}\epsilon_0\mathrm{c}$ factor arises from the substitution of **H** for **E**. From Equation (4.26), n is the refractive index of the medium, ϵ_0 is the permittivity of free space and c is the speed of light. This thesis will assume the definition of intensity as shown in Equation (4.26). The fluence is defined as the total energy striking an area, which can be found from integrating Equation (4.26)

$$\mathrm{F} = \frac{\mathrm{n}\epsilon_0\mathrm{c}}{2}\int_0^\infty |\mathrm{E_t}(\mathrm{t})|^2\,\mathrm{dt}. \tag{4.27}$$

Using Parseval's Theorem, this integral is equivalent to

$$\mathrm{F} = \frac{\mathrm{n}\epsilon_0\mathrm{c}}{2}\int_0^\infty |\mathrm{E_t}(\omega)|^2\,\mathrm{d}\omega = \frac{\mathrm{n}\epsilon_0\mathrm{c}}{2}\int_0^\infty \frac{1}{2}\,|\mathrm{E_0}(\omega)|^2\left|\mathrm{e}^{+\mathrm{i}\phi(\omega)} + \mathrm{e}^{\mathrm{i}(\phi(\omega)+\Delta\phi_\mathrm{D}(\omega))}\right|^2\mathrm{d}\omega. \tag{4.28}$$

where $\mathrm{E_0}(\omega)$ is the common spectral profile of the two pulses. Expanding the mod squared bracket and simplifying, we find that

$$\mathrm{F} = \frac{\mathrm{n}\epsilon_0\mathrm{c}}{4}\int_0^\infty |\mathrm{E_0}(\omega)|^2\left(1+\cos\left(\Delta\phi_\mathrm{D}(\omega)\right)\right)\mathrm{d}\omega. \tag{4.29}$$

Equation (4.29) will give the fluence for a given phase delay between the pulses [95]. By scanning over the phase delay, the function will give a signal, which will vary from a maximum to a minimum. The degree of linear polarisation of this signal gives a metric of how well the pulses have combined, 100% representing perfect combination and 0% representing no combination. The Degree of Linear Polarisation (DOLP) is defined as

$$\mathrm{DOLP} = 100 * \frac{\mathrm{P_{max}} - \mathrm{P_{min}}}{\mathrm{P_{max}} + \mathrm{P_{min}}} = 100 * \frac{\mathrm{F_{max}} - \mathrm{F_{min}}}{\mathrm{F_{max}} + \mathrm{F_{min}}}, \tag{4.30}$$

where $\mathrm{P_{max}}$ ($\mathrm{F_{max}}$) and $\mathrm{P_{min}}$ ($\mathrm{F_{min}}$) are the maximum and minimum powers (fluences) of the pulses combined polarisation state, as measured through a HWP and PBS setup by rotating the HWP through 180° to find the maximum and minimum. For pulses with identical spectra and spectral phases, the metric is expected to give 100%. The metric will not produce 100%,

when the spectral phase acquired by each pulse is different. The acquisition of additional phase for one pulse will change Equation (4.28) into

$$\mathrm{F} = \frac{\mathrm{n}\epsilon_0 \mathrm{c}}{2} \int |\mathrm{E_t}(\omega)|^2 \, \mathrm{d}\omega = \frac{\mathrm{n}\epsilon_0 \mathrm{c}}{4} \int |\mathrm{E_0}(\omega)|^2 \left| \mathrm{e}^{+\mathrm{i}\phi_1(\omega)} + \mathrm{e}^{\mathrm{i}(\phi_2(\omega)+\Delta\phi_\mathrm{D}(\omega))} \right|^2 \mathrm{d}\omega, \tag{4.31}$$

where the phases have subscripts to indicate their difference. Equation (4.31) simplifies to

$$\mathrm{F} = \frac{\mathrm{n}\epsilon_0 \mathrm{c}}{4} \int |\mathrm{E_0}(\omega)|^2 \left(1 + \cos\left(\Delta\phi_\mathrm{D}(\omega) - \Delta\phi_{12}(\omega)\right)\right) \mathrm{d}\omega. \tag{4.32}$$

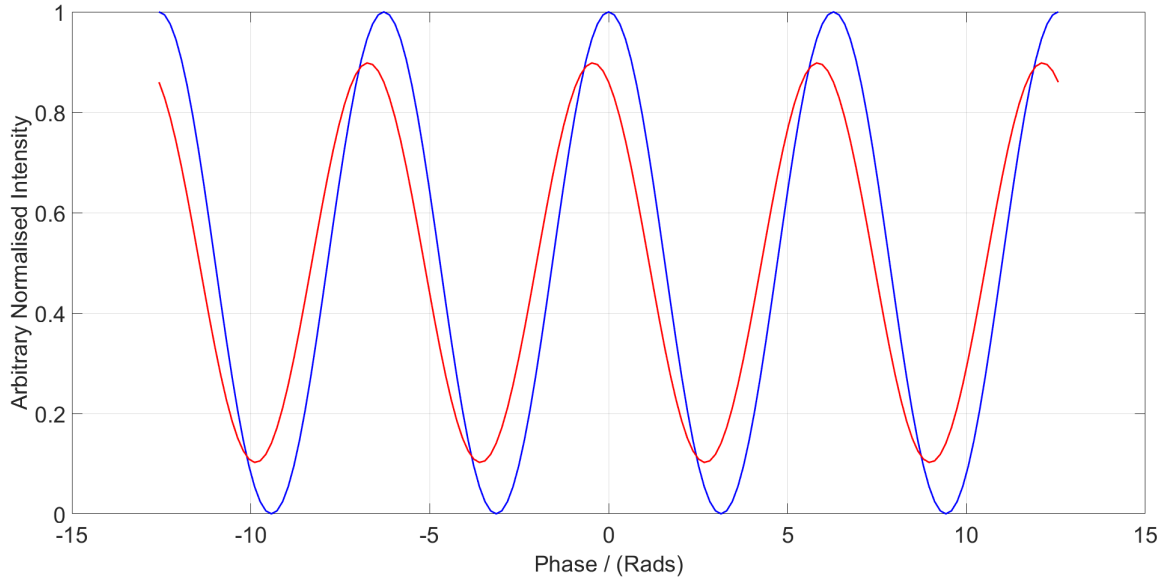


Figure 4.3: The modelled normalised intensity of the combined pulse as would be seen when transmitted through a PBS and observed on a photodiode when the phase delay of the two combining pulses is scanned. Blue line - The combined beam's normalised intensity observed when both pulses propagate through 1 m of fused silica. Red line - The combined beam's normalised intensity observed when one pulse propagates through 1 m of fused silica, while the other propagates through 3 m of fused silica. The extra phase acquired by one pulse has caused some combining frequency components to produce elliptical polarisation states. These states prevent the transmitted intensity through the PBS from dropping to zero. The degree of linear polarisation for the red line, using Equation (4.30), is 80%. Also, in comparison to the blue line, the maximum power obtainable has reduced also.

4.3.2 Non-linear phase compensation

Scanning the phase delay $\Delta\phi_\mathrm{D}(\omega)$ from -2π to 2π, the normalised intensity can be simulated, as seen in Figure 4.3, and the estimated visibility can be calculated. What is seen in Figure 4.3 is the intensity of the combined pulse observed behind a polariser at 45° to two interfering pulses of orthogonal polarisation. The 45° combined polarisation state is a result of the

constructive and destructive interference of components in the diagonal basis. The phase difference between the pulses will determine, whether they combine to form another linear polarisation state, or whether they combine to produce an elliptical polarisation state. The varying normalised intensity in Figure 4.3 is the result of the polarisation state changing from linear to elliptical, to circular, to elliptical (with the major axis of the ellipse orthogonal to that before) and then linear, but orthogonal to the initial state. Integrating Equation (4.32) for different values of ΔL, the spectral phase difference between the two pulses will lead to the minimum value of the normalised intensity being greater than zero, as shown in Figure 4.3. To better visualise the mechanism which causes the degree of linear polarisation to degrade, Figure 4.4 compares the temporal position of the frequency components between the pulses. Assuming two chirped, linearly polarised pulses, the different amount of material dispersion seen by each pulse results in the frequency components having different phase relationships to the central frequency component.

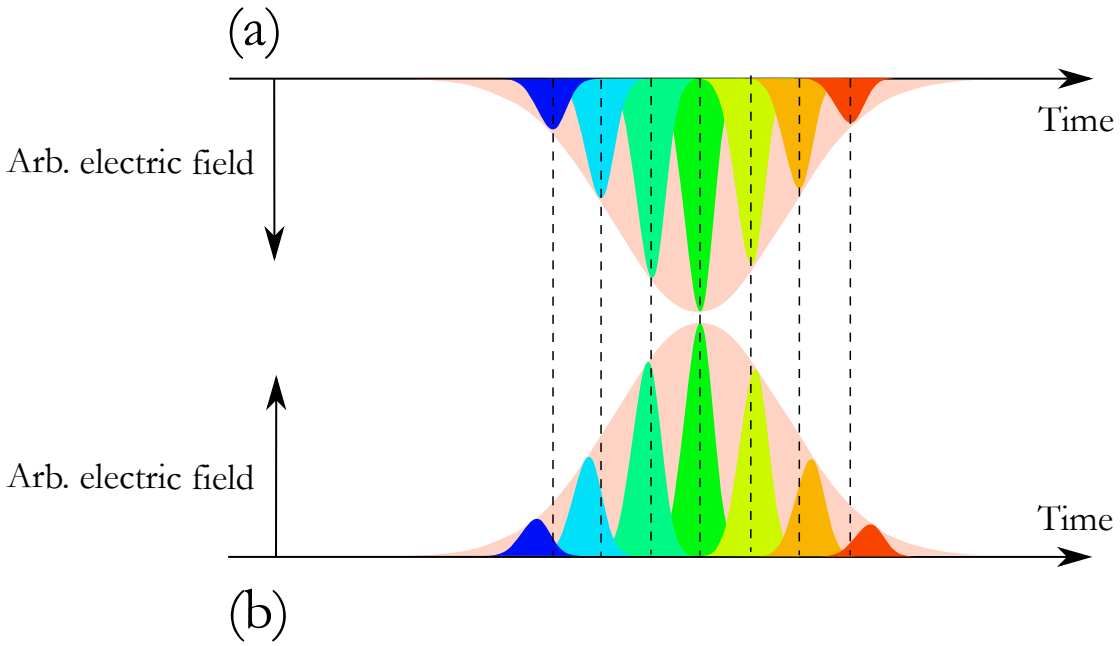


Figure 4.4: Two chirped pulses, which have propagated through different amounts of material. This illustrates and exaggerates how the difference in material dispersion causes pulse (a) and (b)'s frequency components to have a different phase relationship with the central frequency component. For a phase matched central frequency component, the leading and trailing frequency components will see poor overlap between the two pulses due to their arrival at different times at the point of combination.

In Figure 4.4, from overlapping and interfering the frequencies at the leading edge of the pulse (red components), it can be seen that the frequency components further back will overlap less. However, in the context of combining orthogonal linearly polarised pulses originating

from near-identical PCFs, the phase difference accrued from different amounts of material dispersion between the two pulses is small. This means that the offset between complimentary frequency components in each pulse will result in a combined pulse polarisation state, which varies temporally along the pulse as shown in Figure 4.5.

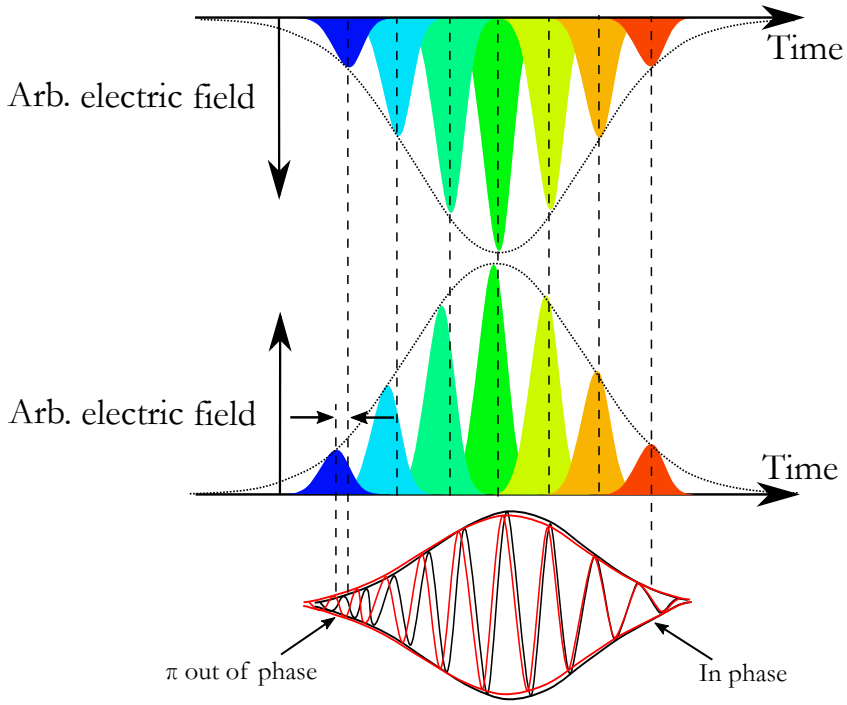


Figure 4.5: A phase difference between the frequency components of two chirped, orthogonal and linearly polarised pulses can result in the combination of those particular components producing elliptical or orthogonal linear polarisation states. This is illustrated above with the leading edge of the pulse (red) combining to produce linear polarisation, but the trailing components combining with incrementally larger phase differences. At the back of the pulse, the phase components are out of phase by π, producing a linear polarisation state orthogonal to that at the leading edge of the pulse. In between the orthogonal linear polarisation states are varying degrees of elliptical and circular polarisation. This model is still over-simplified as, in reality, the phase difference between the two pulses is a much more complicated function of ω.

When this combined pulse, with its time-varying polarisation state, is adjusted for maximum transmission through a polarising beam splitter (PBS), there will be a small amount of rejection from the PBS's dark port. This rejection comes from the elliptical, circular and orthogonal linear polarisation states of the trailing frequency components which have non-zero components orthogonal to the transmission axis of the PBS. This rejected beam power becomes a problem when the optical components after the point of combination are polarisation sensitive. For coherent combination experiments combining just two fibres, the grating compressor is the most likely place for this loss. However, this is not true if more than two fibres are undergoing coherent combination in a filled aperture setup (fibres are combined in

pairs in a cascade-like fashion). Here, greater amounts of beam power are lost through the multiple stages of combination.

As mentioned earlier, the non-linear phase is a quantity that needs to be minimised, avoided or compensated for. For a single pulse, the non-linear phase can be compensated for by adjusting the quantity of third order dispersion added by the compressor [109]. However, for two pulses, any differences in dispersion and non-linear phase will prevent the combined pulse from being compressed back to the transform limit. The level to which the combination would be affected can be visualised by performing a parameter scan of the difference in non-linear phase and the difference in length of the dispersive element seen between the two pulses. From the literature [95], a Figure Of Merit (FOM) identical to that of the degree of linear polarisation of two combined pulses, can be formulated from Equation (4.32). The FOM is valid only for strongly stretched pulses, by setting $\Delta\phi_{\mathrm{D}}(\omega) = 0$ and substituting the phase function $\Delta\phi_{12}(\omega)$ for

$$\Delta\phi_{12}(\omega) = \frac{1}{2}\beta_2 \Delta \mathrm{L}(\omega - \omega_0)^2 + \Delta\phi_{\mathrm{NL}}\mathrm{S}(\omega), \tag{4.33}$$

where ΔL is the difference in length of the dispersive element between the two amplifiers which, with a second order material dispersion coefficient of β_2 and $\Delta\phi_{\mathrm{NL}}$, is the difference in non-linear phase between the two pulses, and $\mathrm{S}(\omega)$ is the dimensionless normalised spectral intensity

$$\mathrm{S}(\omega) = \mathrm{e}^{-4\ln 2\left(\frac{\omega-\omega_0}{\omega_{1/2}}\right)} \tag{4.34}$$

Assuming that the total phase difference is small ($\Delta\phi_{12}(\omega) < \pi/4$), a Taylor expansion of the cosine function in Equation (4.32) can be performed, which after some mathematical manipulation, leads to [95]

$$\mathrm{FOM} = 1 + \left(\frac{1}{\sqrt{2}} - \frac{1}{2} - \frac{1}{2\sqrt{2}}\right)\Delta\phi_{\mathrm{NL}}^2 - \frac{3}{512\log 2^2}\omega_{1/2}^4\beta_2^2\Delta\mathrm{L}^2 \tag{4.35}$$

$$+\frac{1}{16\log 2}\left(1 - \frac{1}{2\sqrt{2}}\right)\omega_{1/2}^2\beta_2\Delta\mathrm{L}\Delta\phi_{\mathrm{NL}},$$

where β_2 is the second order material dispersion, ΔL is the difference in length of the dispersive element between the two amplifiers, $\Delta\phi_{NL}$ is the difference in non-linear phase between the two pulses and $\omega_{1/2}$ is the FWHM bandwidth of the ultrafast laser.

From Equation (4.35), the purely dispersive term has an $\omega_{1/2}^4$ dependence, meaning that the FOM is more sensitive to dispersion as the bandwidth of the laser pulse increases. The ω^2 dependence in the cross term means that the FOM is more sensitive to $\Delta\phi_{NL}$ and ΔL at larger bandwidths. This equation is valid to within 2% of the analytical answer for small phase differences between the pulses of up to $\pi/4$ [95]. Using Equation (4.35), two parameters scans were performed for spectral FWHM bandwidths of $\lambda_{1/2} \sim 7.5$ nm and $\lambda_{1/2} \sim 5$ nm supporting 200 fs and 315 fs pulses respectively. At both bandwidths, ΔL and $\Delta\Phi_{NL}$ were then scanned over a range of $\Delta L \in \pm 1$ m and $\Delta\phi_{NL} \in \pm 1\, rad$, with β_2 being set for fused silica. The results of these scans can be seen in Figure 4.6.

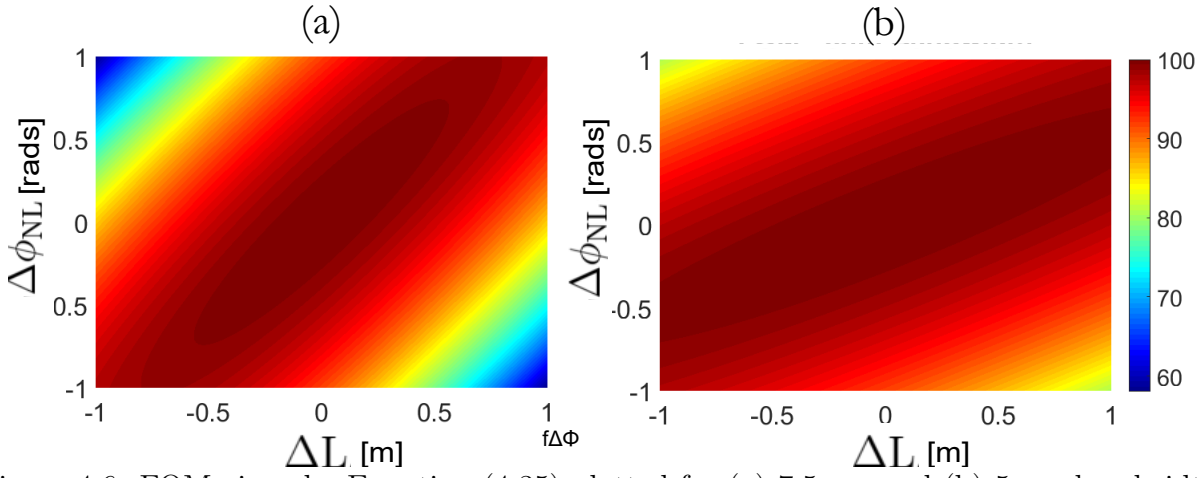


Figure 4.6: FOM given by Equation (4.35) plotted for (a) 7.5 nm and (b) 5 nm bandwidths with $\Delta\phi_{\mathrm{NL}} \in \pm 1$ rads and $\Delta\mathrm{L} \in \pm 1$ m. The colourbar shows the percentage FOM from 60 to 100 %.

Despite a difference between the non-linear phase and length of the dispersive element between the two pulses, certain combinations of non-linear phase difference and length difference of the dispersive element can result in the FOM reaching above 95%. Similarly, certain combinations of non-linear phase difference and length difference of the dispersive element can result in the FOM reducing to the low 60% range. Both bandwidths can achieve a FOM of above 99%, although the FOM of larger bandwidths will be more susceptible to mismatches in ΔL and $\Delta\phi_{NL}$. Even with very different lengths of the dispersive element, the two pulses

can be combined to give a FOM above 99%, by carefully choosing a value for $\Delta\phi_{NL}$ based upon the length difference of the two fibre amplifiers.

4.4 Coherent combination of N lasers

To try to predict the performance of a large facility scale laser system using a simple model, it was assumed that: the type of combination used would be a filled aperture setup; any two fibres with power P from the set of N would combine to give a combined beam power of 2P; and the two beams would combine with a fraction of available beam power of η. For one stage of combination between two fibre lasers

$$\mathrm{P_u} = 2\mathrm{P}\eta, \tag{4.36}$$

where $\mathrm{P_u}$ is the useful beam power available after combination of two fibres. For the combination of many fibres, there will be $\mathrm{N_c}$ combination stages. Equation (4.36) then becomes

$$\mathrm{P_u} = 2^{\mathrm{N_c}}\mathrm{P}\eta^{\mathrm{N_c}} = \mathrm{P}\,(2\eta)^{\mathrm{N_c}}\,. \tag{4.37}$$

From Equation (4.37), as the fraction of available beam power is unlikely to be 100%, the number of combining operations that can be performed to achieve a net gain in pulse energy is limited. Performing three stages of combination (using eight fibres) with a fraction of available beam power of 95% gives a final efficiency of

$$\eta_{\mathrm{t}} = \eta^3 = 0.95^3 \approx 0.86, \tag{4.38}$$

where an assumed 5% loss of power from each stage has resulted in a 14% loss of total beam power. Equation 4.37 will be used again in Chapter 10 to analyse the efficiency of other systems. The fraction of available beam power measured in this thesis and presented in Chapter 8, will be used to predict the future performance of a larger system of four and eight PCFs, which can be directly compared to similar experiments found in the literature.

4.5 Summary

In this chapter, the theory describing an ultrafast laser pulse was introduced and extended to account for the coherent combination of two ultrafast laser pulses. Using this model, the coherent combination of laser pulses could be simulated to give an estimate of the degree of linear polarisation, given that the total non-linear phase acquired by a pulse is less that one radian. Reasons that could prevent the coherent combination of two pules from producing a pure linearly polarised state were discussed. A simple model was used to show the effect of different frequency components arriving at the point of combination at different times. The effect on the figure of merit and degree of linear polarisation, due to the difference in non-linear phase, $\Delta\phi_{\mathrm{NL}}$, and length of the dispersive element, ΔL, between the pulses was discussed also. A small mismatch of ΔL and $\Delta\phi_{\mathrm{NL}}$ between the two pulses was found not to be detrimental to the combination, and control over both of these quantities can be exploited to produce high figure of merits. Finally, the coherent combination of N lasers was discussed, and a simple model produced to predict the useful beam power that could be extracted.

Chapter 5

Experimental Setup

5.1 Overview of the experiment

The experimental setup used in this thesis was a filled aperture, Hänsch-Couillard polarisation locked, coherent combination experiment. With the exception of the seed laser, the full experimental setup was designed and constructed by myself. The full experimental setup is shown in Figure 5.1, which provides a unifying view of the whole experiment. On the back of the page is Figure 5.1a, a larger fold out diagram of the same Figure but with optical components annotated. This should be folded out when reviewing the results chapter for quick identification of the places where power measurements were taken. Two doped photonic crystal fibres (PCF) with different lengths, core diameters and internal structures were placed in separate arms of a Mach-Zehnder interferometer and their outputs combined at the combination polarising beam splitter (PBS). A linearly polarised seed pulse was oriented at 45° to the transmission axis of a PBS, making up the splitting element of the interferometer, and causing the pulse to spilt into two pulses of orthogonal S- and P-polarisations. Each pulse was sent around a different arm of the interferometer, seeding the PCFs, which acted as the power amplifiers in this experiment. The 70 μm and 85 μm PCFs were pumped separately using a 100 W and 500 W diode laser respectively. The amplified linearly-polarised pulses output from the PCFs were combined at another PBS. Half Wave Plates (HWP) were used prior to the combining PBS to rotate the polarisation of the pulses to allow for reflection of both pulses at the PBS.

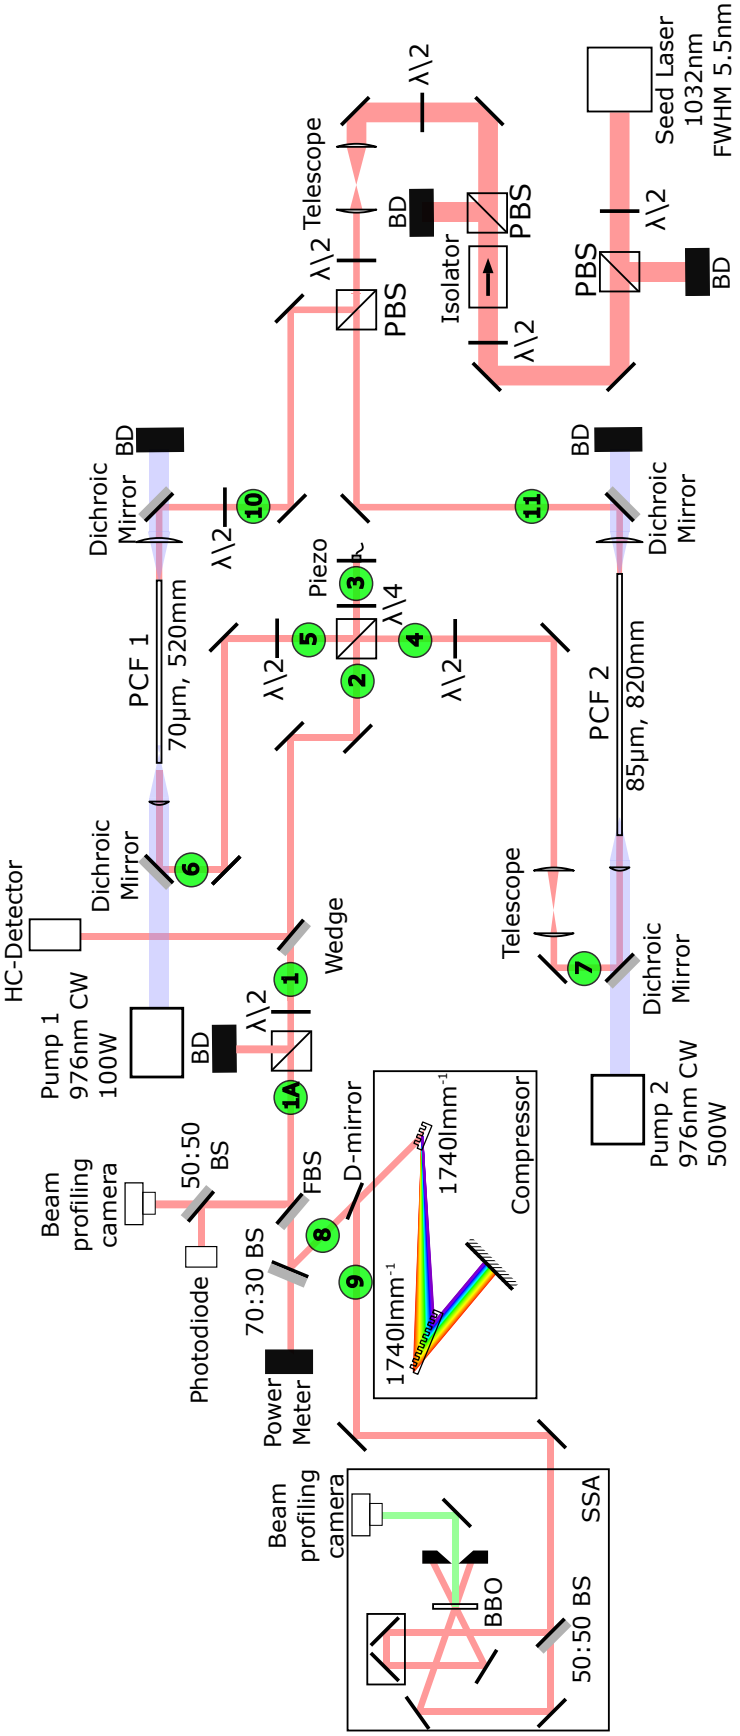


Figure 5.1

The full experimental setup. Red beams are the seed and amplified seed, blue beams are pumps, BD is a beam dump, PBS is a polarising beam splitter, FBS is a Fresnel wedge and BS is an unpolarised beam splitter. Points on the experimental setup where beam power is measured are highlighted with green numbers. **1**: System output power. **1A**: Polarised beam power. **2**: Combination cube output beam power. **3**: 70 µm photonic crystal fibre power prior to combination. **4**: 85 µm photonic crystal fibre power prior to combination. **5**: 70 µm photonic crystal fibre beam power prior to the combination cube. **6**: 70 µm photonic crystal fibre output power. **7**: 85 µm photonic crystal fibre output power. **8**: Compressor input. **9**: Compressor output. **10**: 70 µm photonic crystal fibre seed power. **11**: 85 µm photonic crystal fibre seed power

The pulse emanating from the longer 85 μm PCF was reflected at the PBS and exited the interferometer. The pulse emanating from the shorter 70 μm PCF was reflected from the PBS also to a translation-stage-mounted mirror to add the additional path length needed to balance the interferometer. A Quarter Wave Plate (QWP) was present in the extra path length to rotate the polarisation by 90° for transmission through the PBS upon its second pass. After retro-reflection, the pulse from the 70 μm PCF was transmitted through the PBS and also exited the interferometer. Only on the second pass through the PBS did the beams of the two PCFs get combined. The combined beam was stabilised using Hänsch-Couillard polarising locking, with a Fresnel wedge taking a 1% reflection and directing it into a Hänsch-Couillard detector. The stabilised combined beam had its power reduced by a 70:30 beam splitter; 30% was sent to a compressor and single shot autocorrelator for analysis and the other half was captured by a power meter. This was to avoid damaging the compressor gratings and the single shot autocorrelators's frequency doubling crystal. The seed laser used in this experiment was a commercial system and many of its characteristics were not well known initially. Approximate values were known for its pulse duration, pulse energy and spectrum, but no model existed to describe the pulse which emanated from the seed laser's shutter. Therefore, the seed laser was treated as a black box system in this Chapter, for the purpose of describing the coherent combination experiment. Detailed characterisation of the seed laser is conducted in Chapter 6, and this explains why the pulse propagation model from Chapter 4 needed to be extended.

5.2 Beam transport

From the output of the AS laser, there were several optical components used to manipulate the properties of the beam to make it suitable for splitting and coupling to the fibres. The seed laser produced a train of ultrafast pulses stretched from 210 fs to 290 ps, at a repetition rate of 6.49 MHz and with a pulse energy of 1.6 μJ, creating an average power of $\sim$ 11.4 W. From Figure 5.2, the beam exited the seed laser and was sent through an intensity modulator made from a HWP and a PBS. The intensity modulator allowed the output power of the seed laser to be adjusted without having to adjust the pump lasers, which would require seed laser

to re-thermalise. The next group of components protected the seed laser from any parasitic lasing from the PCF amplifiers. The HWP rotated the seed beam's polarisation state to match the transmission axis of the optical isolator, while the additional PBS was used to clean up the polarisation state of the beam exiting the optical isolator. The final HWP in the beamline was used to control the polarisation of the beam incident upon the splitting PBS. This allowed adjustment of the beam power, to give variable splitting between the two arms.

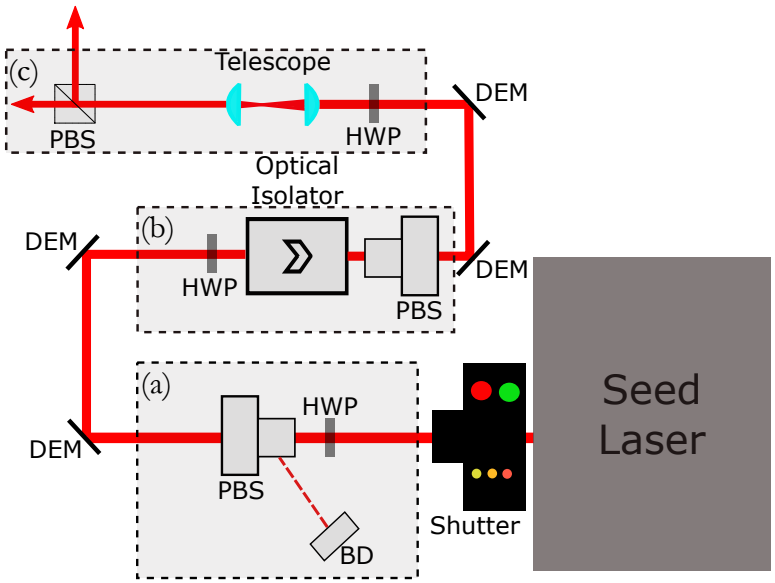


Figure 5.2: The beam transport after the output shutter of the seed laser to the first PBS of the interferometer. The beam was made to propagate through three small systems to allow the beam's properties to be controlled. System (a) was an intensity modulator, used to control the output power of the seed laser, without adjusting the pump beam power driving the seed laser's amplifiers. System (b) was used to protect the seed laser from any parasitic lasing, which would travel back along the beam line from the PCF amplifiers. The HWP matched the seed beam's polarisation to the optical isolator. The additional PBS ensured that the beams polarisation state was pure. System (c) adjusted the beam's size and polarisation state. A HWP changed the power split between the two arms of the interferometer, while the telescope reduced the beam size and collimated the beam. All mirrors shown were dielectric mirrors, and BD was a beam dump.

5.3 The interferometer

Once the beam had been conditioned in the beam transport section of the experiment, the beam was incident upon the splitting element of the interferometer. Two orthogonally polarised beams were created, which propagated along separate arms of the interferometer until met again at the point of combination. For ease of set up and alignment, a Mach-Zehnder interferometer was chosen to embed the amplifiers. Figure 5.3 shows a schematic layout of the two PCF amplifiers embedded in a Mach-Zehnder interferometer.

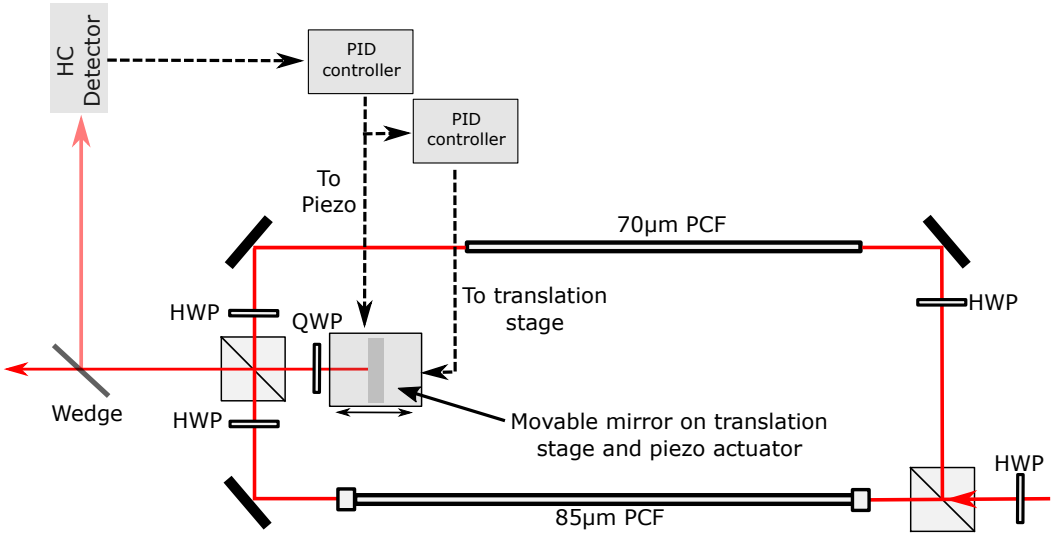


Figure 5.3: A schematic diagram of the Mach-Zehnder interferometer used in the experimental setup. This is a duplicate of figure 3.11 for convenience to the reader. Lenses and extra beam path has been omitted from the diagram to prevent it from becoming cluttered. The full diagram is shown in Figure 5.1

A polarising beam splitter (PBS) with a 10 mm aperture was the first splitting element; the beam was split into two replicas and directed towards the two photonic crystal fibres (PCF). The splitting PBS was preceded by a HWP in the beam transport optics. The beam's linear polarisation state could be rotated with this HWP, to change the splitting ratio between both arms. This allowed the beam power to be split 50:50 or unevenly as required This assumed, however, that there was no leakage. The PBSs used in the interferometer had a 1000:1 extinction ratio for the P-polarised beam but only a 100:1 extinction ratio for the S-polarised beam. Hence, it was expected that there would be a small fraction of leakage in the reflected beam. PBS cubes were chosen over plate PBS because the S-Polarised beam could be reflected at 90° to the direction of propagation. Compare this to plate beam splitters where the S-polarised beam would be reflected near Brewster's angle, which would make the construction of coherent combination setups more challenging.

The 70 µm PCF contained stress rods, which applied stress to the core, making it birefringent. Consequently, only a certain orientation of linear polarisation was allowed to propagate down the core, everything else was rejected into the inner cladding. A HWP preceded the 70 µm PCF to correctly match the beam's linear polarisation state to that accepted by the core. After the PCFs, the linearly polarised beams were incident onto a PBS. For combination, both beams should have orthogonal linear polarisation states, with the direction dependent upon whether the beam was to be transmitted or reflected at the PBS. However, the extra

path length needed to match the interferometer was present at the combining optic, so both beams had their polarisation states set to be S-polarised for reflection at the combining PBS. The beam that emanated from the shorter PCF was made to propagate along the extra path length, undergo a retro-reflection and make a double pass through a QWP. This double pass caused a 90° rotation of the beam's linear polarisation state to P-Polarsaition, which allowed the beam to be transmitted through the PBS on its second pass. It was at this point that the two beams were recombined and stabilised.

5.4 Locking and stabilisation

To combine the ultrafast pulses from the PCFs, the interferometer needed to have arms with exactly matched path lengths to ensure the pulses arrived at the combining element at the same time. In the case of the mismatched PCFs, the arms lengths would not be the same, given the extra 300 mm of Yb-doped silica in the 85 μm PCF. Therefore, the arm containing the 70 μm PCF needed to have extra path length added in order to compensate for this difference. This extra path length was achieved by adding a mirror, mounted on a translation stage set normal to the incident beam. It was positioned at the combining PBS (point 3 on Figure 5.1), so that the first pass through the splitter directed the beam towards the mirror. After propagating along the extra path length, the second pass through the PBS was when the combining occurred. The extra path length could have been added at either the splitting or combining part of the interferometer. The combining PBS was chosen to prevent decoupling the seed beam from the core of the PCFs. At the splitting PBS, accurate mirror alignment would be required to prevent the seed laser from decoupling from the core of the PCF, as the translation stage scanned over its range of travel in an attempt to reacquire the lock. Decoupling the seed laser from the PCF while it was being pumped, particularly at high pump powers, raised the possibility of parasitic lasing, which could have caused damage to components, including the seed laser. Hence, to mitigate the risk of decoupling the seed laser, the retro-reflecting mirror was added to the combination element.

A $\frac{1}{2}$" mirror was used for the retro-reflection, mounted on a piezo-electric ceramic. It was in turn mounted on a small stage driven by a stepper motor, which itself was mounted

onto a larger 25 mm translation stage. The larger translation stage was used for gross path length scanning to find the overlap point. The smaller translation stage and the piezo form part of two PID loops, used to phase lock the pulses. The piezo was used for the sub-λ pulse stabilisation, while the smaller stage was used to keep the piezo near to the centre of its range of travel.

The layout of the Hänsch-Couillard polarisation lock is shown in Figure 5.4. A front surface reflection was taken using a Fresnel wedge and was directed towards a Hänsch-Couillard detector [100]. The Hänsch-Couillard detector analysed the combined beam and provided an error signal to a PID lock box, as discussed in Section 3.4.1. The PID lockbox drove the piezo to change the path length of one of the arms of the interferometer. The change in path length leads to a change in the way the beams combined, which resulted in a new error signal. A closed loop was required to lock the two pulses together. A further PID loop existed within the PID lockbox, where the output signal that drove the piezo was used as an error signal to drive the small translation stage. The translation stage was used to perform gross tuning of the path length of the interferometer, aiming to keep the piezo in the centre of its range, where the stabilisation was the most robust. The combination of the piezo and the translation stage allowed the path lengths of the two interferometer arms to be locked and stabilised over a greater range of path length difference.

5.5 Pulse compression

The seed laser was built originally to produce a compressed seed pulse of $\sim$500 fs. The compressor that was built into the seed laser was removed subsequently to allow further amplification of the stretched pulse. However, the addition of further material dispersion from the two PCFs and all the transmissive optical components, meant that the previous compressor could no longer give optimum compression. The original seed-laser-compressor had gratings, which were fixed to a custom board making it impractical to modify. Therefore, it was necessary for me to design and build a new compressor to make the compressor suitable for further amplification.

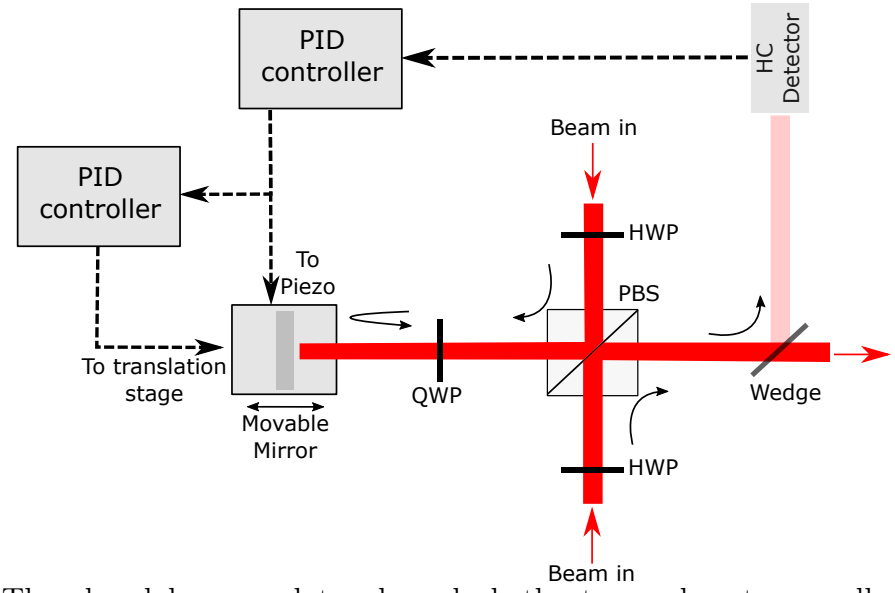


Figure 5.4: The closed loop used to phase lock the two pulses temporally. PBS was a polarising beam splitter, HWP was a half wave plate and QWP was a quarter wave plate. The output of the first PID loop was fed into another PID controller, which drove a translation stage. The translation stage was used to keep the piezo in the centre of its range of travel, where the lock was more robust. Additionally, the large range of the translation stage allowed the system to increase the range of path lengths over which the system could remain locked and stable. The QWP is used to rotate the polarisation of the reflected beam by 90° after a double pass.

5.5.1 Compressor model

To design a compressor, which would compress a pulse optimally, the phase picked up by a pulse from stretching, material dispersion and non-linear effects, needed to be calculated as accurately as possible. As the manufacturer had refused to divulge key details about the stretcher, the phase it contributed to the pulse had to be estimated. This estimation was performed by approximating the incidence angle to the grating (64°, operating at the Littrow angle) and then measuring the distance (measured to be ∼0.480 m) from grating G_2, to the image of the grating G_2' (see Figure 4.2). The distance was measured by using the leakage of the transmitted beam through grating G_2 to find the point at which the angular chirp had been removed at point G_2'. A model was created, which simulated: a pulse being stretched from 210 fs to 290 ps using the simulated stretcher; the material dispersion and non-linear phase picked up from the pre-amp and main-amp; and the effects of spectral narrowing. The model and the characterisation of the seed laser are presented in more detail in Chapter 6. The pulse was split and each pulse propagated through a separate PCF, acquiring further material dispersion and non-linear phase from amplification. Each PCF was modelled with

a different core diameter to scale the non-linear effects correctly according to the intensity and length of each PCF. The pulses were then recombined and the resulting combined pulse was compressed.

5.5.2 Compressor design

It was intended that the compressor for the coherent combination experiment would be different from the fixed compressor design used by the seed laser. Both gratings were mounted on rotation mounts to allow the angles of incidence and diffraction of the beam to be tuned, while the second grating of the compressor was further mounted on a motorised translation stage to allow remote adjustment of the grating separation. The ability to adjust and tune the angle and grating separation allowed the user of the compressor to adjust the amount of second, third and fourth order phase removed from a chirped pulse, as defined in Equations (4.14 - 4.16). The additional benefit of being able to adjust the grating parameters quickly was that non-linear phase could be compensated for by introducing third-order phase [109]. As the amplification in the coherent combination experiment increased, any non-linear phase acquired could be compensated for by tuning the compressor's grating separation and the angle of both gratings.

The seed-laser-compressor made use of a pair of 1750 l/mm blazed (for 1030 nm) transmission gratings. Due to a desire to keep the seed-laser-compressor intact, it was decided that additional gratings should be purchased. 1740 l/mm blazed (for 1030 nm) reflective diffraction gratings were considered initially as candidates for the replacement compressor, but calculations showed that the angle of incidence and angle of diffraction required to remove all of the phase from the pulse were closely matched (near Littrow angle). Similar incident and diffraction angles are not ideal for a compressor design using reflective diffraction gratings, as the position of the second grating would be in the path of the incident beam. Hence, it became obvious why the designers of the seed-laser-compressor had chosen to use transmission gratings. Therefore, the pair of 1750 l/mm blazed transmission gratings from the seed laser were extracted and remounted for use. The new compressor design is depicted in Figure 5.5.

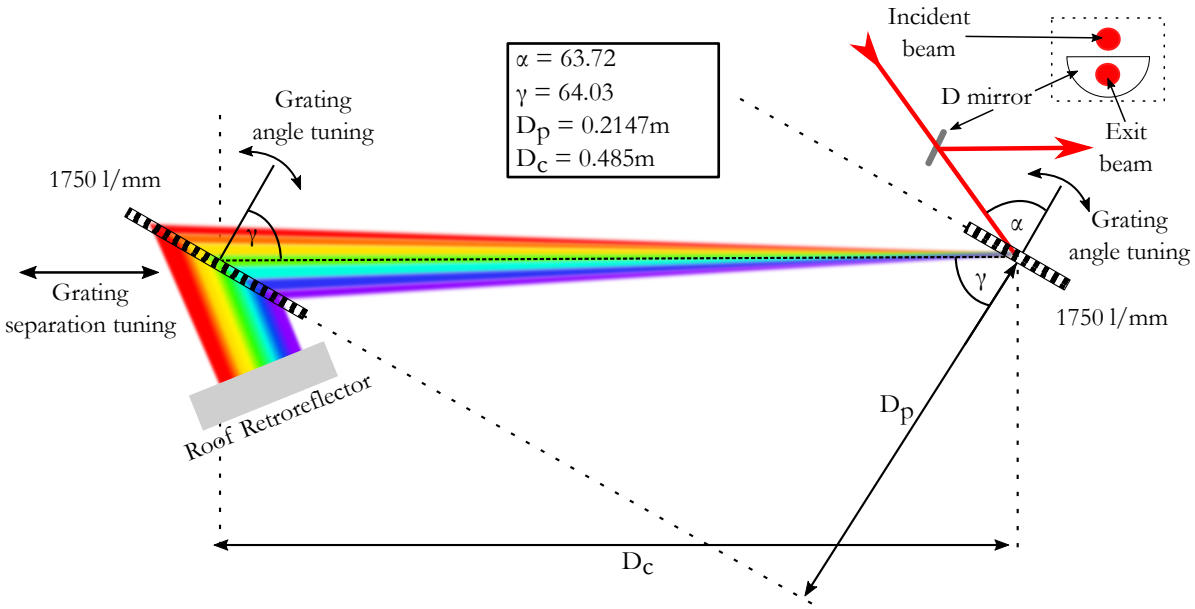


Figure 5.5: A top down view of the compressor layout. α is the incident angle and γ is the diffracted angle. D_c is the grating to grating distance measured between their centres and D_p is the parallel grating distance suitable for use in Equations (4.14 - 4.16). Changing D_p directly would have affected α and γ of the two gratings, so the parameter D_c was introduced to decouple D_p from α and γ.

5.5.3 Compressor testing

The compressor was built on the optical bench and aligned to optimise the calculated parameters. The compressor layout can be seen in Figure 5.6. The first tests of this compressor produced autocorrelations which showed that the pulse was sub-optimally compressed. Therefore, the grating distance and grating angle were adjusted to find the optimum pulse compression. This was a slow process as the metric used to determine the impact of an adjustment was a pulse's autocorrelation, measured using a scanning intensity autocorrelator. A small adjustment would be made to the compressor, and the autocorrelator would be scanned, taking a few seconds to complete the scan. If many adjustments were required to optimise the compressor, this could potentially take hours. Therefore I decided to provide some automation to speed up the process.

The scanning autocorrelator and second compressor grating were mounted on motorised stages to facilitate quick changes of position. The motor controllers had a USB interface, which could be controlled with ActiveX controls. MATLAB contained software libraries that enabled easy interfacing with the ActiveX controls, allowing me to write a MATLAB script to enact computer control over both of these translation stages. The MATLAB script was

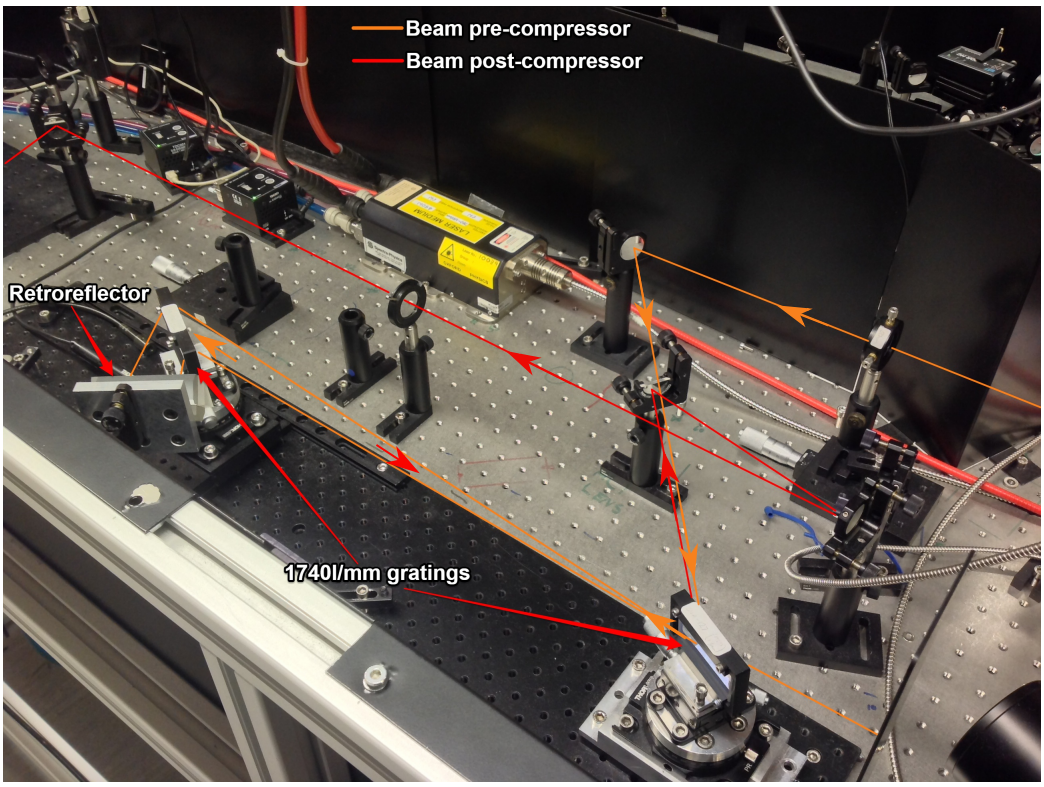


Figure 5.6: The layout of the compressor with the beam line annotated.

used also to control an oscilloscope, via Ethernet, and then transmit the recorded data back to the PC. The translation stages, oscilloscope and the MATLAB script all formed part of a closed loop. The MATLAB script scanned the compressor's grating distance in small 0.1 mm steps. After every increment, an autocorrelation was taken, and the data analysed to extract the autocorrelation's FWHM. The autocorrelation's FWHM, along with the position in which it was recorded, was saved into temporary variables. After the next compressor increment and autocorrelator scan, the program compared the autocorrelation's FWHM to the previous value and, if any improvement on the autocorrelation's FWHM was observed, saved the new FWHM and position into the temporary variables. This process was repeated until the MATLAB script had iterated over the entire specified range of the second compressor grating. Once the end was reached, the best autocorrelation and position were recalled, and the stage moved directly back to the optimum position. Plotting all of the autocorrelation FWHM values against position produced a graph, as seen in Figure 5.7. This control loop could be left to scan the parameter space and find the optimal pulse compression, given the current grating angles. An example of the optimal autocorrelation measured from the grating distance scan is shown in Figure 5.8.

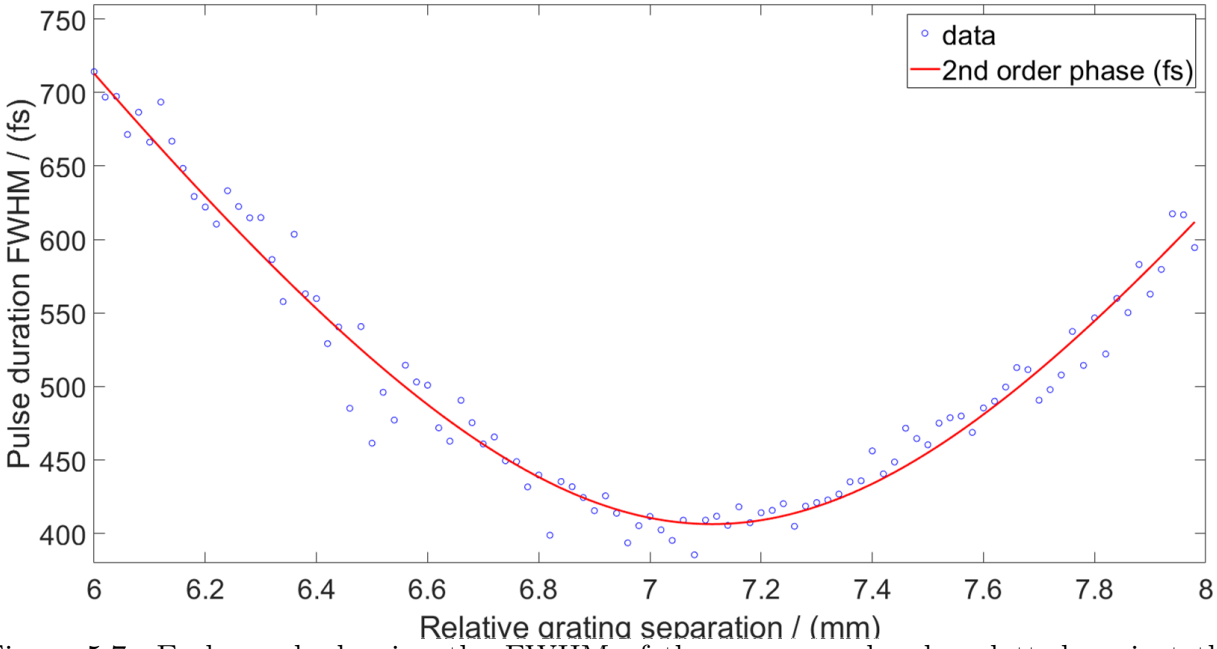


Figure 5.7: Early work showing the FWHM of the compressed pulse plotted against the second grating's relative position for a given grating angle of ∼63.5°. The measured minimal pulse durations are ∼10% longer than the ∼ 370 fs seen earlier in the chapter. This is because the optimisation work was done using a model that made assumptions about the phase added by the stretcher. Subsequently, this led to the optimum compressor parameters being different to those seen in Figure 5.5. The grating compressor second-order-phase Equation (4.14) (RMS uncertainty ∼ 25 fs) is shown fitted to the data.

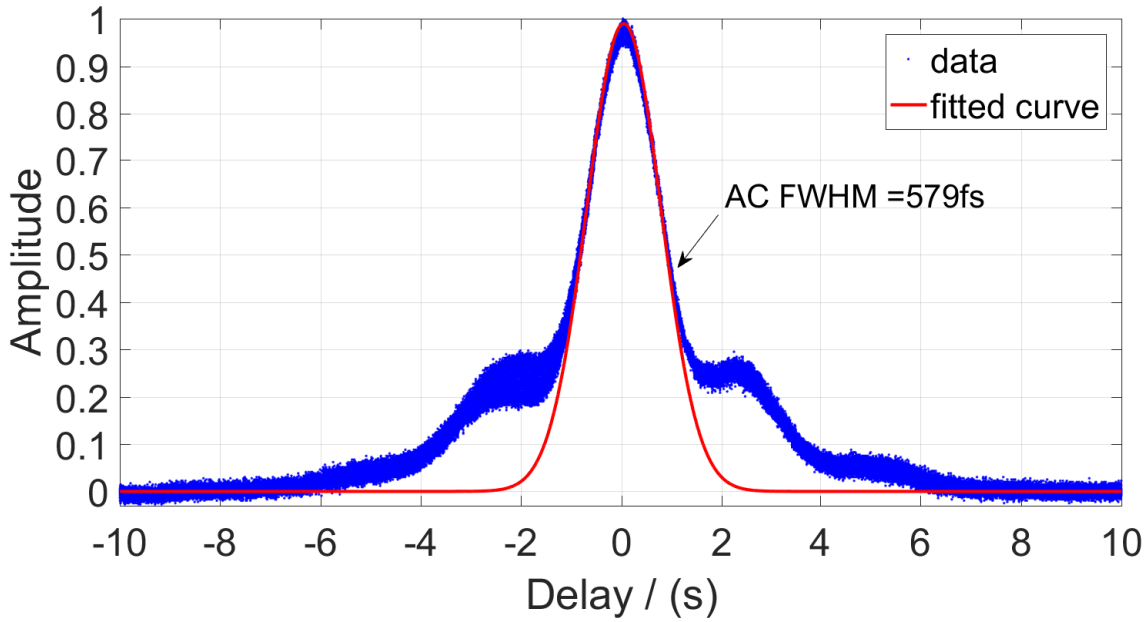


Figure 5.8: The optimal autocorrelation found from earlier work where the compressor grating distance was scanned over range of positions with autocorrelations recorded at each point. The autocorrelation FWHM was measured to be ∼ 580 fs, with the pulse duration deconvolution as a Gaussian to give ∼ 410 fs. Because of noise in the signal seen by the oscilloscope, there was some uncertainty in the underlying autocorrelation trace. Therefore, fitting to the autocorrelation could give slightly different results between measurements.

Small incident and diffracted angle variations could be trialled, and the system left again to scan and optimise the grating distance. The compressor used manual rotation stages to control the grating angle, which prevented me from conducting further development of the compressor auto-alignment work. However, the intention was to control the grating angle, so that the computer could optimise over the grating distance and grating angle concurrently. For this, two algorithms were considered. The first was a simple hill climbing algorithm [110], which would make a small change to the grating angle and grating distance in all directions, measure the autocorrelation for each, and calculate which direction led to the most substantial improvement on the compression. The compressor would be adjusted in the best direction before the grating angle and distance were probed again. This algorithm would be simple to implement, but it does have the drawback that it could get stuck in a local maximum. External influences or multiple layers of control would be needed to recognise this and move it to a different part of the parameter space. The second algorithm that could be used was simulated annealing [110, 111].

Simulated annealing would make a random change to the grating angle and distance and measure the FWHM of the resulting autocorrelation. The autocorrelation FWHM is the parameter that would be optimised. A temperature variable is used to weight a probability of new values being accepted, even if they are worse than the best found so far. Initially, at high temperatures, the algorithm will more than likely accept solutions, which produce longer autocorrelations than those that have been found previously, as well as those that improve the FWHM. As the algorithm progresses, the temperature variable gradually reduces. As this happens, the solutions that produced longer autocorrelations are rejected more often by the algorithm, in favour of solutions that reduce the autocorrelation FWHM. At high temperatures, the algorithm explores the global parameter space, and, at low temperatures, the algorithm enters into a hill climbing regime. Hence, the algorithm is less likely to get stuck in a local maximum. However, despite the speed of either algorithm, they were limited by the speed of the scanning autocorrelator, which could result in a single run taking hours. Consequently, I decided to construct a single shot autocorrelator to speed up these processes.

5.6 Single shot autocorrelator

The final part of the beam line was the single shot autocorrelator. Previously, experiments had made use of a scanning autocorrelator, which I had built, and could be used either as an interferometric or an intensity autocorrelator. In both of these methods, the autocorrelator was limited by the speed at which it could take an autocorrelation. The process of making a small change to a compressor parameter, and then waiting for the autocorrelator to scan, became prohibitively slow. In addition, there was some variation in the measured autocorrelation on a scan to scan basis as the stage and fitting routine gave slightly different results each time. An example of this variation can be seen in Figure 5.9, where 2000 ACs were taken and binned into a histogram. Here, the pulse duration had been optimised as best as could be achieved. Therefore, I decided it would be beneficial to build a single shot autocorrelator to allow more rapid compressor optimisation and data acquisition.

5.6.1 Single shot autocorrelator theory

A single shot autocorrelator differs from a scanning autocorrelator as it is able to take a full autocorrelation from a single shot. Similar to the intensity autocorrelator, a pulse is split into two which then interact as they propagate through a frequency doubling crystal. The delay for the single shot autocorrelator is mapped spatially onto the crystal resulting in a frequency doubled beam, whose spatial distribution correlates to the delay [112]. Note that the delay is not the same as the pulse duration. In order to recover the pulse duration from the autocorrelation, the autocorrelation should be deconvolved. Different pulse shapes have different deconvolution factors, whether they would be Gaussian, Sec^2 or super-Gaussian. For this thesis, a Gaussian pulse shape has been assumed since the oscillator autocorrelation, shown in Section 6.1.2, shows excellent agreement with a Gaussian fit. The choice of a Gaussian pulse shape dictates that the deconvolution factor is $\sqrt{2}$. Hence, any pulse durations can be recovered from the autocorrelation by dividing the delay by $\sqrt{2}$. The layout for a single shot autocorrelator is shown in Figure 5.10. The schematic diagram showing the crossed beam geometry and their interaction in the frequency doubling crystal is shown in Figure 5.11.

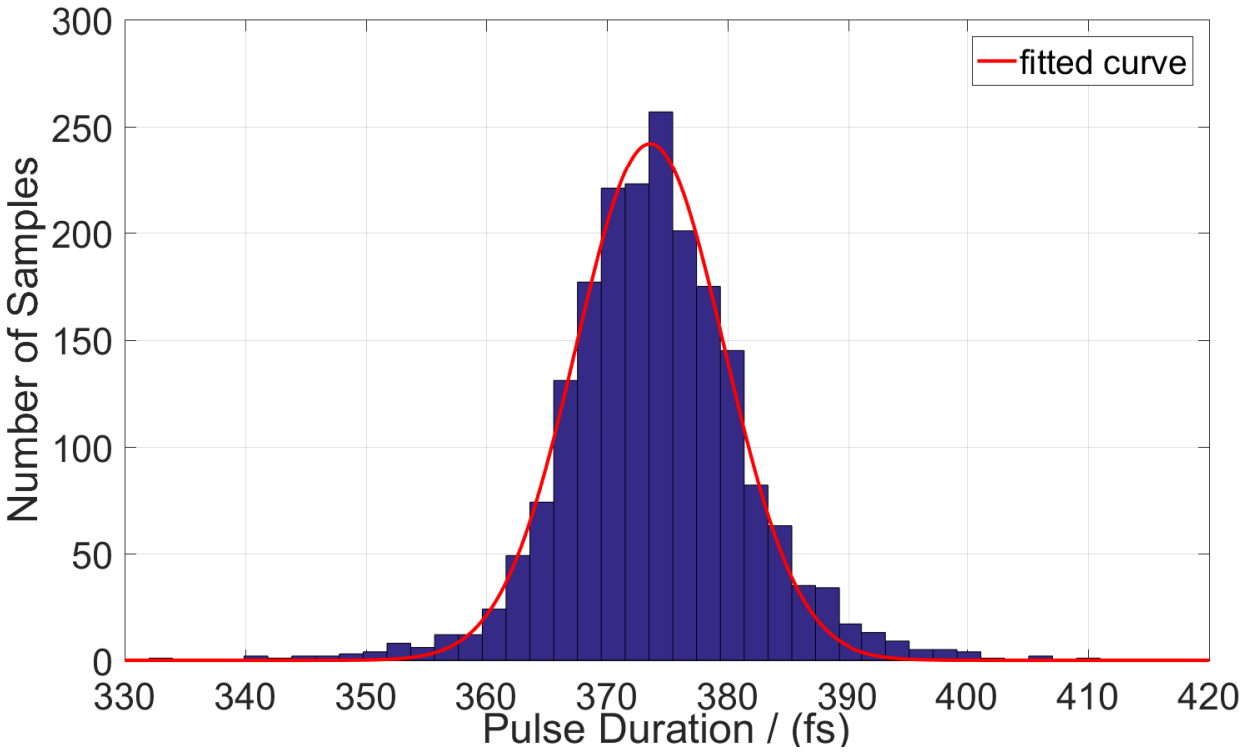


Figure 5.9: The variation in the autocorrelation (deconvolved to give the pulse duration) using the scanning autocorrelator over 2000 scans. The distribution is centred at 374 fs with sigma (FWHM) of 8.7fs (20 fs). Each autocorrelation scan took $\sim$ 4 seconds to measure and analyse.

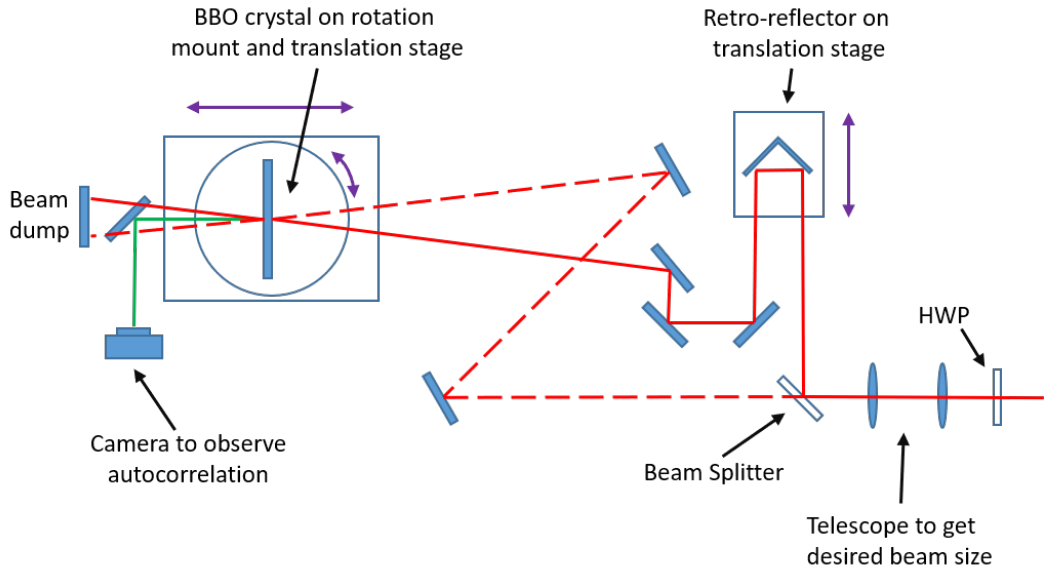


Figure 5.10: The general layout of the single shot autocorrelator.

The interaction of the two pulses leads to a third beam of frequency doubled light, whose intensity profile is dependent upon the interaction of the two crossing pulses, and whose direction bisects the angle between the two beams. Thicker crystals give the beams more

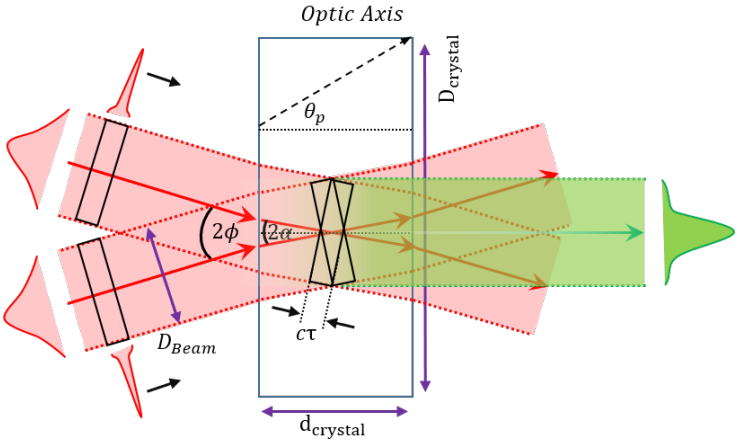


Figure 5.11: Two beams interacting inside the frequency doubling crystal. Beams enter from the left and exit to the right. Here ϕ is the angle of incidence of the beam to the normal of the crystal face, α is the reduced angle according to Snell's law, τ is the length of the temporal window and θ_p is the phase matching angle. The two pulses overlap spatially and temporally in the crystal producing frequency doubled light where they interact. The spatial profile of the frequency doubled beam is then used to infer the pulse duration of the input pulse.

distance over which to interact and can produce higher intensities of second harmonic light. The maximum crystal length is determined by the spectral acceptance of the crystal, which determines the passband of the laser bandwidth. The passband determines which frequency components could still undergo frequency doubling despite the phase matching only being optimised for the central frequency component of the laser spectrum [113]. As the spectral acceptance scales inversely with distance, crystals cannot be made arbitrarily thick. A crystal whose length is too long could result in clipping of the laser bandwidth, producing a broadening effect on the measured pulse durations. Therefore, to choose the correct crystal for the single shot autocorrelator, the spectral acceptance must first be considered. Equation (5.1) is used to calculate the passband supported by a crystal,

$$\mathrm{S} = \frac{\gamma}{\mathrm{d_{crystal}}}, \tag{5.1}$$

where S is the passband of the crystal, γ is the spectral acceptance, with units of $\mathrm{nm \cdot cm}$ and $\mathrm{d_{crystal}}$ is the length of the crystal. The passband defines the region where a laser with frequencies between the lower and upper cutoff frequencies, can undergo frequency doubling. Outside of this range, a laser would see no frequency doubling, causing the bandwidth of the frequency doubled light to be clipped. Typical values for γ are shown in Table 5.1, as taken from the CASTECH® website [114]. So, taking the spectral bandwidth of the seed laser used

Crystal	BBO	LBO (Type-1 SHG)	LBO (Type-2 SHG)	KTP
Spectral Acceptance (nm · cm)	1.1	1.0	1.3	0.56

Table 5.1: The spectral acceptance of a few common non-linear crystals, as taken from the CASTECH® website [114].

in the experiments in this thesis, discussed in Chapter 6, the 3% to 3% bandwidth of the pulse is measured to be ≈ 12 nm. The zero to zero bandwidth was taken from where the spectral intensity was measured to be 3% of the peak intensity. Using Equation (5.1) and the value for the spectral acceptance for BBO, the maximum crystal length that could be used was

$$\mathrm{d_{crystal}} = \frac{1.1}{12} \approx 0.092\,\mathrm{cm} = 92\,\mu\mathrm{m} \tag{5.2}$$

For TW class Ti:Sapphire systems with large bandwidths, typically from 40-60 nm, the necessary crystal thickness to prevent spectral clipping, using BBO, is approximately 50µm. This puts an upper limit on the crystal length and explains why crystals for large bandwidth systems are ∼ 100 µm thick or less. The diameter of the crystal (crystal aperture) is dependent upon the beam diameter, which itself is dependent upon the temporal window [112]

$$\mathrm{D_{beam}} = \frac{c\tau}{\tan(\phi)}, \tag{5.3}$$

where $\mathrm{D_{beam}}$ is the beam diameter, ϕ is the angle of incidence of the beam to the normal of the crystal face and τ is the temporal window, the maximum duration of time that can be resolved. The crystal aperture is defined as:

$$\mathrm{D_{crystal}} \geq \frac{\mathrm{D_{beam}}}{\cos(\phi)} + \mathrm{d_{crystal}}\tan(\alpha), \tag{5.4}$$

where $\mathrm{D_{crystal}}$ is the diameter of the crystal and all other variables are as previously defined. The phase matching angle changes depending upon the type of geometry used for the single shot autocorrelator and the desired value for α. The generalised equation for the phase matching angle in a uniaxial crystal is defined as [112]

$$\theta_\mathrm{p} = \arcsin\left(\frac{(\mathrm{n_0^\omega}\cos(\alpha))^{-2} - (\mathrm{n_0^{2\omega}})^{-2}}{(\mathrm{n_e^{2\omega}})^{-2} - (\mathrm{n_0^{2\omega}})^{-2}}\right), \tag{5.5}$$

where n_0^{ω} is the ordinary refractive index at frequency ω, $n_0^{2\omega}$ is the ordinary refractive index at frequency 2ω and $n_e^{2\omega}$ is the extraordinary refractive index at frequency 2ω.

The phase matching angle is the angle between the optical axis and the direction in which $n_o(\omega)$ and $n_e(2\omega)$ are equal, as demonstrated in Figure 5.12. The aim of this exercise is to make the propagation velocity of a beam at frequency ω, equal to that of a beam at frequency 2ω. The crossing angle α is introduced, so that the beams produce 2nd harmonic light on the axis, which can then be separated from the two input beams and imaged onto a CCD. The crossing angle of the beams inside the crystal is denoted as α. Snell's law will need to be applied to calculate the necessary beam crossing angle exterior to the crystal as shown in Figure 5.11.

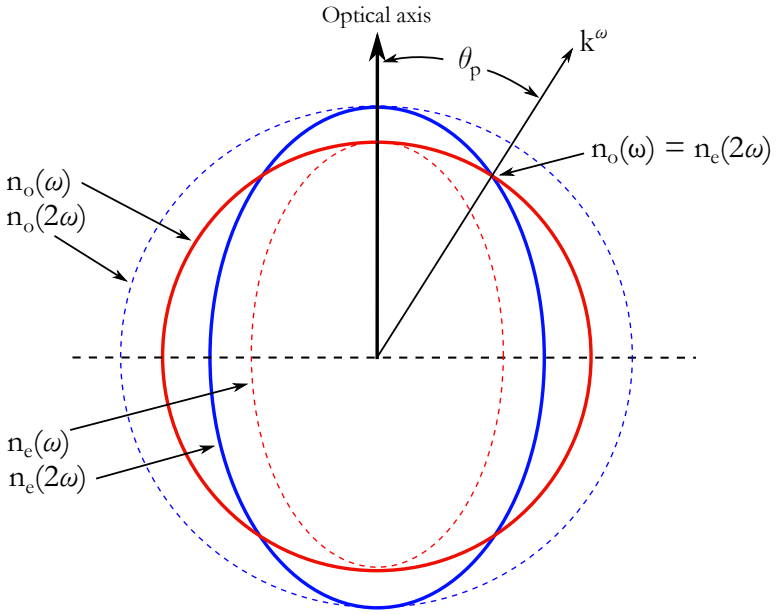


Figure 5.12: The index ellipsoid of a uniaxial crystal. A uniaxial crystal has a refractive index, which is dependent upon the direction of propagation through the crystal. For a beam to 'see' this changing refractive index, a beam's polarisation must be parallel to the plane in which the index variation occurs. A beam with a polarisation state perpendicular to this plane will see no refractive index variation. The circles in the diagram represent the refractive index seen by beams of frequency ω (red lines) and frequency 2ω (blue lines) polarised orthogonal to the plane (direction independent). The ellipses represent the refractive index seen by beams of frequency ω (red lines) and frequency 2ω (blue lines) polarised parallel to the plane (direction dependent). To perform frequency doubling, the beams of ω and 2ω must be phase matched, which means that they must travel at the same propagation velocity through the crystal. It is possible to achieve phase matching by propagating a beam of frequency ω at an angle of θ_p from the optical axis through the crystal, as defined by the wave vector k_{ω}.

τ (ps)	4 ps Window	5 ps Window
$c\tau$ (mm)	1.2	1.5
D_{beam} (mm)	4.5	5.6
D_{beam} ($\times$2.3) (mm)	10.3	12.9
$d_{crystal}$ (mm) (max)	0.92	0.92
$D_{crystal}$ (mm)	10.9	13.6
γ (nm $\cdot$ cm)	11.98	11.98

Table 5.2: The parameters required of a BBO crystal for a single shot autocorrelator, given the requirement for a temporal window of either 4 or 5 ps. ($\times$2.3) is the multiplication factor for D_{beam} required to ensure that 99.9 % of the beam's energy passes through the crystal aperture and avoids clipping.

5.6.2 Single shot autocorrelator setup

The compressed pulse duration for the laser system in this thesis was approximately $\sim$400 fs. To observe the pulse, while not optimally compressed, a temporal window of 5 ps was chosen, along with a beam half crossing angle of $\alpha = 8\,\text{deg}$. BBO was the crystal chosen as it had a similar damage threshold to LBO crystals, a larger spectral acceptance than an LBO crystal, and was a uniaxial crystal. This meant that the calculations for the phase matching angle were much simpler. Using Equations (5.1-5.5), Table 5.2 shows the calculated single-shot autocorrelator parameters. These parameters were calculated for both a 4 ps and a 5 ps window for comparison.

The single shot autocorrelator was constructed using the 5 ps temporal window and aligned to image the autocorrelator's frequency doubled beam onto a beam profiling camera. However, before a lineout could be taken of the autocorrelation, the delay per pixel had to be calibrated. This was done by scanning the translation stage of the single shot autocorrelator a known amount, and measuring how far the peak moved across the CCD, as seen in Figure 5.13. Once the delay per pixel has been determined, a lineout could be taken from the observed beam profile and the FWHM of the central peak, calculated in terms of pixels, as shown in Figure 5.14. The measured FWHM, in terms of pixels, was converted into delay (fs) and a deconvolution factor of $\sqrt{2}$ was used to convert from time delay to pulse duration.

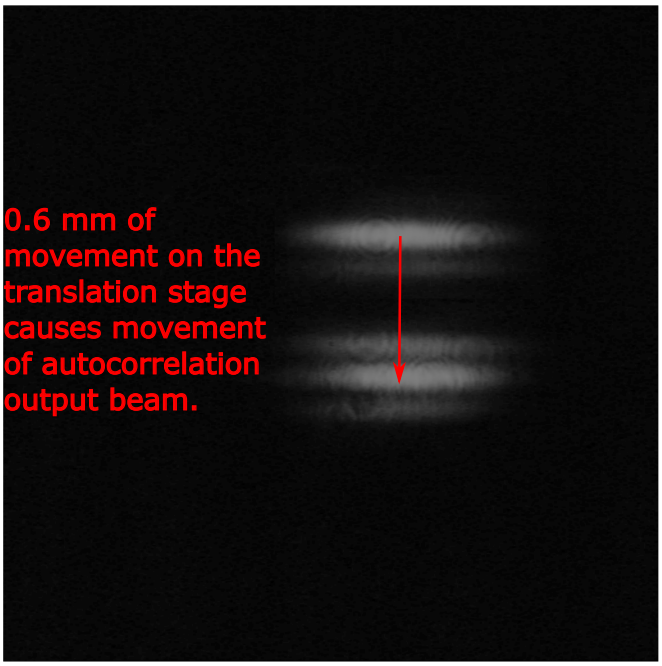


Figure 5.13: The autocorrelation on the camera from the single shot autocorrelator before (top beam) and after (bottom beam) the movement of the translation stage. The images of the two beams have been merged into one photograph for ease of demonstration.

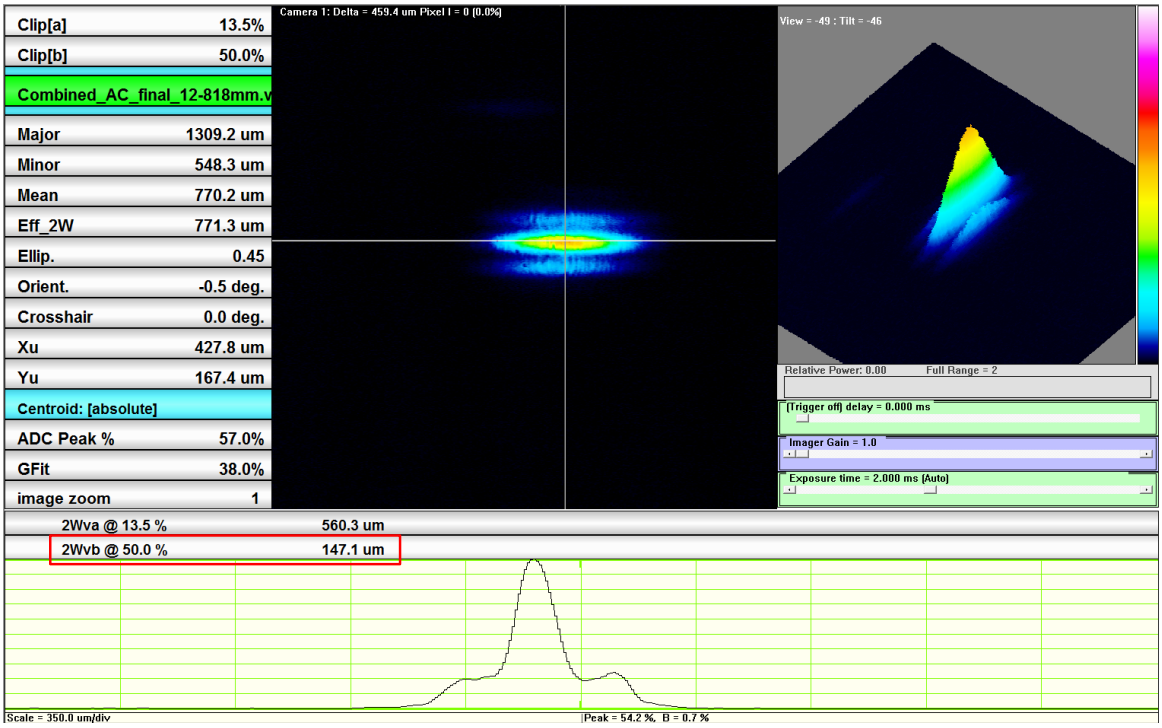


Figure 5.14: A line out is taken of the autocorrelation, and the FWHM is determined in terms of pixels. The DataRay software has already performed the conversion to microns, based upon its measured FWHM, boxed in red.

5.7 Beam combination

The crucial part of this experiment was the coherent combination of the beams from two fibres into a single beam [83, 85]. This meant that the pulses from each fibre would need to be overlapped spatially, temporally and spectrally, and to have the same frequency components present. For two pulses with identical spectra, flat phases and S- and P-polarisations, when overlapped and locked together (with zero phase difference), all of the frequency components will interfere to generate another linear polarisation state at 45° to both S- and P-polarisations. Placed out of phase by π, the linear polarisation state generated by the combined beam would be orthogonal to the previously generated state. When viewed through a HWP and polariser, the combined beam's polarisation state can be analysed. Rotating the HWP rotates the combined beam's polarisation state and should allow the transmitted beam power to go from a maximum value to complete extinction. For unlocked pulses, the combined polarisation state will vary depending on the phase difference between the two pulses. Observed through an analyser, the signal seen on a photodiode was that seen in Figure 5.15. Figure 5.15 assumes that the pulses were already in temporal and spatial overlap. Without this, there would be no varying polarisation state seen on the photodiode. Instead, only a steady voltage would be observed. The minimum power level seen in a free-running system could be due also to unequal power contributions from the combining beams. When both in, and out of phase, the combining beams would not produce polarisation states that were orthogonal to one another. The acute angle between the combined in, and out of phase polarisations states would prevent complete extinction of the combined beam when viewed through a PBS. This is not a problem for a combined and phase stabilised beam, as the polarisation state is analysed by rotating a HWP. The HWP is able to fully rotate the linear polarisation state to make it orthogonal and this forces the complete extinction of the beam. Any leakage seen now will be due to unpolarised components leaking through the PBS.

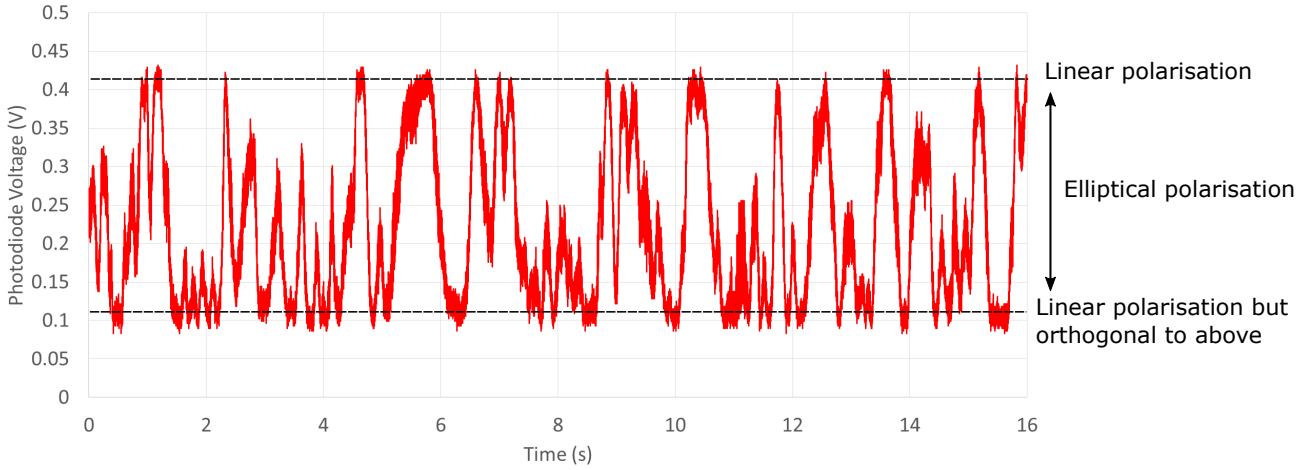


Figure 5.15: The signal seen on a photodiode behind a polariser, when the phase of the two beams is left to drift over time. The signal going from maximum to minimum represents the combined beam's polarisation state evolving to and from orthogonal linear polarisation states. The signal does not go to zero, as there was some leakage through the PBS, caused by the combined beam not being 100% linearly polarised. This non-perfect polarisation state was due to the wavefront differences between the beams, a subject that will be covered in Chapter 7.

5.8 Diagnostics

To gauge the quality of the combined beam, I used several diagnostics, in addition to the single shot autocorrelator. These involved a pickoff wedge and two beam splitters, to allow multiple diagnostics to be used concurrently. The layout is shown schematically in Figure 5.16. The polarisation analysis stage of the diagnostics showed how much of the beam remained in a linear polarisation state, and how much of the beam was lost into a rejected beam. The two most important figures of merit used in this thesis, to assess the quality of coherent combination, were the degree of linear polarisation and the fraction of available beam power in the linear polarisation state. Both of these were determined from the polarisation analysis. The degree of linear polarisation is important, when considering the power loss of a combined beam, when it encounters polarisation sensitive optics. The fraction of available beam power is important when assessing how much power is available for further stages of coherent combination.

For a two fibre system, the next polarisation sensitive component seen by the combined beam would be a grating compressor. The compressor throughput efficiency could reduce

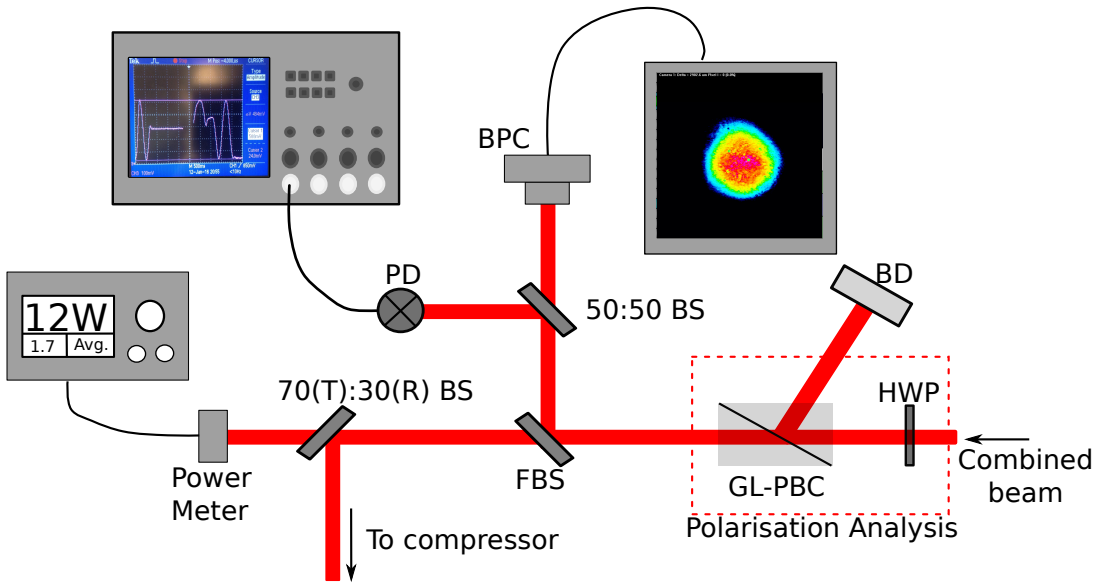


Figure 5.16: The diagnostic setup. FBS is a Fresnel beam spitter and BPC is a beam profiling camera, BD is a beam dump, BS is an unpolarised beam splitter, PD is a photodiode and GL-PBC is a Glan-Laser Polarising Beam Cube.

significantly if the beam's polarisation state was not linear and perpendicular to the grating lines. At best, the rejected beam could be transmitted or reflected into the zeroth order. At worst, the rejected power could be absorbed by the grating and result in material damage. Any power lost by a beam propagating through the compressor could not be used for an experiment further along the chain. Hence, it would have an impact on the wall plug efficiency of the whole system. A loss of wall plug efficiency could negate the advantages offered by fibre lasers and their high wall plug efficiencies over Ti:Sapphire systems. So, it would be in the best interests of the experiment to ensure that as much of the beam was as linearly polarised as possible. To measure the polarisation state of the combined locked beam, the beam was sent through the polarisation analysis stage. By rotating the HWP, the transmitted power for the polarisation state of the desired and rejected beams could be measured by a power meter. The power meter recorded the average power of the beam over 20 s. This is seen in Figure 5.16.

The average power of the combined beam prior to the polarisation analysis was used for determining the pulse energy. To measure the output power of the system, the same power meter was moved to Position 1 in Figure 5.1, immediately after the pickoff wedge, which directs a small amount of beam power to the Hänsch-Couillard detector. This defines the output power of the entire laser system post combination and, therefore, defined the

system power. In Figure 5.16, after the Fresnel beam splitter is the compressor and the single shot autocorrelator. To enable the simultaneous use of the single shot autocorrelator and the polarisation analysis diagnostic without moving the optics, a 70:30 BS was added. The transmitted beam was directed into the power meter and rescaled to show the power that would have been measured directly in front of the 70:30 BS. The reflected beam was directed towards the compressor and single shot autocorrelator. Once the measurements had been performed with the locked system, the system was unlocked, and power measurements were taken at Positions 2-7, 10 and 11 in Figure 5.1. These unlocked power measurements allowed the degree of linear polarisation of each PCF to be determined, along with the system efficiency of the experiment, and the gain of each PCF. These power measurements are presented later in Chapter 8.

In addition to the power meter, a beam profiling camera and a photodiode (wavelength range of 350-1100 nm with a bandwidth of 25 MHz) were located in the diagnostics area. These were useful during the alignment phase of the experiment to help find and improve the spatial and temporal overlap of the two pulses. During a running experiment, measurements of the beam profiles from each PCF, before and after combination, could be taken and the stability of the Hänsch-Couillard polarisation lock could be assessed. Any phase differences seen between the two pulses due to path length variations led to fluctuations in the photodiode signal, which were recorded. By measuring the photodiode signal of the combined locked and unlocked beams over an extended period of time, the stability of the system could be characterised by comparing the two photodiode signals. The photodiode signal of the locked state allowed calculation of the RMS stability of the locking system. The photodiode signal of the unlocked state served to highlight the sources of noise affecting the interferometer.

5.9 Summary

The experimental setup presented here was a filled aperture, Hänsch-Couillard polarisation locked, coherent combination experiment, designed and built by myself for two doped PCFs with different lengths, core diameters and internal structures. The combined beam was stabilised using Hänsch-Couillard polarising locking, with a Fresnel wedge taking a 1% reflection and directing it into a Hänsch-Couillard detector. The stabilised combined beam

had its power reduced by a 70:30 beam splitter, just before Point 8 on Figure 5.1, with 30% of the beam power being sent to the compressor and single shot autocorrelator for analysis, while the other 70% was captured by the power meter. The power sent to the compressor and single shot autocorrelator was reduced to avoid damaging the compressor gratings and the single shot autocorrelator's BBO crystal. For higher power systems, larger beam sizes and diffraction gratings could be used to mitigate the damage experienced by a compressor, but non-linear crystals used in single shot autocorrelators have costs that scale non-linearly with size so are kept small. Hence, to avoid damage, power reductions are necessary before the beam is directed towards a single shot autocorrelator.

Chapter 6

Seed Laser Characterisation

The seed laser used for the experiments in this thesis was not well understood initially. Being a commercial system, many of the details of the seed laser internals were proprietary, so only a rough estimate for the stretched pulse duration was known (which later turned out to be wrong). Many details such as the design of the stretcher and the pulse energy before and after the amplifiers required reverse engineering, measurement and inference to work out how the laser system functioned. Details of the seed laser output were provided in the specification, but given that the compressor had been removed to allow further amplification, the compressed pulse shape and duration would no longer match the seed laser's original specification.

To build a compressor to compress the amplified, coherently combined pulses optimally, the phase of the system should be known. This means that the phase acquired by a pulse should be known for both the photonic crystal fibre (PCF) amplifiers and the seed laser itself. Furthermore, the experiment aimed to show that two mismatched PCF amplifiers could be coherently combined to give high values for the degree of linear polarisation and the fraction of available beam power, while not degrading the pulse's temporal profile. This is to say that the pulse that enters the splitting polarising beam splitter (PBS) of the interferometer should be closely matched to the pulse leaving the combining PBS of the interferometer. The coherently combined pulse should be compared to the seed pulse that exits the seed laser system, and is presented to the interferometer, not the pulse which emanates from the oscillator. Therefore it is necessary to know what the seed pulse looks like upon exit from the seed laser system.. This chapter introduces the seed laser and discusses what was done to characterise its output.

6.1 Seed laser system

6.1.1 Seed laser diagram

The commercial system used as the seed laser was quoted as being capable of producing output pulses with 2 μJ of energy, pulse durations of ∼500 fs, at a wavelength of 1032 nm. The system used chirped pulse amplification (CPA) to reduce the potential for damage and non-linearities, as the pulse passed through two stages of pre-amplification. The CPA system was made up of a grating stretcher and grating compressor. This compressor was removed from the system to enable further amplification of the pulse. This allowed the seed laser to generate pulses with a stretched pulse duration of $\sim$ 290 ps. The layout of the seed laser is shown in Figure 6.1.

6.1.2 Oscillator

The oscillator was constructed from a tuneable length cavity with a Yb^{3+}:YAG crystal gain medium. The oscillator produced a 25 nJ pulse at a repetition rate of 51.92 MHz. Its spectrum was centred on 1032 nm and had a bandwidth of 7.5 nm FWHM, as shown in Figure 6.2. This produced a transform-limited pulse of $\sim$ 210 fs. The pulse duration was verified by taking an autocorrelation, which showed a FWHM of 301 ± 23 fs, as seen in Figure 6.3. This provided a FWHM deconvolved pulse duration of 213 ± 16 fs.

6.1.3 Stretcher

The stretcher was a four pass, single grating stretcher using a 1750 l/mm transmissive grating to stretch pulses from $\sim$ 210 fs to $\sim$ 385 ps. The setup is shown in Figure 6.4. In a typical two grating stretcher, the presence of the lenses allowed an image of the first grating to be created behind the second, as shown in Figure 4.2. The distance between the image of the first grating and the second grating gave the stretcher length (negative D). This was used in Equations (4.14-4.16) to work out the acquired phase.

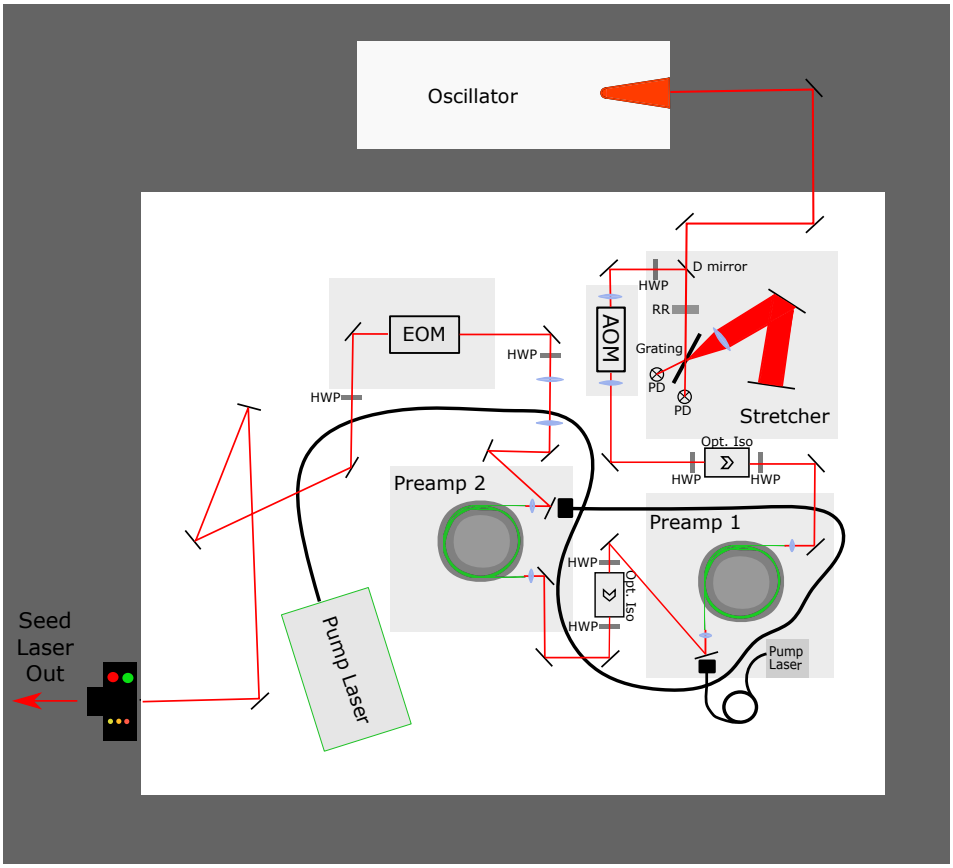


Figure 6.1: The layout of the seed laser, a bespoke design by Amplitude Systême. RR is a roof retro-reflector, PD is a photodiode, HWP is a half wave plate and 'Opt. Iso' is an optical isolator. The oscillator produced $\sim$ 210 fs pulses at a central wavelength of 1030 nm, with 4.5 nJ of energy and at a repetition rate of 51.92 MHz. These pulses were fed into a single 1750 l/mm transmission grating stretcher, which used a roof retro-reflector and a large mirror to fold the beam over four times, within the same beam path. This stretcher is discussed in more detail in Figure 6.4. The pulses were chirped to $\sim$ 385 ps, before they exited the stretcher and were sent into the acoustic optical modulator, which was used to reduce the repetition rate of the laser from 51.92 MHz to 6.49 MHz. An optical isolator was the next component in the beam line, and was used to protect the acoustic optical modulator from any parasitic lasing, caused by the pre-amplifier. The output of the pre-amplifier fed into the second pre-amplifier, known as the main-amplifier, via another optical isolator to protect the pre-amplifier from the main-amplifier. The pre-amplifier was pumped with a 3 W pump laser and increased the pulse energy from 4.5 nJ to 182 nJ. The second pre-amplifier, which performed the majority of the amplification, was pumped with a 50W pump laser and increased the pulse energy from 110 nJ to 2.7 µJ. Finally, the amplified pulses were sent through an electric optical modulator, which enabled the pulse train to be electronically gated. Pulses leaving the EOM were directed towards the output shutter.

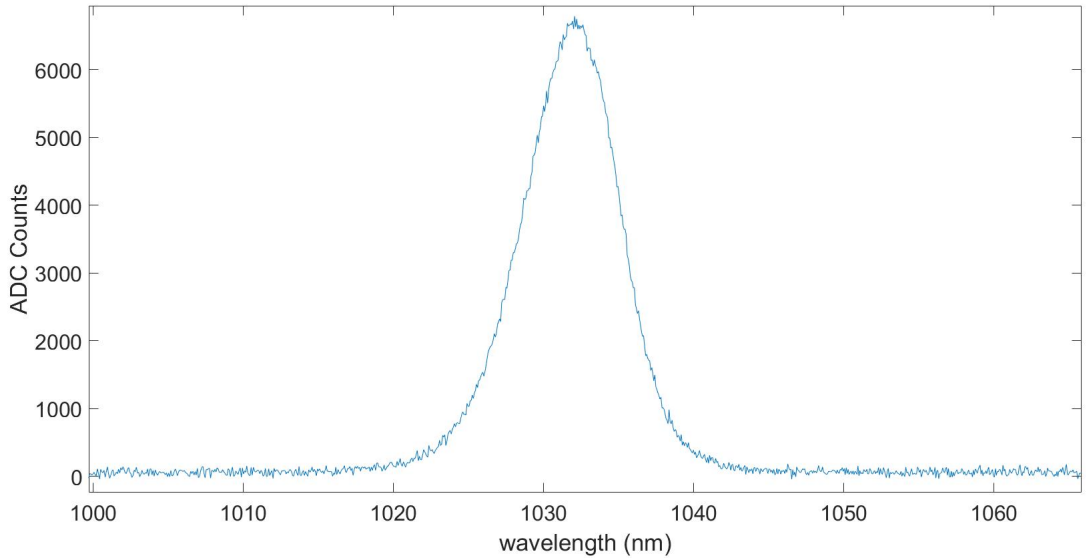


Figure 6.2: Spectral output of the oscillator, measured with an Ocean Optics USB2000 spectrometer supporting a resolution of 0.1 nm per pixel.

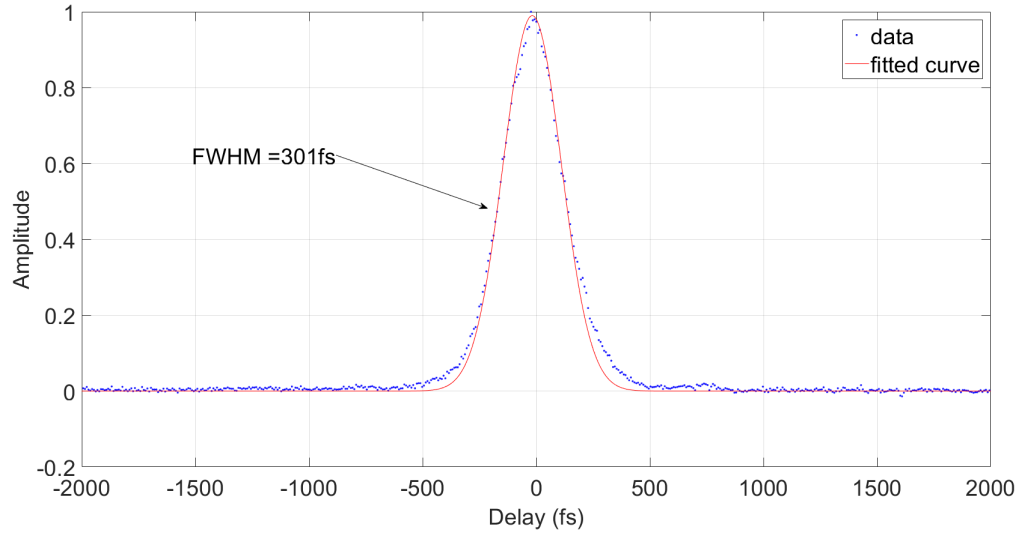


Figure 6.3: The autocorrelation of the pulse with the spectrum as seen in Figure 6.2, taken using the single shot autocorrelator, as described in Section 5.6. A Gaussian shows a good fit to the autocorrelation, with a FWHM of 301 ± 23 fs. The Gaussian fit meant that the autocorrelation could be deconvolved using the deconvolution factor of $\sqrt{2}$, producing a pulse duration of 213 ± 16 fs.

A folding mirror was added in-between the gratings in Figure 4.2 to retro-reflect the beam back towards the first grating. However, the beams would have to pass through the grating at four different beam heights to allow them to be treated separately, as seen in Figure 6.4. Photodiodes, were positioned behind the grating to observe the reflected and transmitted beams. The seed laser controller used these as diagnostics to prevent the amplifiers from being pumped before they were seeded. They also provided a reference signal to synchronise the acoustic optical modulator to the oscillator.

6.1.4 Acoustic Optical Modulator

The acoustic optical modulator (AOM) was present in a pulse picker configuration to reduce the repetition rate of the system. The AOM was driven by a Radio Frequency (RF) generator which produced a 250 MHz RF signal modulated by a 0-1V TTL signal at 6.49 MHz. The seed laser was set up to reduce the system frequency by a factor of 8 from 51.92 MHz down to 6.49 MHz. This frequency was a legacy of a previous experiment [115] where the laser system was configured to match the repetition rate of an early design of the International Linear Collider (ILC) [116].

6.1.5 Amplifiers

The seed laser system made use of two amplifiers. Both of these were fibre amplifiers, using polarisation-maintaining Yb-doped step-index fibres. Each fibre had a core diameter of 25 μm, an inner cladding diameter of 200 μm and stress rods. Figure 6.5 shows a schematic diagram of the internal structure of the fibre amplifiers.

The first amplifier, the pre-amplifier, was approximately 1.4m long and was seeded with linearly polarised, ∼4.5 nJ pulses at a repetition rate of 6.49 MHz. This amplifier increased the pulse energy to ∼ 180 nJ. The amplified seed beam was then guided towards the second amplifier, the main-amplifier, through an optical isolator to prevent back-propagating parasitic pulses from causing damage to the pre-amplifier. The main-amplifier was seeded with linearly polarised, ∼110 nJ pulses, which had seen some loss due to the optical isolator present between the pre-amplifier and main-amplifier, see Figure 6.1. The pulse energy was boosted by the main-amplifier to ∼2.6 μJ, when running at full pump power and measured at

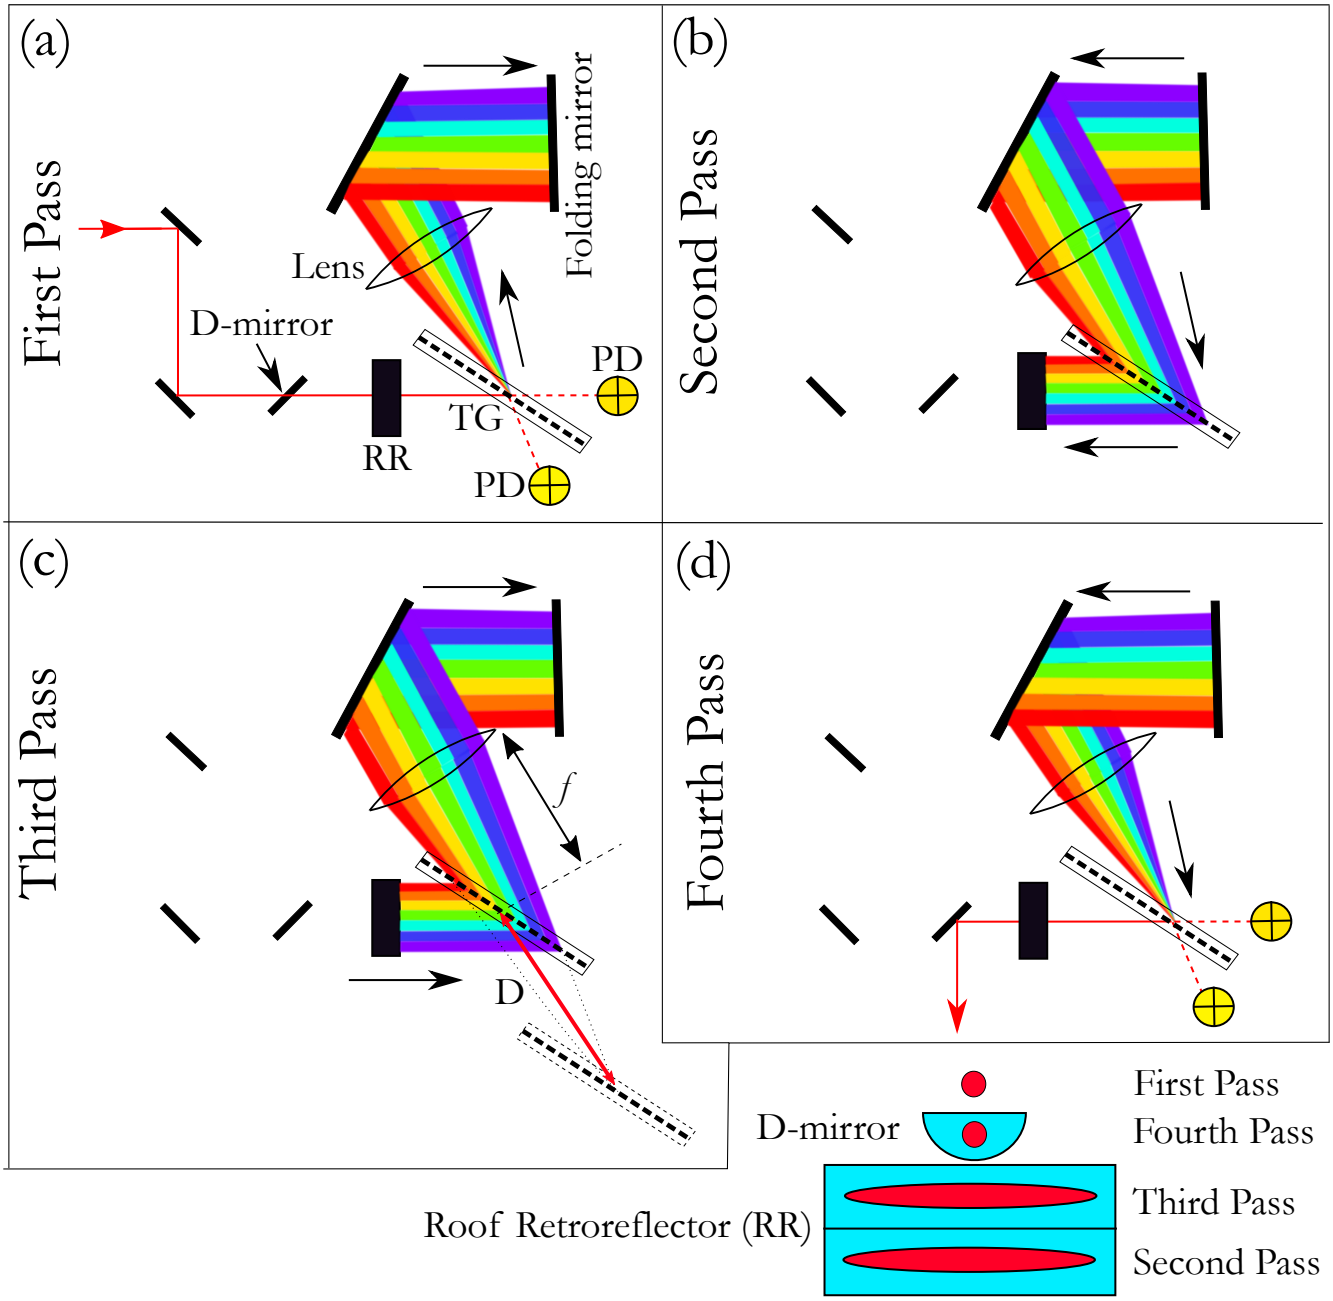


Figure 6.4: The schematic layout of the single grating stretcher and the path taken by the beam. TG is a transmissive grating, RR is a roof retro-reflector and the two photodiodes, PD, are present to synchronise the AOM with the oscillators repetition rate and to interlock the amplifiers to prevent pumping without seeding. (a) The first pass of the beam through the stretcher is at the highest beam height, passing over the top of the D-mirror, diffracting from the grating and propagating to the folding mirror. (b) The second pass of the beam is from the folding mirror, back towards to the roof retro-reflector and descending to the lowest beam height. (c) The third pass is from the roof retro-reflector back towards the folding mirror, with an increased beam height. (d) The fourth pass through the stretcher is from the folding mirror, over the roof retro-reflector and towards the D-mirror. The image of the grating used to ascertain the grating separation, D, and the lens-grating separation, f, has been illustrated in Figure (c).

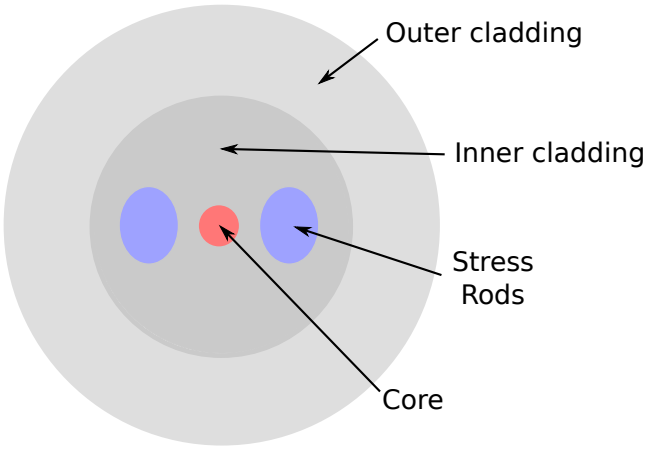


Figure 6.5: A schematic diagram of the internal structure of the pre-amplifier and main-amplifier fibres in the seed laser system. This is a Yb-doped polarisation-maintaining, step index fibre.

the end of the fibre. The main-amplifier pump power could be varied by altering the current driving the diode pump laser. Currents from 0-19A could be chosen, but for consistency, every experiment using this laser system, regardless of the intended average output power, had the pump laser current set to 19A. The output power of the seed laser was adjusted externally to the laser using a HWP and polariser. This ensured that the seed laser was operating at the same temperature and, therefore, the same conditions, every time it was used. Figure 6.6 shows a laboratory photograph of the two fibre amplifiers, while being seeded and pumped.

6.1.6 Electro Optical Modulator

The electro-optical modulator (EOM) allows the beam to be temporally modulated for short pulse train generation. It was not used for this experiment, but it was in the beam path of the laser, so it is covered here for completeness.

6.1.7 System pulse characteristics

The system was terminated with an interlocked shutter, and this point defined the output of the seed laser. Pulses with energies of $\sim$1.6 μJ, durations of $\sim$290 ps, degrees of linear polarisation of $\sim$ 99% and at repetition rates of 6.49 MHz were measured from the output of the shutter. The amount of non-linear phase acquired by the pulse was estimated by measuring the input and output pulse energies of the pre-amplifier and main-amplifier and

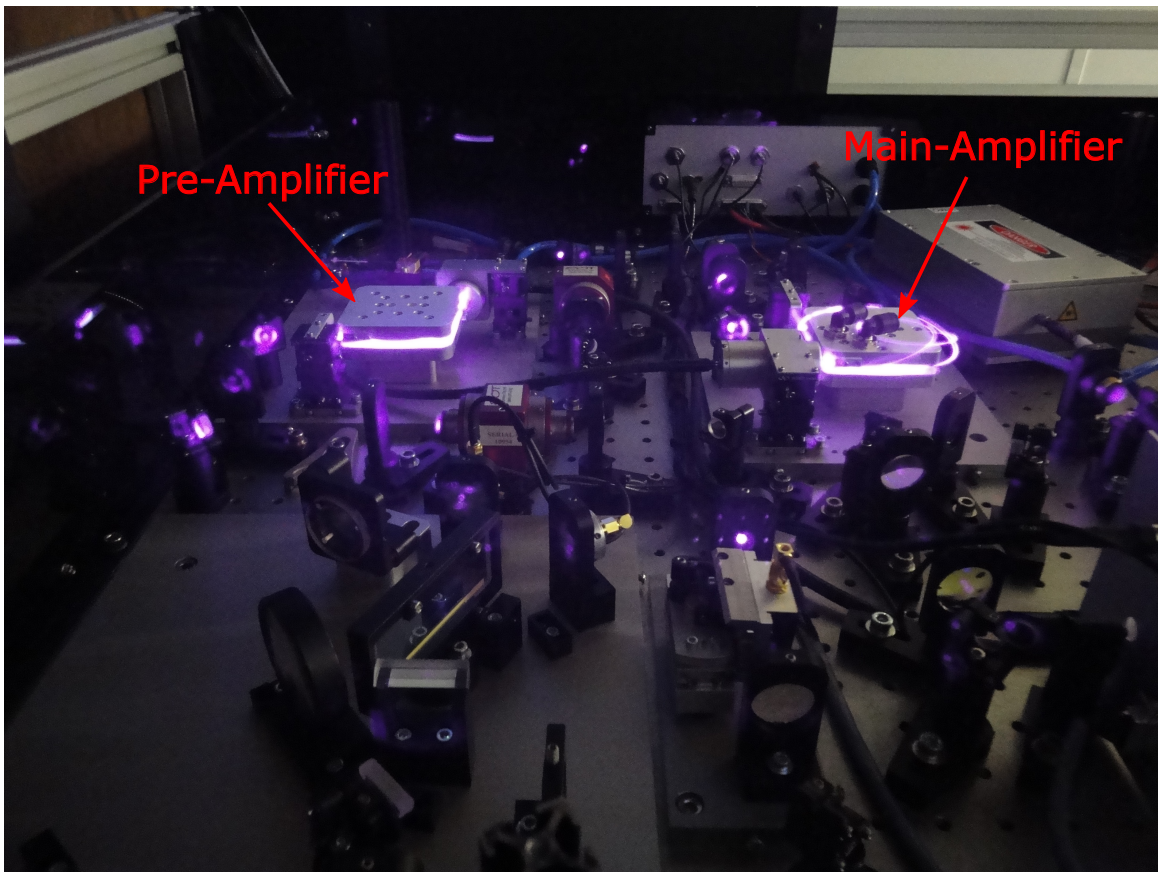


Figure 6.6: The two amplifiers in the seed laser. The pre-amplifier is shown on the left and the main-amplifier is shown on the right.

calculating the pulse peak intensity, when coupled into the fibre's core. The peak intensity is defined as

$$I_p = \frac{2P_p}{A_{mode}}, \tag{6.1}$$

where A_{mode} is the mode field area and P_p is the peak power, which can be defined in terms of the pulse energy

$$P_p = \frac{U}{\tau_p}\sqrt{\frac{4\ln(2)}{\pi}}, \tag{6.2}$$

where U is the total pulse energy and τ_p is the FWHM pulse duration. The square root term comes from the Gaussian integral of the pulse over time and serves as a correcting factor for the simpler top hat approximation of the peak power. The intensity can then be defined as

$$I_p = \frac{2U}{A_{mode}\tau_p}\sqrt{\frac{4\ln(2)}{\pi}}. \tag{6.3}$$

A simple small-signal exponential gain was used to model the growth of intensity as the pulse was amplified by the fibres [103]

$$\mathrm{I_p}(\omega, \mathrm{z}) = \mathrm{I_p}(\omega, 0)\mathrm{e}^{\alpha(\omega-\omega_0)\mathrm{z}}, \tag{6.4}$$

where $\mathrm{I_p}(\omega, \mathrm{z})$ is the peak intensity of the pulse after travelling a distance z through the amplifier, $\mathrm{I_p}(\omega, 0)$ is the initial peak intensity, $\alpha(\omega - \omega_0)$ is the gain coefficient and ω_0 is the central frequency. Use of this simple gain model is justified as long as the system is operating below the saturation energy of the amplifier [75]. The saturation energy is defined as the incident pulse energy required to reduce the gain to 1/e of its value. It comes from the idea that to cause stimulated emission of one photon, dA/σ_{21} photons are required, where dA is a small element of area. Integrating dA over the mode field area gives the total number of photons in a pulse, which can be converted into an energy by multiplying by the photon energy

$$\mathrm{E_{sat}} = \frac{\mathrm{A_{mode}}\hbar\omega_0}{\sigma_{21}}, \tag{6.5}$$

where $\mathrm{A_{mode}}$ is the beam mode area, ω_0 is the lasing frequency and σ_{21} is the stimulated emission cross-section, which for $\mathrm{Yb^{3+}}$ has been measured to be $0.6\,\mathrm{pm^2}$ [117]. Below saturation, a small-signal-gain model for amplification could be used, whereas values above saturation should use a linear-gain-model instead [103]. Using a wavelength of 1032 nm and a beam mode radius of 10 μm, the saturation energy for both seed laser amplifiers was calculated to be 101 μJ. Given the total pump power available to the seed laser main-amplifier was 50 W, and the maximum output pulse energy of the main-amplifier was 2.7 μJ, the system was operating well below the saturation energy. Hence, modelling an exponentially increasing pulse energy was acceptable for this simple model. To calculate the maximum non-linear phase acquired by the pulse, the intensity of the central wavelength component taken at the centre of the spatial intensity distribution was used. By setting $\omega = \omega_0$, this simplifies Equation (6.4) to

$$\mathrm{I_p}(\omega_0, \mathrm{z}) = \mathrm{I_p}(\omega_0, 0)\mathrm{e}^{\alpha\mathrm{z}}. \tag{6.6}$$

	Pre-amplifier	Main-amplifier
Fibre length (m)	1.4	2.7
Pulse energy in (nJ)	4.5	110
Pulse energy out (μJ)	0.180	2.7
Calc. α (at ω_0)	2.6508	1.1596
non-linear phase (at ω_0) (rads)	0.16	7

Table 6.1: Characteristics of the two amplifiers within the seed laser system. The system operates at a central wavelength of 1032 nm. Both fibre amplifiers have 25 μm core diameters with inner cladding diameters of 200 μm.

Rearranging this for α gives

$$\alpha = \frac{\ln(\frac{\mathrm{I_p}(\omega_0,\mathrm{z})}{\mathrm{I_p}(\omega_0,0)})}{\mathrm{z}}. \tag{6.7}$$

Substituting Equation (6.3) into Equation (6.7) gives

$$\alpha = \frac{\ln\frac{\mathrm{U}(\omega_0,\mathrm{z})}{\mathrm{U}(\omega_0,0)})}{\mathrm{z}}. \tag{6.8}$$

So, by measuring the pulse energy before and after amplification, α could be calculated. This was substituted into Equation (4.1), using $\mathrm{n_2} = 2.7 \times 10^{-27}\ \mathrm{W^{-1}\,m^2}$ (which was assumed to be constant for all wavelengths over the bandwidth of the pulse [118]) and integrated over z. The results are shown in Table 6.1. The non-linear phase seen by the pulse at ω_0 in the main amplifier was significant and it was likely that it could not be compensated for by the compressor. Uncompensated phase results in the compressed pulse not having the shortest pulse duration possible for the given bandwidth. In addition, the spectrum of the seed pulse, after propagating through the pre-amplifier and main-amplifier, saw a significant amount of spectral narrowing. The reduced bandwidth available meant that the pulse duration would not be as short as the pulses emanating from the oscillator. This spectral evolution can be seen in Figure 6.7.

A Gaussian fit to the central portion of the spectrum after the main-amplifier gave a FWHM of ∼3.5 nm. The transform-limited pulse duration expected was ∼ 440 fs. However, the sidelobes present in the spectrum contributed to the pulse duration also, causing the measured pulse duration to be 373 ± 11 fs, as seen in Figure 6.8. Given the significant quantity of non-linear phase acquired by the pulse in the main-amplifier, leading to the spectral features seen in Figure 6.7, it was deemed necessary to use a more sophisticated

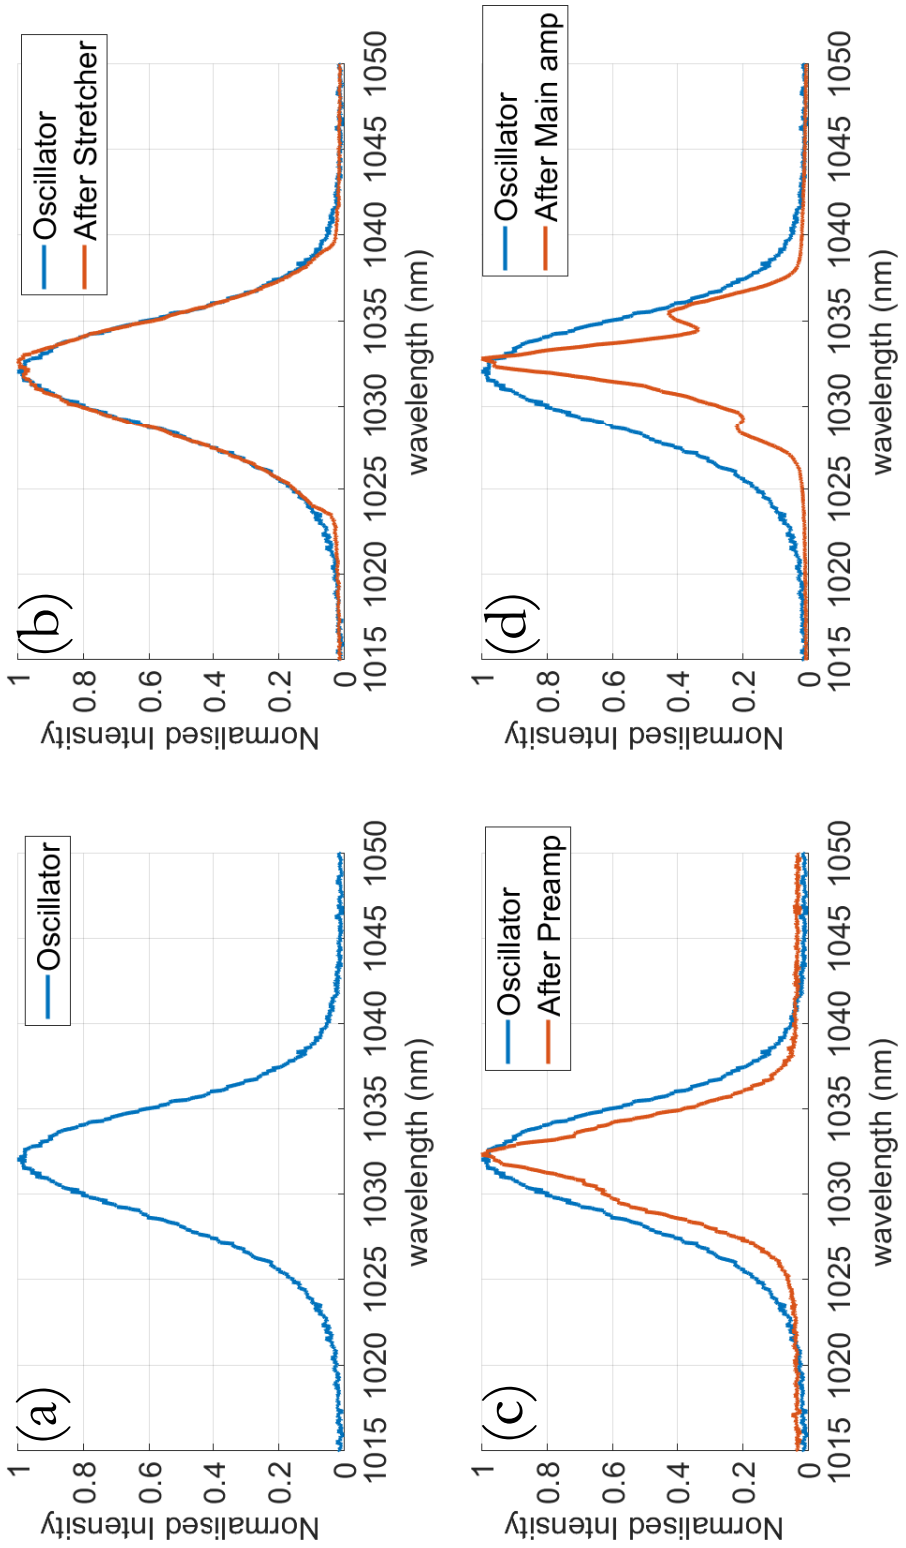


Figure 6.7: The normalised, seed laser output spectrum was seen to be narrowing as it propagated through the seed laser. (a) The oscillator spectrum. (b) The spectrum seen after the stretcher. Some clipping had occurred due to the finite size optics used in the stretcher. (c) The spectrum seen after the pre-amplifier. (d) The spectrum seen after the main-amplifier, including the side lobes caused by self-phase modulation of the pulse. The spectrum was measured with an Ocean Optics USB2000 spectrometer, supporting a resolution of 0.1 nm per pixel.

method to model the pulse evolution through the pre- and main-amplifier, and to reproduce the pulse autocorrelation as seen in Figure 6.8.

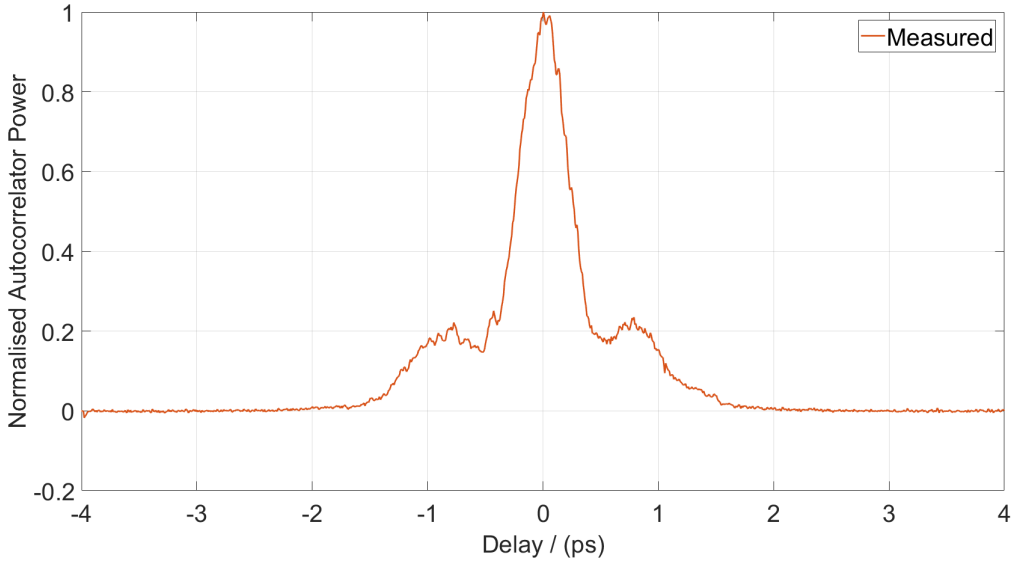


Figure 6.8: The autocorrelation of the seed laser after it had been compressed optimally. The autocorrelation FWHM was calculated to be $\sim 526 \pm 16$ fs, which was deconvolved to a pulse duration of $\sim 373 \pm 11$ fs. assuming a Gaussian pulse shape in time.

6.2 Split Step Fourier method

Given that there was a significant amount of non-linear phase at ω_0 acquired by a pulse propagating through the seed laser's main-amplifier, the simpler model, which applied the dispersion and non-linear phase consecutively, would no longer be valid. To account for the effects of the significant non-linear phase acquired by the seed laser pulse, the non-linear phase and dispersion should be calculated concurrently rather than in a piecewise fashion. For a laser pulse to exist, its mathematical form must be a solution to the wave equation. Therefore, the description of the pulse in terms of an electric field is

$$\mathrm{E}(\mathrm{z}, \omega) = \mathrm{A}(\mathrm{z}, \omega - \omega_0)\mathrm{e}^{\mathrm{i}\beta_0 \mathrm{z}}, \tag{6.9}$$

where β_0 is the wave vector at ω_0, z is the propagation distance, ω_0 is the central frequency of the pulse's spectrum and A is the amplitude of the pulse's electric field, which varies slowly in z. Substitution of Equation (6.9) into the wave equation gives

$$\nabla^2 \left(\mathrm{A}(\mathrm{z}, \omega - \omega_0) \mathrm{e}^{\mathrm{i}\beta_0 \mathrm{z}} \right) + \beta^2 \mathrm{A}(\mathrm{z}, \omega - \omega_0) \mathrm{e}^{\mathrm{i}\beta_0 \mathrm{z}} = 0. \tag{6.10}$$

Some mathematical manipulation is required, to turn Equation (6.10) into that of Equation (6.14) [119], where Equation (6.14) is the general non-linear Schrödinger equation. To get from Equation (6.10) to Equation (6.14), the slowly varying envelope approximation is invoked, which sets $\partial_{\mathrm{z}}^2 = 0$ and is valid in the case where the envelope of the forward travelling wave varies slowly over one period of oscillation of the carrier. The wave vectors β and β_0 are assumed to be similar in magnitude, allowing some simplification of the mathematical form. Then, β is substituted for a linear expansion of terms

$$\beta = \beta_0 + \Delta\beta_{\mathrm{NL}} + \beta_1(\omega - \omega_0) + \beta_2(\omega - \omega_0)^2 + ..., \tag{6.11}$$

where

$$\Delta\beta_{\mathrm{NL}} = \gamma|\mathrm{A}|^2, \qquad \beta_1 = \frac{\partial\beta}{\partial\omega} = \frac{1}{\mathrm{V_g}(\omega)} \qquad \text{and} \qquad \beta_2 = \frac{\partial^2\beta}{\partial\omega^2} = \frac{1}{\mathrm{V_g}(\omega)^2}\frac{\partial \mathrm{V_g}(\omega)}{\partial\omega}. \tag{6.12}$$

$\Delta\beta_{\mathrm{NL}}$ describes the non-linear interactions of the pulse with the medium, where γ is the non-linear factor, while β_1 and β_2 describe the dispersive aspects of the medium, where $\mathrm{V_g}(\omega)$ is the group velocity. γ is defined as

$$\gamma(\omega_0) = \frac{\omega_0 \mathrm{n}_2}{\mathrm{c}\mathsf{A}_{\mathrm{eff}}} \qquad \mathsf{A}_{\mathrm{eff}} = \frac{(\int\int_{-\infty}^{\infty} |\mathrm{F}(\mathrm{x},\mathrm{y})|^2 \mathrm{dxdy})^2}{\int\int_{-\infty}^{\infty} |\mathrm{F}(\mathrm{x},\mathrm{y})|^4 \mathrm{dxdy}}, \tag{6.13}$$

where n_2 is the non-linear refractive index, as defined in Equation (4.2), $\mathrm{A_{eff}}$ is the effective area of the fibre core and $\mathrm{F}(\mathrm{x},\mathrm{y})$ is the transverse spatial profile of the laser pulse propagating down the fibre core. The transverse spatial profile is used in the calculation of the non-linear factor only. Some further mathematical manipulation leads to Equation (6.14), the general non-linear Schrödinger equation

$$\left[\frac{\partial \mathrm{A}}{\partial \mathrm{z}}\right]_1 + \left[\beta_1 \frac{\partial \mathrm{A}}{\partial \mathrm{t}}\right]_2 + \left[\sum_{\mathrm{n}=2}^{\infty} \frac{\mathrm{i}^{\mathrm{n}}}{\mathrm{n}!}\beta_{\mathrm{n}} \frac{\partial^{\mathrm{n}} \mathrm{A}}{\partial \mathrm{t}^{\mathrm{n}}}\right]_3 + \left[\frac{\alpha}{2}\mathrm{A}\right]_4 = \left[\mathrm{i}\gamma|\mathrm{A}|^2\mathrm{A}\right]_5 - \left[\frac{\gamma}{\omega_0}\frac{\partial}{\partial \mathrm{t}}(|\mathrm{A}|^2\mathrm{A})\right]_6 - \left[\mathrm{i}\gamma \mathrm{T_R} \mathrm{A}\frac{\partial|\mathrm{A}|^2}{\partial \mathrm{t}}\right]_7. \tag{6.14}$$

The general non-linear Schrödinger equation is widely used for modelling the pulse propagation through different media, as the equation is able to account for different effects. From Equation (6.14); Term 1 describes the change of the amplitude, as the pulse propagates along z; Term 2 describes the propagation of the envelope; Term 3 describes n orders of material dispersion, which affect the shape of the envelope; Term 4 describes the effect of absorption within the medium; Term 5 describes the non-linear phase contribution (self-phase modulation (SPM)) due to the fibre non-linearity; Term 6 describes the effect of self-steepening (SS); and Term 7 describes the effect of intrapulse Raman scattering (IRS). β_{n} is the nth dispersion coefficient, α is the absorption parameter and T_R is the Raman timescale [60]. A has been scaled, such that $|\mathrm{A}|^2$ is in units of power and t is measured in a frame of reference moving with the pulse at the group velocity.

Depending upon the characteristics of the pulse and the fibre in which it propagates, simplifications can be made to the general non-linear Schrödinger equation, if the effects of some of the non-linear terms (Terms 5, 6 and 7) can be neglected. The non-linear terms (Terms 5, 6 and 7) each depend upon γ which, for an effective area of 160 $\mu\mathrm{m}^2$ (a fibre with a 25 μm core, as used for the seed laser pre-amplifier and main-amplifier) gives a value of $\gamma \approx 0.001$ at a wavelength of $\lambda = 1032$ nm. The non-linear terms of self-steepening and intrapulse Raman scattering depend upon the gradient of the pulse envelope. The gradient of the intensity can be used to find the magnitude of Terms 5 and 7, but for Term 6, some mathematical manipulation is required to get it into a more useful form. Term 6 can be rewritten as [120]

$$\frac{\partial}{\partial \mathrm{t}}(|\mathrm{A}|^2\mathrm{A}) \equiv \mathrm{A}\frac{\partial}{\partial \mathrm{t}}|\mathrm{A}|^2 + |\mathrm{A}|^2\frac{\partial}{\partial \mathrm{t}}\mathrm{A}, \tag{6.15}$$

where $|\mathrm{A}|^2 \equiv \mathrm{A}\bar{\mathrm{A}}$ (where $\bar{\mathrm{A}}$ is the complex conjugate) to allow a factor of A to be pulled out

$$\mathrm{A}\frac{\partial}{\partial \mathrm{t}}|\mathrm{A}|^2 + |\mathrm{A}|^2\frac{\partial}{\partial \mathrm{t}}\mathrm{A} = \mathrm{A}\left[\frac{\partial}{\partial \mathrm{t}}|\mathrm{A}|^2 + \bar{\mathrm{A}}\frac{\partial}{\partial \mathrm{t}}\mathrm{A}\right]. \tag{6.16}$$

where the size of the self-steepening term can now be described as the sum of the gradient of the intensity and that of the electric field. To determine the the impact of Terms 5, 6 and 7, a simple model was built, where a transform-limited Gaussian pulse with a pulse energy of 2.7 μJ (maximum seed laser pulse energy), and a pulse duration of 200 fs, was stretched to

290 ps using a simulated stretcher, with an incident angle of 64°, a grating separation of 240 mm, and a grating line density of 1750 l/mm. The pulse was propagated down 2.4 m of fibre, with an effective area of 160 μm^2, giving a non-linear factor of $\gamma \approx 0.001$ at $\lambda = 1032$ nm. The gradient of A and $|A|^2$ at all points could be found and substituted into the self-phase modulation (SPM), self-steepening (SS) and intrapulse Raman scattering (IRS) terms

$$\mathrm{SPM} = \mathrm{L}\gamma|\mathrm{A}|^2\mathrm{A}, \tag{6.17}$$

$$\mathrm{SS} = \frac{\mathrm{LA}\gamma}{\omega_0}\left[\frac{\partial}{\partial \mathrm{t}}|\mathrm{A}|^2 + \bar{\mathrm{A}}\frac{\partial}{\partial \mathrm{t}}\mathrm{A}\right], \tag{6.18}$$

and

$$\mathrm{IRS} = \mathrm{L}\gamma \mathrm{T_R A}\frac{\partial|\mathrm{A}|^2}{\partial \mathrm{t}}. \tag{6.19}$$

where L is the length of the fibre. Using Equations (6.17-6.19), the quantities of self-phase modulation, self-steepening and intrapulse Raman scattering were calculated and normalised against the most significant contributor of the three. Figure 6.9 shows the comparison of the three non-linear contributions plotted on a logarithmic scale. From Figure 6.9, the impact of self-steepening and intrapulse Raman scattering for a stretched pulse of 290 ps are negligible compared to the impact of self-phase modulation. Therefore, it is justified to drop Terms 6 and 7 from Equation (6.14), giving the simplified non-linear Schrödinger equation

$$\left[\frac{\partial \mathrm{A}}{\partial \mathrm{z}}\right]_1 = -\left[\beta_1\frac{\partial \mathrm{A}}{\partial \mathrm{t}}\right]_2 - \left[\sum_{\mathrm{n}=2}\frac{\mathrm{i}^\mathrm{n}}{\mathrm{n}!}\beta_\mathrm{n}\frac{\partial^\mathrm{n}\mathrm{A}}{\partial \mathrm{t}^\mathrm{n}}\right]_3 - \left[\frac{\alpha}{2}\mathrm{A}\right]_4 + \left[\mathrm{i}\gamma|\mathrm{A}|^2\mathrm{A}\right]_5. \tag{6.20}$$

Unless specific cases are chosen for Equation (6.20), where there is either zero contribution from the dispersion or zero contribution from the non-linear effects, Equation (6.20) cannot be solved analytically, so numerical methods are required [60]. To model the pulse propagation, the split-step Fourier method was chosen [60]. The split step Fourier method approaches the problem by solving the equation in piecewise fashion over small increments of z distance. The method calculates the evolution of the pulse envelope, due to the effects of material dispersion and self-phase modulation in an alternating fashion. This is to say that over a small increment of z, the non-linear contribution is assumed to be zero allowing the above equation to be solved

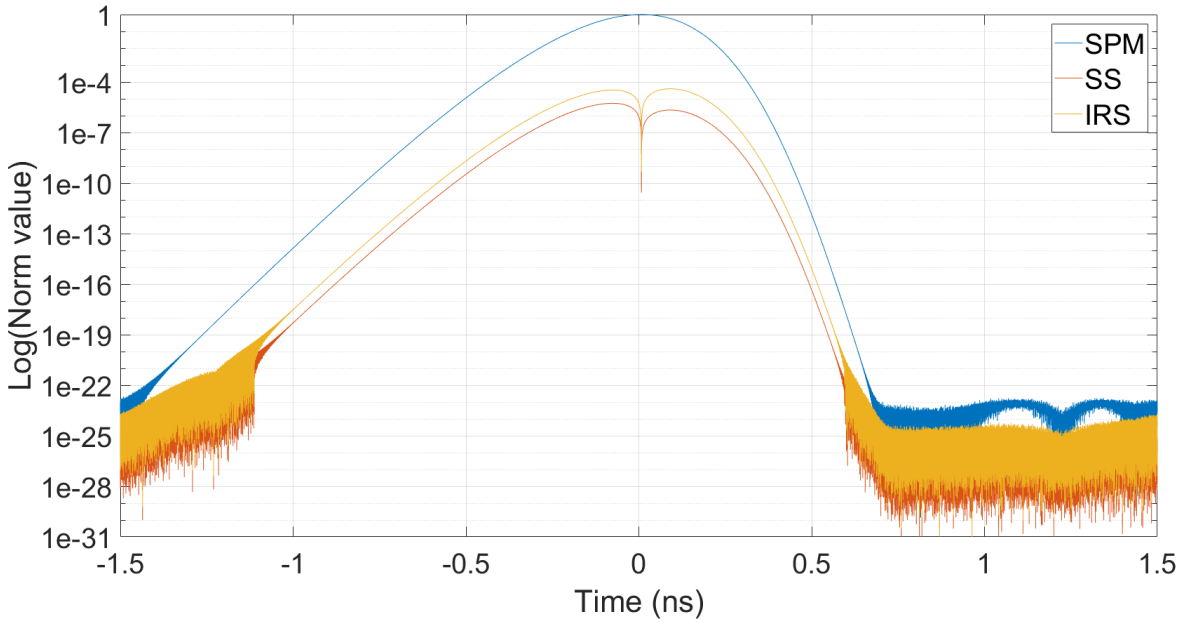


Figure 6.9: The comparison of the self-phase modulation (SPM), self-steepening (SS) and intrapulse Raman scattering (IRS) acquired by a stretched pulse of 290 ps travelling down the core of a 2.4 m fused silica fibre with an effective area of 160 μm^2. As can be seen, the self-steepening and intrapulse Raman scattering provided contributions, which were orders of magnitude smaller than that of the self-phase modulation over the length of the main-amplifier fibre. The pulse was simulated as a transform-limited, 200 fs FWHM Gaussian pulse with a pulse energy of 2.7 µJ, stretched to 290 ps using a simulated stretcher. The simulated stretcher had an incident angle of 64°, a grating separation of 240 mm, and a grating line density of 1740 l/mm. The non-linear factor of the fibre was $\gamma \approx 0.001$ at $\lambda = 1032$ nm.

analytically in the frequency domain. Then over another small increment of z, the effects of self-phase modulation are applied to the pulse in the time domain, while the effects of dispersion are assumed to be zero. This process alternates back and forth between the time and frequency domains, until the end of the fibre has been reached. To make Equation (6.20) more suitable for the split-step Fourier method, Equation (6.20) can be decomposed into a dispersive and non-linear term, which simplifies Equation (6.20) to

$$\left[\frac{\partial \mathrm{A}}{\partial \mathrm{z}}\right]_1 = \left(\hat{\mathrm{D}} + \hat{\mathrm{N}}\right) \mathrm{A}, \tag{6.21}$$

where,

$$\hat{\mathrm{D}} = -\left[\beta_1 \frac{\partial}{\partial \mathrm{t}}\right]_2 - \left[\sum_{\mathrm{n}=2} \frac{\mathrm{i}^{\mathrm{n}}}{\mathrm{n}!} \beta_{\mathrm{n}} \frac{\partial^{\mathrm{n}}}{\partial \mathrm{t}^{\mathrm{n}}}\right]_3 - \left[\frac{\alpha}{2}\right]_4, \tag{6.22}$$

and

$$\hat{\mathrm{N}} = \left[\mathrm{i}\gamma |\mathrm{A}|^2\right]_5. \tag{6.23}$$

The general solution to Equation (6.21) is

$$A(z+h,t) = e^{h(\hat{D}+\hat{N})}A(z,t), \tag{6.24}$$

where $A(z,t)$ is the envelope of the pulse prior to a step, $A(z+h,t)$ is the envelope of the pulse after a step of h, and h is the size of the step to employ. However, as previously mentioned, there is no way to calculate the dispersive and non-linear effects simultaneously. Therefore, it is useful to approximate Equation (6.24) as

$$A(z+h,t) = e^{h\hat{N}}e^{h\hat{D}}A(z,t), \tag{6.25}$$

where the effects of $\hat{D}$ and $\hat{N}$ are assumed to be independent over a small step, h. The $\hat{D}$ operator is applied in the frequency domain, with $\hat{N} = 0$ followed by a Fourier transform to the time domain to apply $\hat{N}$ with $\hat{D} = 0$. Including the Fourier transforms, one complete step can be represented as

$$A(z+h,t) = e^{h\hat{N}}F_T^{-1}\left[e^{h\hat{D}(i(\omega-\omega_0)}F_T\left[A(z,t)\right]\right], \tag{6.26}$$

where F_T is the Fourier transform and F_T^{-1} is the inverse Fourier transform. The accuracy of Equation (6.26) is to the second order in h, but can be improved to the third order in h by changing the way the operators are implemented [60]. Instead, the dispersion term is applied over h/2, and the non-linear term is applied at h/2 for a step size of h. Then another step of h/2 is made to apply the remaining dispersion for that particular step. This is mathematically written as

$$A(z+h,t) = e^{\frac{h}{2}\hat{D}}e^{h\hat{N}}e^{\frac{h}{2}\hat{D}}A(z,t), \tag{6.27}$$

and with the Fourier transforms included

$$A(z+h,t) = F_T^{-1}\left[e^{\frac{h}{2}\hat{D}(i(\omega-\omega_0))}F_T\left[e^{h\hat{N}}F_T^{-1}\left[e^{\frac{h}{2}\hat{D}(i(\omega-\omega_0)}F_T\left[A(z,t)\right]\right]\right]\right]. \tag{6.28}$$

Using Equation (6.28), I built a model of the laser system seen in Figure 5.1, covering the commercial seed laser, amplification via the mismatched photonic crystal fibres (PCFs), coherent combination and compression. The model began with a 4.45 nJ, 210 fs transform-limited pulse, chosen to match the measured pulse duration and pulse energy of the oscillator. Assuming a flat phase, these two quantities could be used to reconstruct the electric field of the oscillator pulse. With a known starting point, the pulse was stretched by a simulated single grating stretcher and then amplified by a simulated pre-amplifier, followed by the simulated main-amplifier. Given the simplicity involved in extending the model, the split-step Fourier method was extended to include the two PCF amplifiers, even though the non-linear phase was expected to be small. The pulse was divided into two copies, which shared the total energy available, before they were rescaled according to the core diameter of their respective PCF. The transverse change in the beam diameter impacts the non-linear factor, which directly affects the amount of non-linear phase that each pulse could acquire in each PCF. After the pulses were amplified in their respective PCFs, their amplitudes were again rescaled to represent beam size matching before the pulses were recombined. Then, the stretched combined pulse was subject to compression, using a simulated compressor, and analysis was performed to view the simulated electric field and intensity autocorrelation. This is shown schematically in Figure 6.10.

The model shown in Figure 6.10 allowed the entire experiment to be simulated from oscillator pulse to amplified combined pulse. One aspect not included in the model was the non-linear spectral generation due to self-phase modulation, leading to the side lobe features seen in Figure 6.7(d). To replicate these spectral features simply, it was assumed that the phase acquired by each frequency component in the model, due to dispersion and non-linear phase, mirrored what was happening experimentally to the laser pulse. However, the electric field envelope, remained a Gaussian in the model which did not match what was observed experimentally. Therefore, I decided to correct the electric field envelope, such that all of the frequency components in the spectrum would carry the correct phase and the correct magnitude, with respect to the central frequency component. This process describes the spectral shaping operation, shown in Figure 6.10, and was performed by masking the simulated, spectral electric field, with the shape of the root measured spectrum.

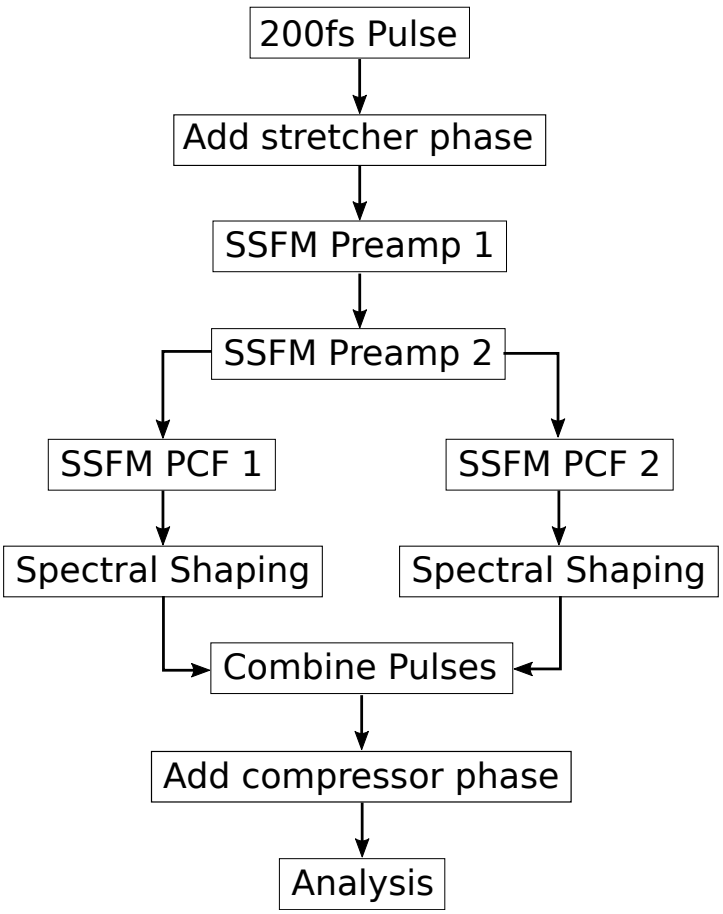


Figure 6.10: The schematic layout describing the flow of the split step Fourier model for the complete laser system.

Initially, this spectral shaping was performed after the seed laser main-amplifier in the model to reflect the spectrum measured after the seed laser. However, once the PCFs were added to the model, it suffered from some numerical instabilities, when performing the split-step Fourier method in each of the PCFs. After further investigation, I found that the spectral shaping operation, performed before the pulse propagation down each PCF, was responsible. To overcome these numerical instabilities, I moved the spectral shaping operation after the split-step Fourier method loops of the two PCFs (SSFM PCF 1/2), as shown in Figure 6.10. It was assumed that the effect of shaping the spectrum after the PCFs, rather than before, would have a minimal impact because of; the low pulse energies used; the use of chirped pulse amplification; and the very large core diameters offered by the PCFs to help combat non-linear phase.

We now move back to discuss the spectral shaping with respect to the seed laser, as was done when I initially characterised the seed laser pulse shape. The discussion which follows assumes only the seed laser is present and not the additional PCFs. Mathematically, the

spectral shaping was performed by taking the simulated electric field, $E_s(\omega)$, and masking it with the measured, normalised version of the spectral electric field, $\sqrt{S_m(\omega)}$. The mask was created by calculating

$$M(\omega) = 1 - \left(\frac{|E_s(\omega)| - \sqrt{S_m(\omega)}}{|E_s(\omega)|} \right), \tag{6.29}$$

with the mask applied to $E_s(\omega)$ by

$$E_{mask}(\omega) = M(\omega)E_s(\omega). \tag{6.30}$$

The masking operation had the effect of taking the simulated Gaussian spectrum and modifying it to what was observed by the spectrometer. Figure 6.11 shows the spectral electric field of the simulated pulse, the derived electric field from the measured spectrum and the spectrally shaped, simulated electric field. Because of the noisy background of the spectral data, I had to perform several smoothing operations to make the mask suitable for use in the model. This background noise was partly responsible for the numerical instabilities. A Butterworth filter [121] was applied to the derived, spectral electric field to smooth out the high-frequency noise. This was again applied in the time domain to smooth the background noise to zero.

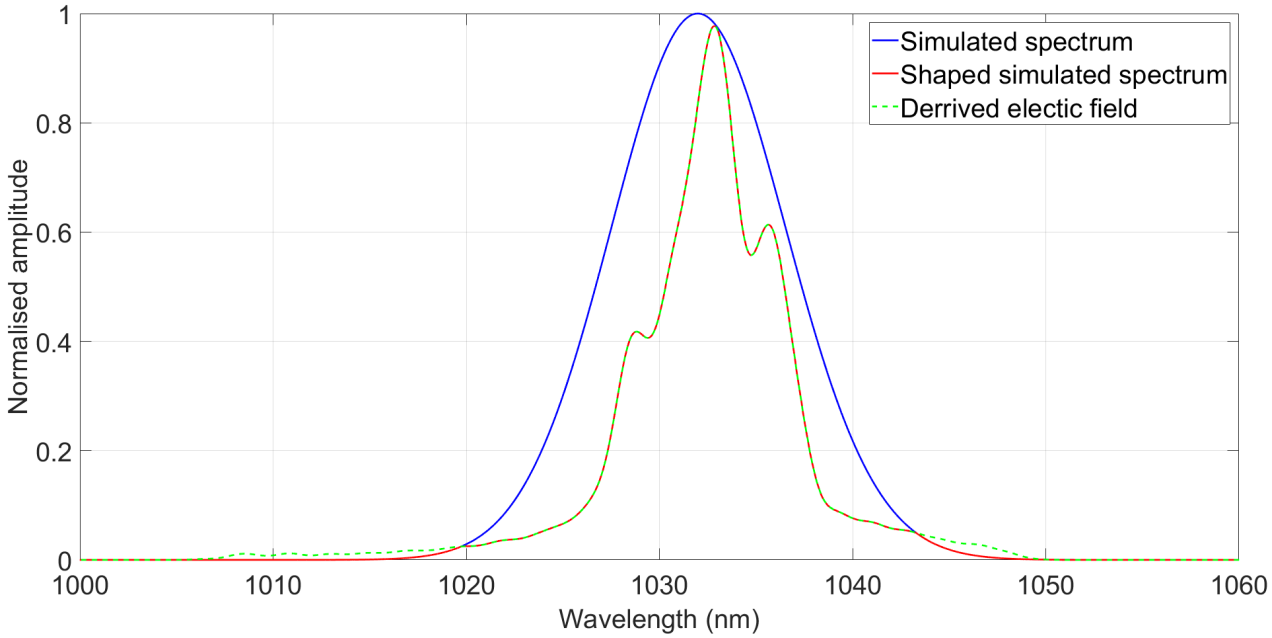


Figure 6.11: The electric field of the simulation before masking, is shown in blue The electric field derived from the measured spectrum, is shown in dashed green. The simulated spectrum shown after spectral shaping, is shown in red. The measured spectrum was taken with an ocean optics USB2000 spectrometer supporting a resolution of 0.1 nm per pixel.

The spectrally shaped pulse was sent to the simulated compressor, which had its parameters optimised to produce the largest peak intensity. An autocorrelation of the simulated intensity profile was generated, and compared against the measured seed pulse autocorrelation as seen in Figure 6.12. In terms of the central peak width, there was good agreement between the FWHM of both the simulated and measured autocorrelations. There was a small discrepancy between the wings of the two pulses, with the measured autocorrelation showing clear bumps on either side of the main central peak.

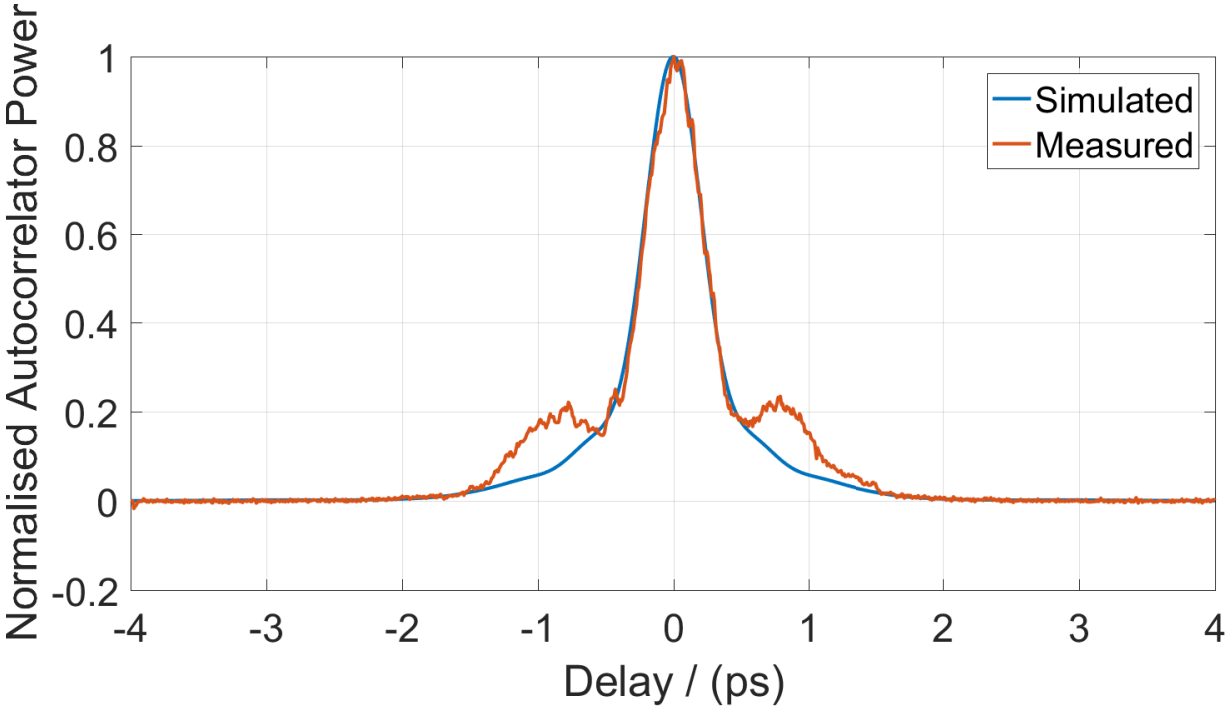


Figure 6.12: The measured autocorrelation of the seed laser, compared to the simulated autocorrelation. The measured autocorrelation FWHM was calculated to be 526 ± 16 fs, which was deconvolved to 373 ± 11 fs. The simulated autocorrelation FWHM was calculated to be 520 fs, which was produced by an intensity pulse duration of ∼370 fs.

Referring to Figure 6.13, it could be seen that a small pre-pulse was present. I hypothesised that the simple amplification model used in the split-step Fourier method code, did not amplify the pre-pulse properly, as was observed experimentally. However, without access to Frequency-Resolved Optical Gating (FROG) or Spectral Phase Interferometry for Direct Electric-field Reconstruction (SPIDER), I could not reconstruct the pulse, so this hypothesis could not be verified.

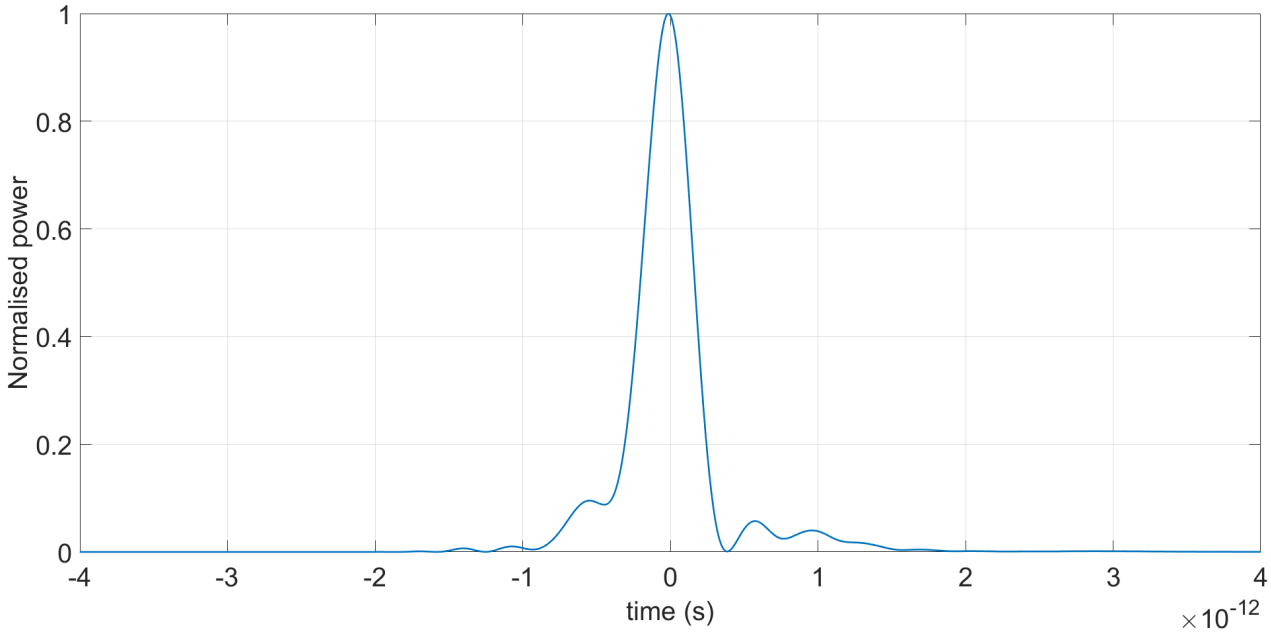


Figure 6.13: The simulated, normalised intensity profile of the seed laser. The model predicts that there would be a small pre-pulse leading the main pulse.

Despite the wings of the autocorrelations not matching between simulation and measurement, the model does predict that the pulse duration of the measured autocorrelation was the shortest pulse that could be achieved, given the amount of spectrum available. The end goal was to reproduce the seed pulse accurately, such that the phase of the frequency components matched that observed in the experiment. The only tool available to assess this was the single shot autocorrelator. The ability to produce simulated pulses, with pulse durations similar to that seen experimentally, gave confidence in the model. Furthermore, given that the compressed seed pulse autocorrelation was characterised as looking like that seen in Figure 6.12, this is shortest pulse duration that could be realistically achieved without non-linear compression. Any seed pulse which is subject to amplification, coherent combination and compression, could be no better than the compressed seed pulse seen in Figure 6.12. The loss of spectrum, as seen in Figure 6.7, prevented further compression towards the transform limit of the oscillator.

After the seed laser section of the model had been developed sufficiently, I used the model as a tool to predict the behaviour of two coherently combined pulses from two mismatched photonic crystal fibres (PCFs). The calculated pulse energies input to and output from each PCF in the high power regime (discussed in Section 8.2), were used to correctly scale the non-linear factors, γ, for each PCF. The pulse from the seed laser was split into two, and fed

into the split-step Fourier model for each fibre. Using the simulated pulses produced from each PCF, an estimate was made of the expected degree of linear polarisation and fraction of available beam power. Figure 6.14 shows the result of running the simulation. The model predicted that the coherent combination of the two PCFs should be excellent, with the degree of linear polarisation and fraction of available beam power estimated to be above 99%.

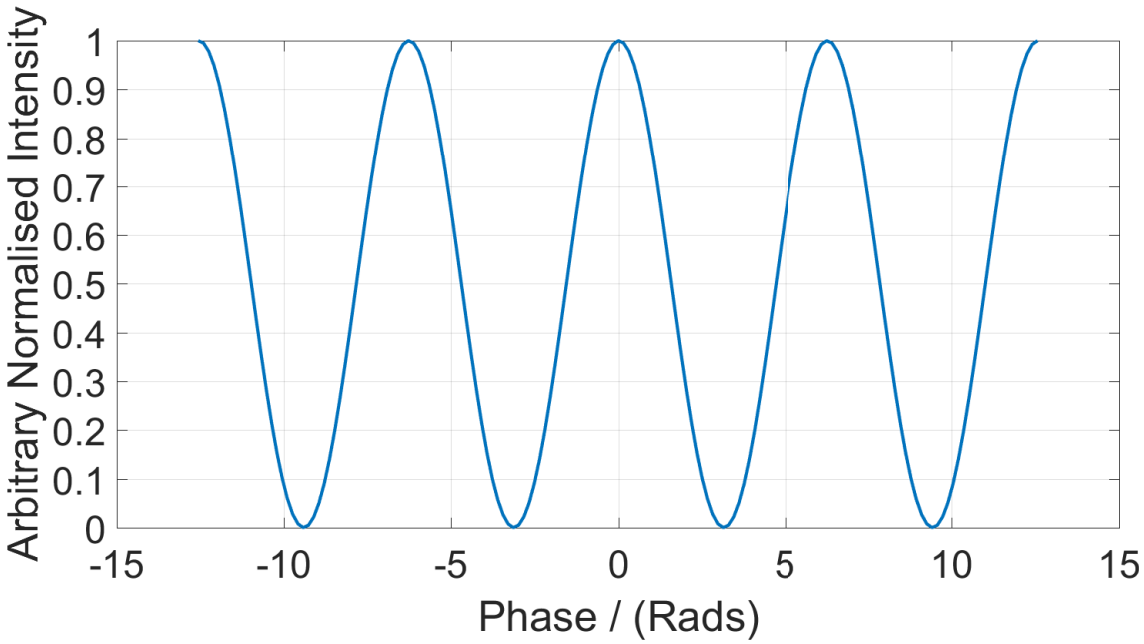


Figure 6.14: The visibility predicted by the split step Fourier method, after modelling two pulses amplified by the 70 μm and 85 μm PCFs in the high power regime (see Section 8.2). The degree of linear polarisation was predicted to be better than 99%, so despite the difference in lengths between the PCFs, excellent combination was expected.

To estimate the maximum amplification using mismatched PCFs, I set up the model to use input and output pulse energies of 3 µJ and 600 µJ respectively. These pulse energies were as reported in the coherent combination experiment, which used 8 PCFs to produce 12 mJ pulses [42]. Using the same experimental parameters as before (oscillator pulse duration, pulse stretching, spectral shaping), the chirped pulse from the seed laser had its energy increased to 6 µJ and was split equally between the two PCFs, before being amplified to 600 µJ. Once the pulse had been amplified, the non-linear phases predicted by the model for the 70 µm and 85 µm PCFs were 20.7 and 19.2 radians respectively. This was a significant amount of non-linear phase, which had an impact on the optimally compressed pulse autocorrelation, as seen in Figure 6.15. However, the combined pulse has still achieved a degree of linear polarisation of 92%.

The large amount of non-linear phase, which caused the simulated compressor difficulty in compressing the pulse was reduced to match the acquired non-linear phase of the 8 PCF

experiment (5 radians). The non-linear phase was reduced by decreasing the pulse energy acquired by the pulses in their respective PCFs. This energy reduction was done incrementally, until the non-linear phase acquired by each PCF was approximately 5 radians. The reduction to 5 radians was achieved by using an amplified pulse energy of 100 µJ. This non-linear phase reduction also led to the pulse autocorrelation better resembling a pulse, as shown in Figure 6.16. The non-linear phases, calculated at an amplified pulse energy of 100 µJ for the 70 µm and 85 µm PCF, were 5.1 and 4.7 radians respectively. The degree of linear polarisation and fraction of available beam power of the combined pulse at the reduced pulse energy, were both calculated to be over 99%.

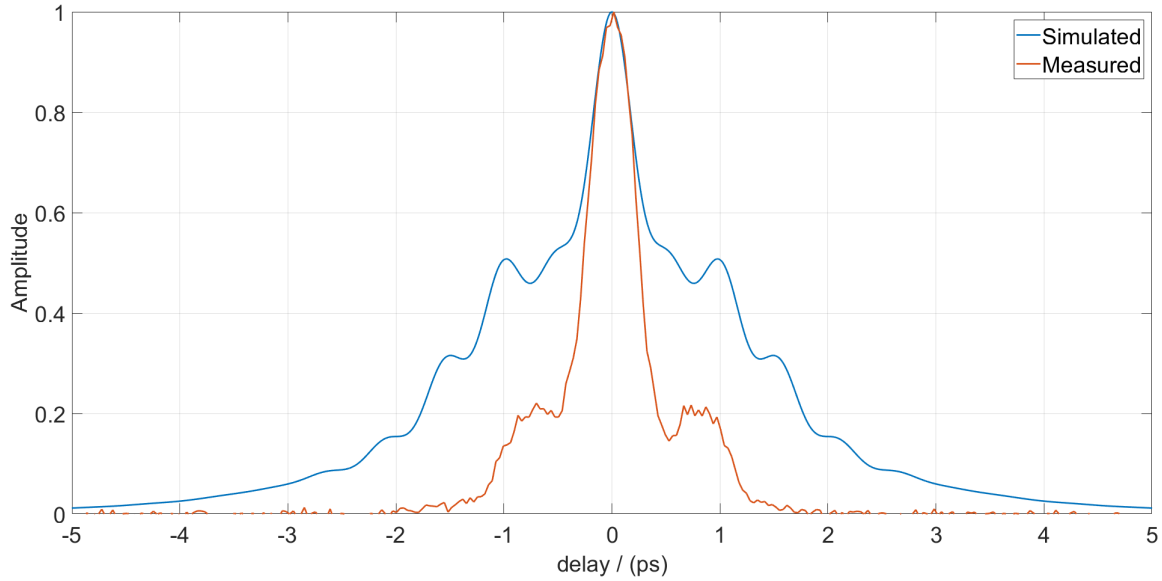


Figure 6.15: The predicted pulse autocorrelation after coherently combining two pulses with 600 µJ of energy. The non-linear phases calculated for the 70 µm and 85 µm PCFs were 20.7 and 19.2 radians respectively. Shown in red, is the the measured combined pulse autocorrelation, taken from the high-power regime experiment in this thesis (Section 8.2).

Assuming the pulse energy of 100 µJ per fibre and the combined pulse duration of 400 fs, the peak power possible for a 200 µJ pulse, with 75 % of the energy under the central peak, was calculated to be 375 MW for two PCFs.

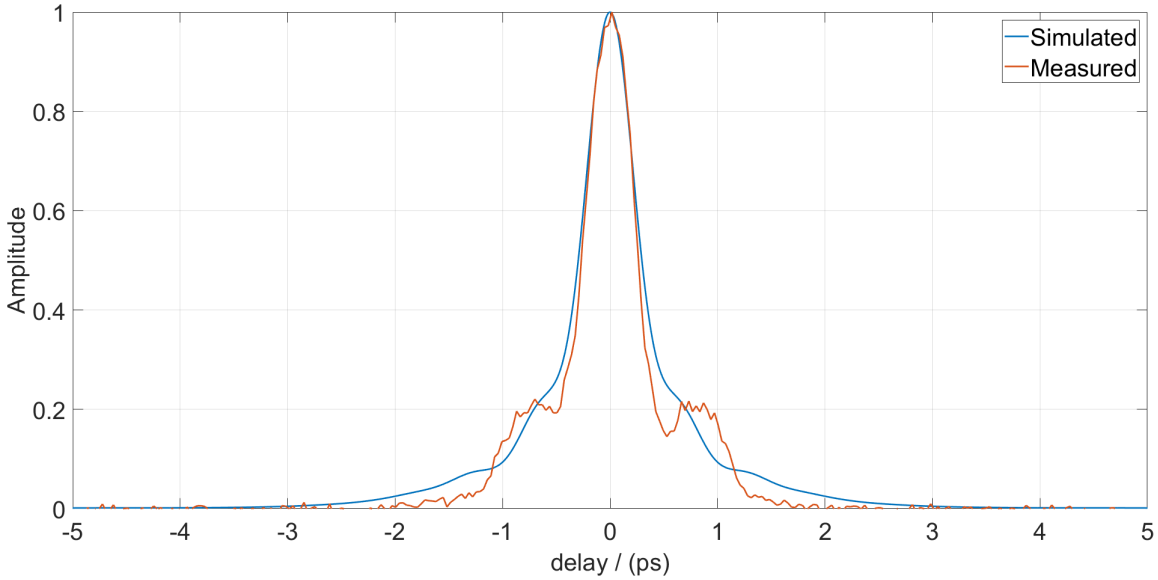


Figure 6.16: The predicted pulse autocorrelation after coherently combining two pulses with 100 µJ of energy. The non-linear phases calculated for the 70 µm and 85 µm PCF were 5.2 and 4.8 radians respectively. Additionally, the deconvolved pulse duration was calculated to be $\sim$ 400 fs. Shown in red, is the the measured combined pulse autocorrelation taken from the high-power regime experiment in this thesis (Section 8.2). The degree of linear polarisation and fraction of available beam power of the simulated combined pulse was calculated to be over 99%.

6.3 Summary

In this chapter, the component parts of the seed laser were described in detail. I was able to identify that the cause of the seed laser's broadened pulse-duration and compressed pulse shape was due to its spectrum. The impact of gain narrowing affected the width of the seed laser's spectrum and resulted in a broadening of the pulse duration from 200 fs to $\sim$ 370 fs. The effect of gain narrowing can be seen in the spectra of both fibre pre-amplifiers in figure 6.7. Self phase modulation was experienced by the pulse in the main amplifier, with a non-linear of phase of $\phi_{\mathrm{NL}} = 7$ being estimated. This led to fine structure appearing in the spectrum and the creation of side lobes in the autocorrelation, which I was able to partially confirm later, using the split step Fourier method model of the seed laser. After the pulse had exited the simulated seed laser, the simulated pulse's spectrum was masked with the shape of the measured spectrum of the seed laser, to compensate for the features the model lacked.

The simulated autocorrelation produced by the model showed good agreement to the measured autocorrelation in terms of its pulse duration, but there was a discrepancy between

the size of the side lobes, which remained unresolved. However, as the aim of the experiments in this thesis was to show that the outputs of two mismatched PCFs could be coherently combined, replicating the exact pulse shape of the seed laser was not necessary. The important point is that the seed laser, not the mismatched PCFs, was responsible for the fine structure and pulse broadening. Any additional effects seen in the pulse duration or pulse shape after amplification in the PCFs, could be attributed to the PCFs.

With some confidence in the split step Fourier method model of the seed laser section, I then extended the model to include the pulse propagation through the two mismatched PCFs. Using this extended model, shown schematically in figure 6.10, the degree of linear polarisation and fraction of available beam power of the combined pulse were predicted to be above 99 %. Further predictions were then made to identify how far the coherent combination of mismatched PCFs could be taken before the coherent combination would begin to suffer. Adjusting the pulse energy to achieve similar non-linear phase of 5 radians, as reported in the eight PCF experiment [42], the maximum pulse energy for each fibre was estimated to be 100 μJ. The degree of linear polarisation and fraction of available beam power of the 100 μJ simulated combined pulse, was calculated to be over 99%. Assuming a combined pulse duration of 400 fs, the peak power possible for a 200 μJ combined pulse, with 75 % of the energy under the central peak, was calculated to be 375 MW for two PCFs.

Chapter 7

Wavefront and Thermal Effects

7.1 Wavefront considerations

During the initial stages of setting up the combination experiment, plate polarising beam splitters (PBS) were used for the splitting and combining elements. With typical, combined degree of linear polarisation (DOLP) values of approximately 10% or less being common, it seemed to indicate that it was an unpolarised, combined beam. However, each of the fibre's output beams showed an average DOLP of 70% and 95% for the 70 μm PCF and 85 μm PCF respectively. It would have been assumed that the combined polarisation state of these two beams, should have been at least as good as the lowest DOLP figure.

The plate PBSs in the interferometer were identified as the source of this DOLP loss, so they were exchanged for two 10 mm PBS cubes. The first benefit of the PBS cubes was the ease of setup. The plate PBSs required an acute angle between the incident and reflected beams, and this complicated the layout of components on the optical bench. The 10 mm PBS cubes could be positioned to give a reflected beam at 90 degrees to the incident beam, making component layouts easier to design. Following this, the most time intensive parts of setting up the experiment were balancing the arms and searching for temporal overlap. Once temporal overlap had been found, the beams were spatially overlapped using an IR card in the near and far field, to attain co-propagating beams.

The second benefit of the PBS cubes was the improvement in the degree of linear polarisation. Swapping the plates for the cubes, increased the combined DOLP to 20%, which was

encouraging, but still not the 80%-90% that would be expected from combining individual beams with high polarisation purity. An IR card was used to inspect the combined beam behind the HWP and polariser, and this revealed some structure in the beam which did not show Gaussian-like beam characteristics. Therefore to investigate this effect further, I set up a series of diagnostics to monitor the different beam parameters. These diagnostics are described in Section 5.8. A beam profiling camera was part of these diagnostics, and was positioned after a PBS to inspect the transmitted beam profile of the combined beam. An example of the transmitted and rejected beams can be seen in Figure 7.1. From Figure 7.1,

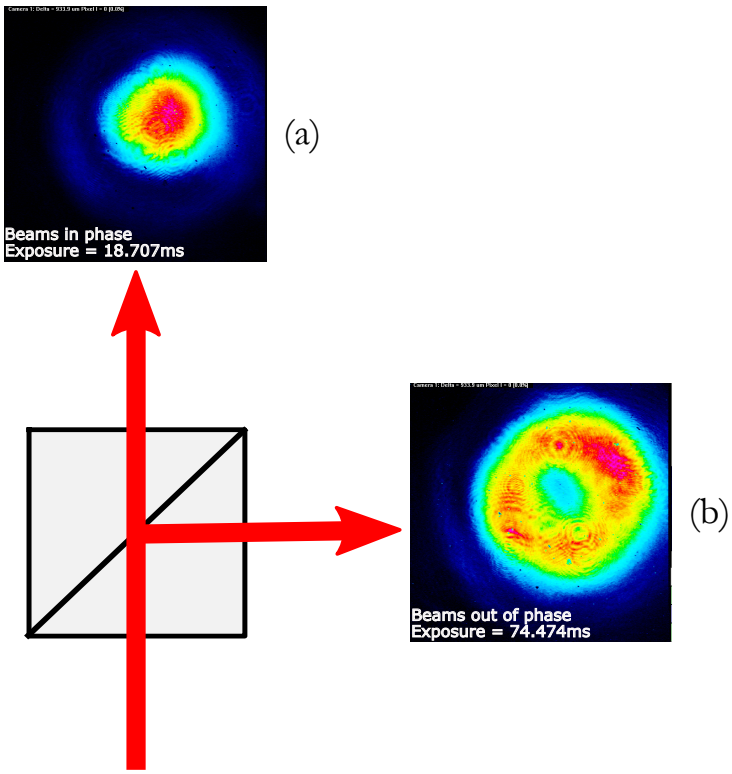


Figure 7.1: The combined beam, as seen through a PBS with a beam profiling camera. The images show the same combined beam observed in orthogonal (a) P-polarisation (transmitted) and (b) S-polarisation (reflected) states. The transmitted beam, shown in (a), is the useful beam that could be used for subsequent stages of combination. The rejected beam, shown in (b), is the power lost at the point of combination.

it can be seen that the intensity distribution was not random. Transmitted through the PBS, a beam with a central spot and a faint outer ring was observed. The reflected beam (from the PBS) had a defined ring structure, with a minimum in the centre. This ring pattern pulsated in and out of the centre, as the path length difference changed. A Half Wave Plate (HWP), placed before the PBS could be rotated, causing the beams to exit from different ports. This observation suggested that the ring structure was related to the polarisation state of the

combined beam. Interfering S- and P-polarised beams with different wavefronts, had created a spatially dependent polarisation state across the combined beam profile. These wavefront differences can be represented as small, spatially dependent phase differences across the beam profile, as shown in Figure 7.2.

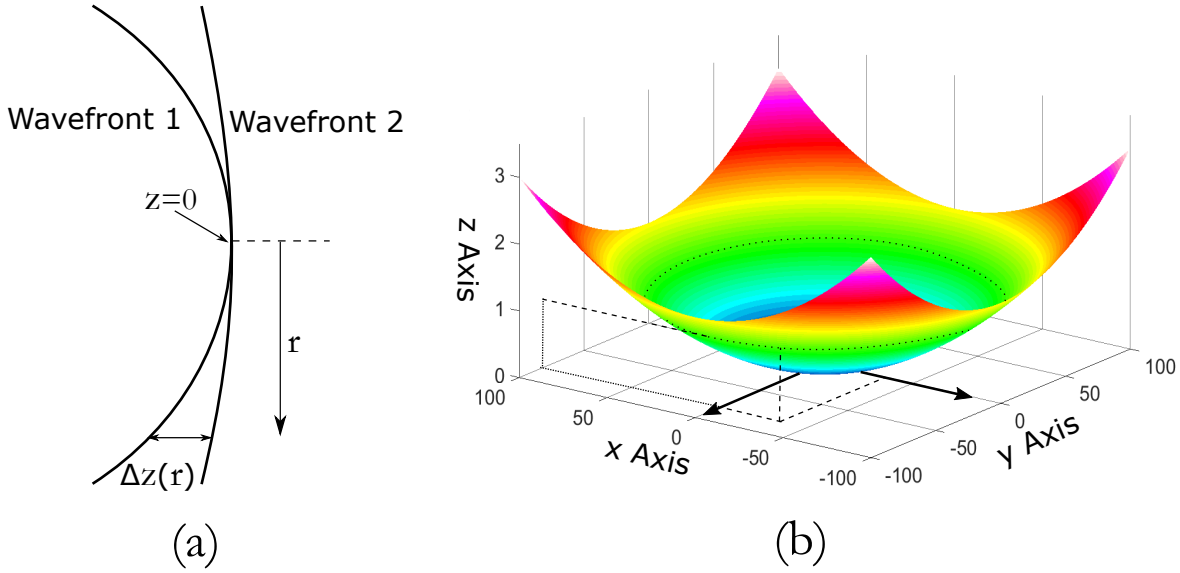


Figure 7.2: (a): A schematic of the wavefronts of two beams. Here, z is the the distance of the wavefront from the tangent plane, and r is the radial distance from the centre of the wavefront. When the wavefronts are temporally overlapped at their centres, there is a growing difference between the wavefronts, as one moves to larger values of r. (b): In 3D, a point on the wavefront can be defined in Cartesian coordinates.

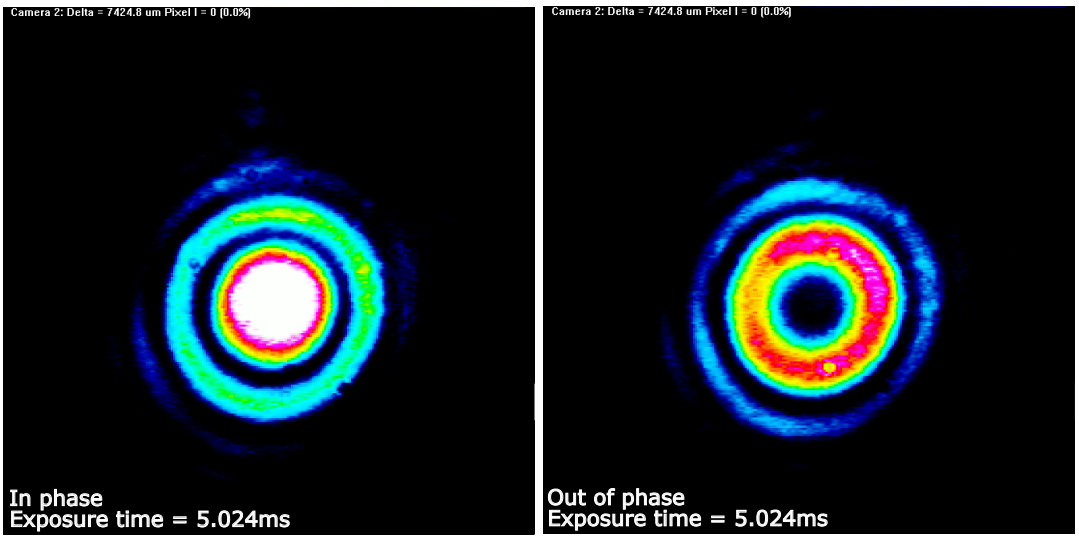
To see how detrimental the mismatching of wavefronts could be, I placed a $f = 200$ mm focal length lens into the same arm as the 85 μm PCF. This ensured that the wavefronts were vastly different, whilst keeping the beam diameters matched at the diagnostic PBS. With the beams placed in a co-propagating arrangement, the interference was as seen in Figure 7.3.

The difference in the wavefronts appeared to be causing detrimental effects on the polarisation state. To verify if this was true, I constructed a mathematical model using MATLAB to simulate beams with Gaussian intensity distributions, and wavefronts with different radii of curvature (ROC). A pair of electric fields with 2D Gaussian spatial beam profiles, $F(x, y)$, were assumed and multiplied with a phase

$$\mathrm{E_1(x, y) = F_1(x, y)e^{-i\phi_1}}, \tag{7.1}$$

and

$$\mathrm{E_2(x, y) = F_2(x, y)e^{-i\phi_2}}, \tag{7.2}$$



(a) (b)

Figure 7.3: The combined beam observed (as per Figure 7.1) when an $f = 200$ mm lens is used in one of the beam lines to exaggerate the difference between the wavefronts. Rings are clearly visible and the difference between the orthogonal polarisation states shows that the wavefront effect is affecting the polarisation. Shown is (a) a P-polarised beam and (b) an S-polarised beam

where x and y are the transverse coordinates of the spatial beam profile. The wavefront is defined as a spherical surface with a ROC R, and its tangent shifted to the (x,y) plane. The difference in path length is calculated by taking the distance between the wavefront and a tangent plane at the centre of the wavefront at z=0. z(x,y) at all points is calculated by

$$\mathrm{z(x,y)} = \sqrt{\mathrm{R^2 - x^2 - y^2}}, \tag{7.3}$$

which can then be used to find the path difference between the wavefront and the (x,y) plane at any point

$$\mathrm{P_d(x,y)} = \mathrm{R} - \sqrt{\mathrm{R^2 - x^2 - y^2}}. \tag{7.4}$$

To calculate the phase ϕ, the path difference is scaled by the laser wavelength λ and converted into radians

$$\phi(\mathrm{x,y}) = \frac{2\pi \mathrm{P_d(x,y)}}{\lambda}, \tag{7.5}$$

giving the phase of the wavefront relative the tangent plane. For beams with different ROC, each electric field should be multiplied by its characteristic phase, using Equation (7.2). To find the resulting intensity distribution, they are summed together and then multiplied by

their conjugate

$$I_{arb}(x, y) = \frac{n\epsilon_0 c}{2} * (\hat{E}_1 + \hat{E}_2) \times (\hat{E}_1 + \hat{E}_2)^*, \tag{7.6}$$

where the $\frac{1}{2}$ comes from assumption that E_1 and E_2 are S- and P-polarised electric fields and the $n\epsilon_0 c$ factor appears from Equation (4.26). Using the fraction of available beam power as a metric to assess the quality of the combination, it was calculated by

$$FABP = \frac{I_P}{I_P + I_S}, \tag{7.7}$$

where I_P is the intensity of the P-polarised light, which is useful and could go on to further stages of combination, and I_S is the power lost at this stage of combination. The P-polarised light was defined as the two combining beams being completely in phase, whereas the S-polarised light was defined as the combining beams being out of phase by π. Using two beams with 5mm beam diameters ($1/e^2$) and matched ROC of 2000 mm, the model was used to predict the fraction of available beam power (FABP), and the result is shown in Figure 7.4. Keeping all parameters the same, except for reducing the ROC of one beam's wavefront to 1500 mm (ROC of 2000 mm and 1500 mm), the ring structure could be seen starting to form in Figure 7.5. Finally, keeping all parameters the same, except for reducing the ROC of one beam's wavefront to 800 mm (2000mm and 800 mm), a significant number of rings could be seen forming in Figure 7.6.

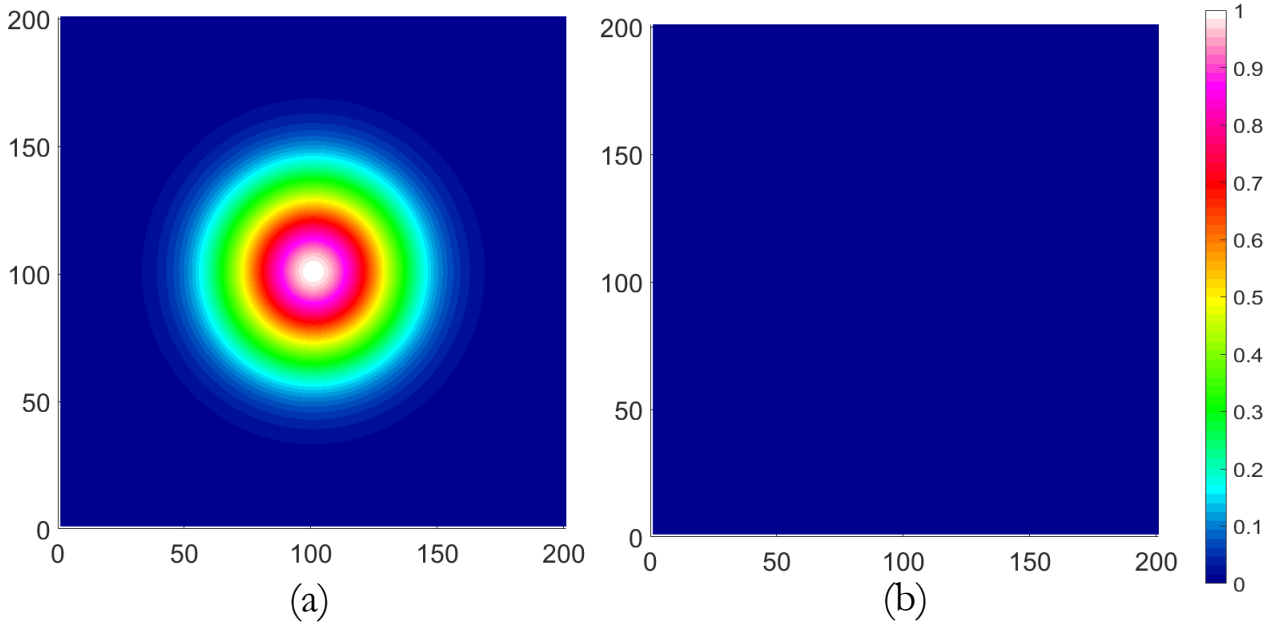


Figure 7.4: The combined beam profile that would be seen through a polariser, when two individual beams of 5 mm beam diameter, with wavefront ROCs of 2000 mm, and with perpendicular polarisation states are combined in phase. The outputs of the polariser are the (a) P-polarised transmitted beam, and (b) S-polarised rejected beam. The estimated FABP of this combined beam is 100%.

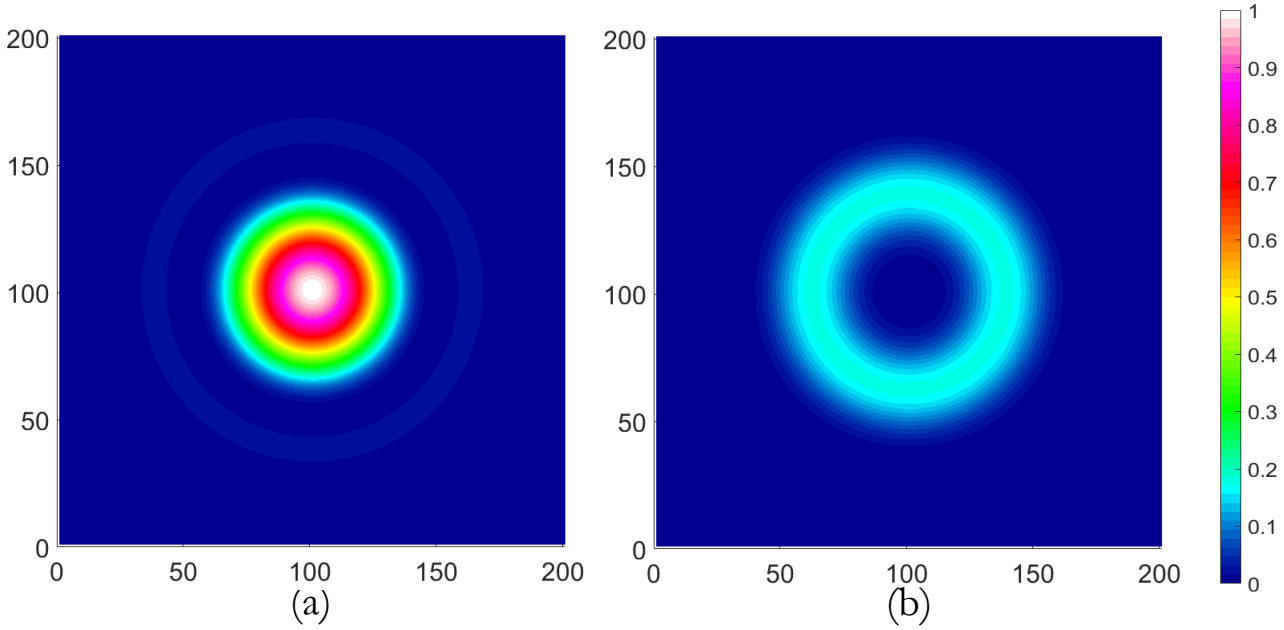


Figure 7.5: The combined beam profile that would be seen through a polariser, when two individual beams of 5 mm beam diameter, with wavefront ROCs of 1500 mm and 2000 mm, and with perpendicular polarisation states, are combined in phase. The outputs of the polariser are the (a) P-polarised transmitted beam and (b) S-polarised rejected beam. The estimated FABP of this combined beam is $\sim 64\%$.

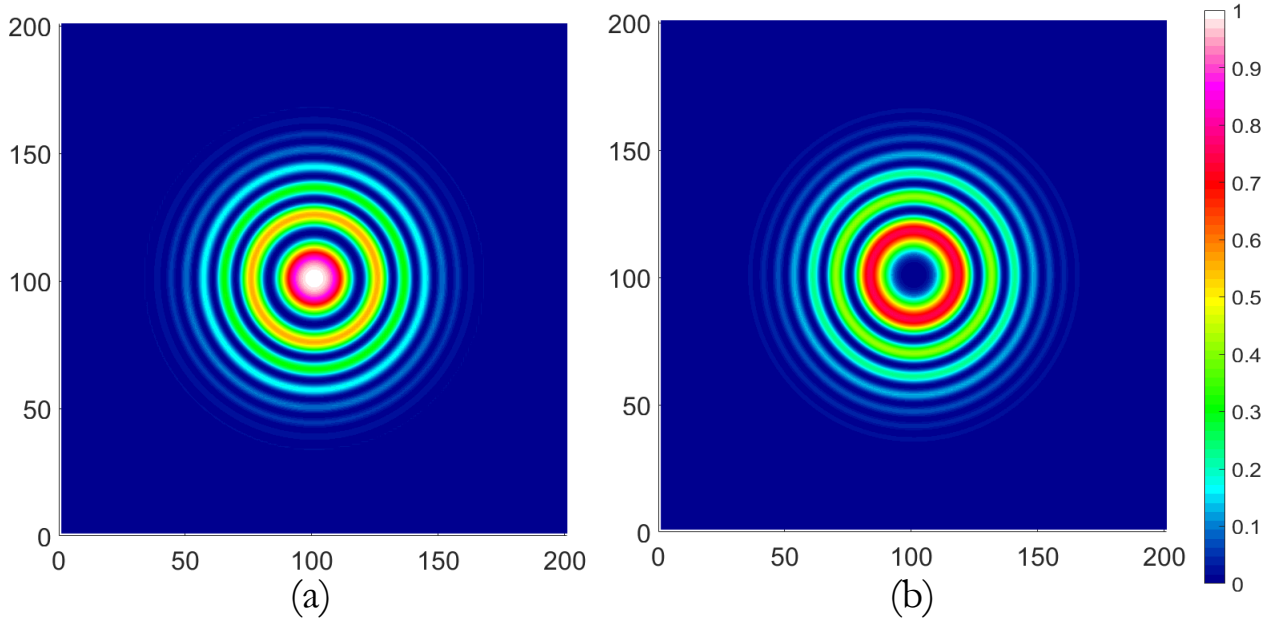
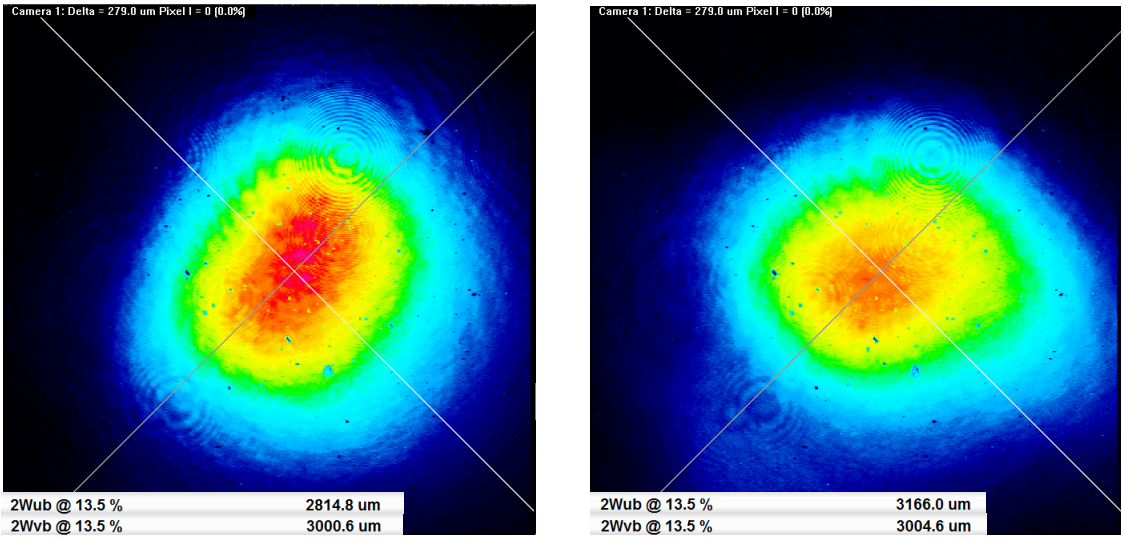


Figure 7.6: The combined beam profile that would be seen through a polariser, when two individual beams of 5 mm beam diameter, with wavefront ROCs of 800 mm and 2000 mm, and with perpendicular polarisation states, are combined in phase. The outputs of the polariser are the (a) P-polarised transmitted beam and (b) S-polarised rejected beam. The estimated FABP of this combined beam is $\sim 51\%$.

From Figure 7.6, the FABP was ~ 51 %, which seemed to indicate that this combined beam was no better than an unpolarised beam. Comparing with real data, Figure 7.7 presents the beam profiles that were recorded for the 70 µm PCF and the 85 µm PCF at the point of interaction. The beam diameters were extracted from the beams seen in Figure 7.7 using dataRay, a beam profiling software package, and used to model the Gaussian intensity profiles in the theoretical model. For the wavefront, I used a Thorlabs wavefront sensor to estimate the ROC of the beams, at a resolution of 28×35 points across the entire beam. The measured ROC were input as additional parameters into the theoretical model, in order to fully characterise the simulated beams. The resulting in-phase and out-of-phase beams can be seen in Figure 7.8.

Figures 7.8(a) to 7.8(d) showed good agreement for the modelled and measured intensity distributions. The in-phase images (a and c) both possess the faint outer ring around the central spot, and the out of phase images (b and d) both show a pale blue ring. In addition, the simulated and measured FABP values of $\sim 59\%$ and $\sim 61\%$ respectively, also showed good agreement. The model used the same colour map to replicate the beam profiles generated in



(a) (b)

Figure 7.7: An example of the beam profiles taken from each fibre and viewed through a PBS. A HWP has been used to maximise the signal transmitted through the PBS. (a) the 70 µm PCF. (b) the 85 µm PCF.

the DataRay software. The similarities between the theoretical model and the experimental data suggested that the wavefront differences between the two pulses could not be disregarded, as they would contribute to significant power loss in the rejected beam.

Another contribution to the degradation of the polarisation state was the effect of non-parallel beams. Beams with the same linear polarisation state and which are spatially overlapped on a CCD, but propagate at some angle relative to each other, will produce interference fringes [91]. Assuming plane wavefronts, the fringes will have a thickness inversely proportional to the size of the angle between them

$$\mathrm{d_f} = \frac{\mathrm{m}\lambda}{\sin\theta}, \tag{7.8}$$

where $\mathrm{d_f}$ is the fringe spacing and θ is the angle between the beams. When the phase of one of the beams changes with respect to the other, the fringes will be seen to move back and forth across a beam profiling camera. For combining beams with orthogonal polarisation states, a polariser is required with the transmission axis set at 45° to both beams. The transmitted combined beam will also manifest these fringes, as result of the polarisation state varying across the beam. The smaller the angle between the two beams, the larger the fringe spacing will be, with perfectly co-propagating beams providing an infinite fringe spacing. In this case, the combined beam would have a uniform polarisation state across the whole spatial profile,

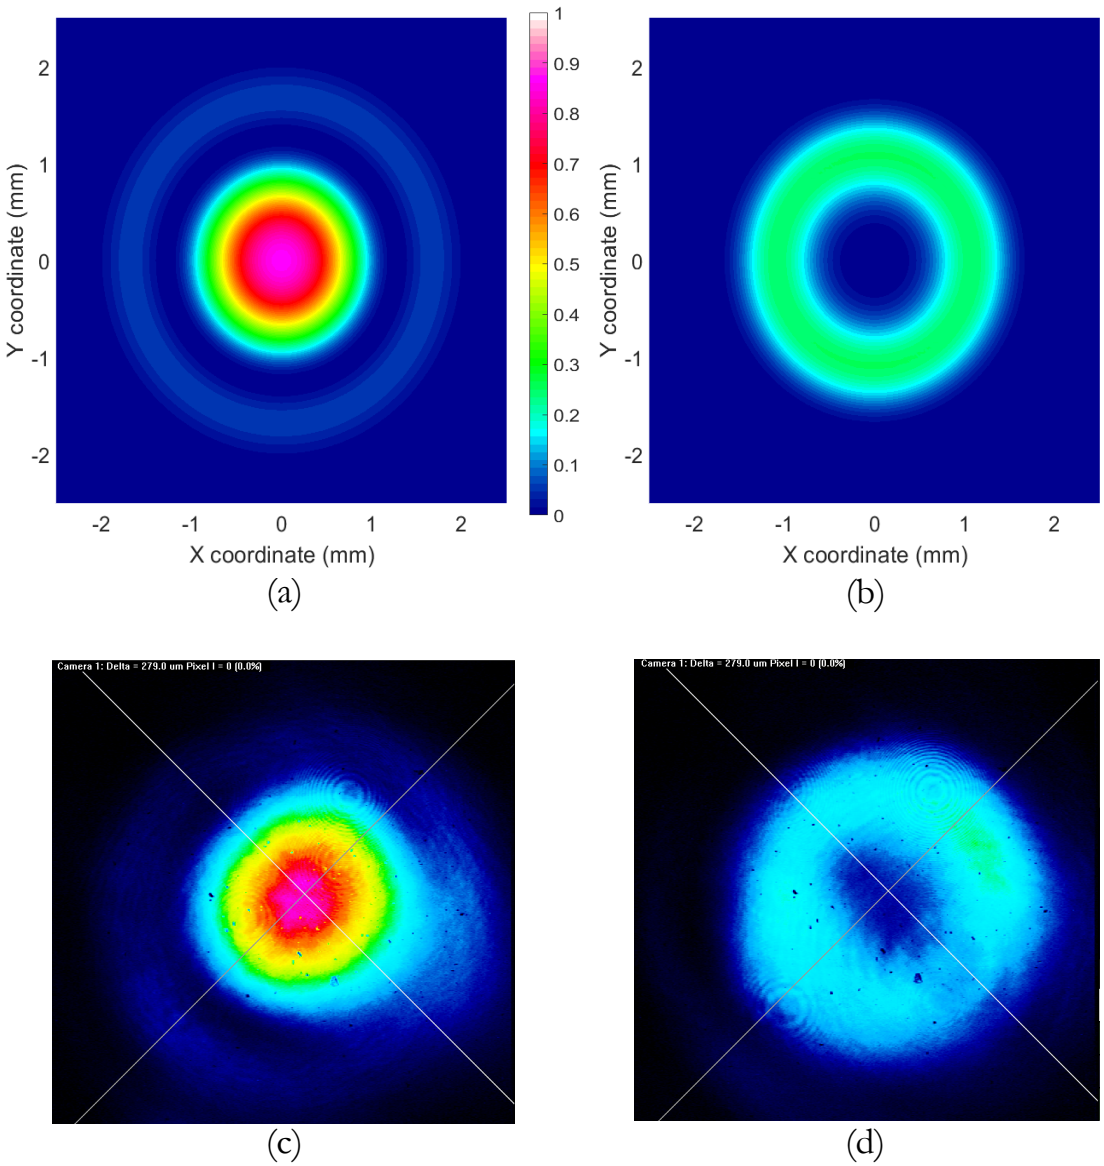


Figure 7.8: Using the measured beam diameters from Figure 7.7, and the measured wavefront ROCs, the interfering beams were modelled showing the (a) P-polarised transmitted beam and (b) S-polarised rejected beam. This gives an estimated FABP of $\sim 59\%$. The beams from Figure 7.7 interfering and viewed behind a HWP / PBS setup, gave a measured FABP of $\sim 61\%$. Shown is (c) the P-polarised transmitted beam and (d) the S-polarised rejected beam.

allowing for either complete transmission or reflection through a PBS. To reconstruct these observations, the model was extended to combine beams with plane wavefronts to interfere at a finite angle. Figure 7.9 shows two plane-wavefront beams combining initially at an angle with respect to each other, which was subsequently reduced to zero.

Figure 7.9 shows that any angle subtended between beams with flat wavefronts give parallel interference fringes. If I gave the beams mismatched radii of curvature, as well as an angular mismatch, the fringes exhibited curvature as seen in Figure 7.10. This could be seen experimentally also, as shown in Figure 7.11.

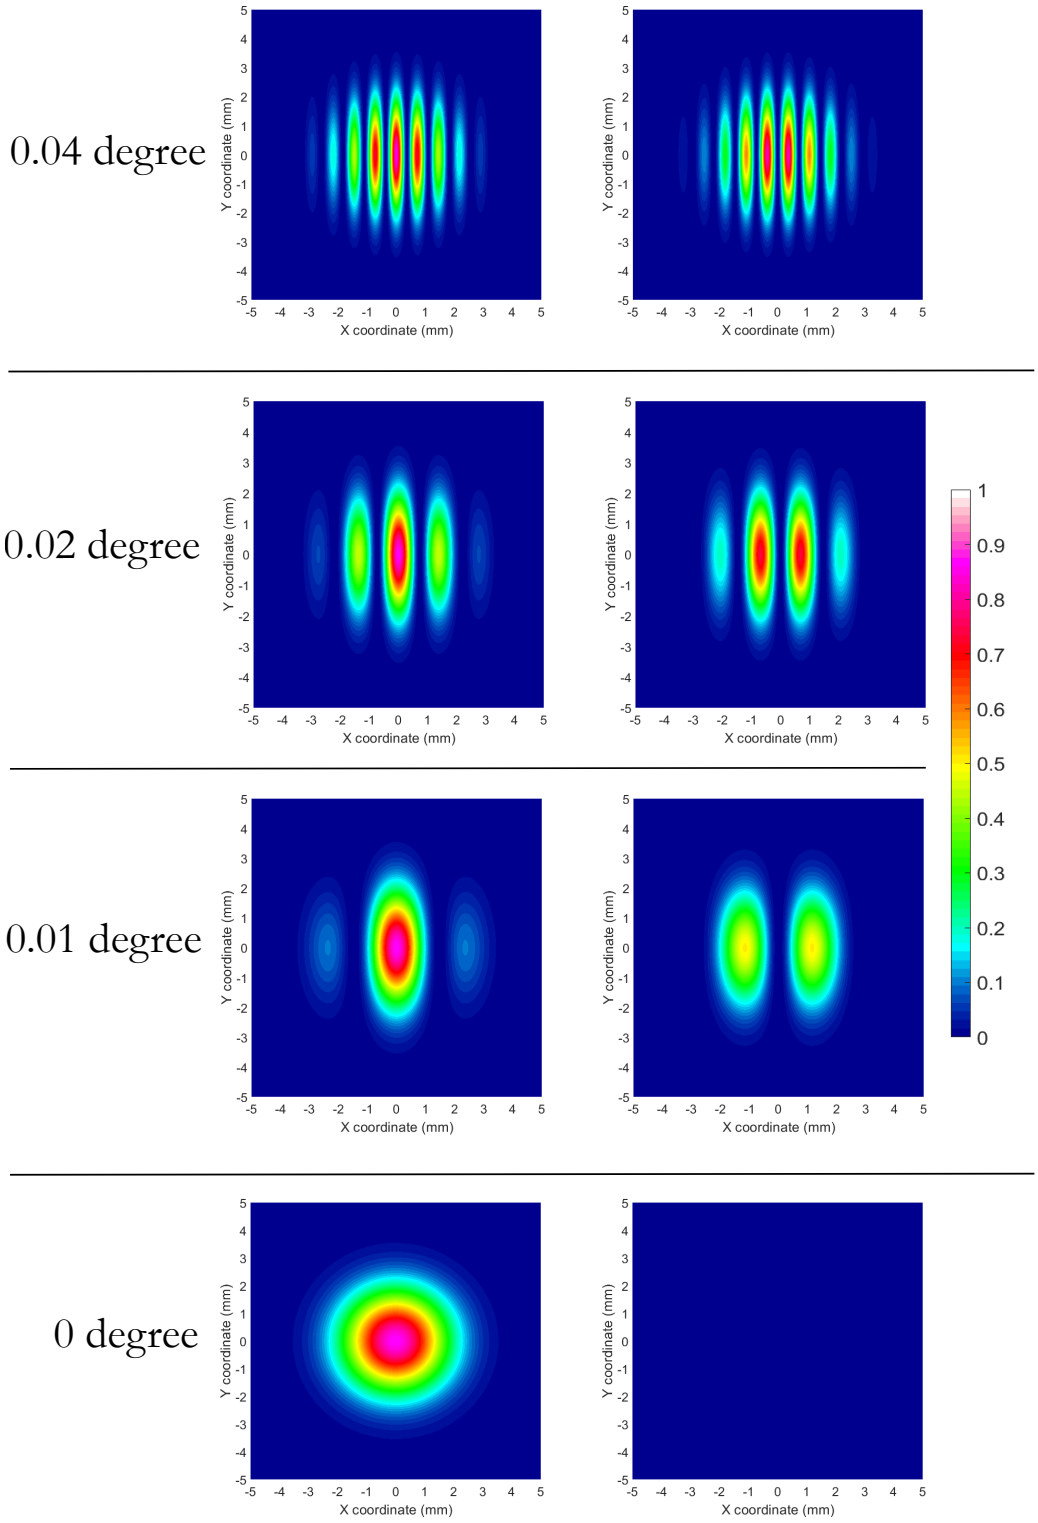


Figure 7.9: The effect seen if beams are spatially overlapped at the point of observation, but arrive at an angle relative to each other. The angle between the two beams (shown on the left) is large initially (top), but reduces to zero (bottom). (a) P-polarised beams. (b) S-polarised beams.

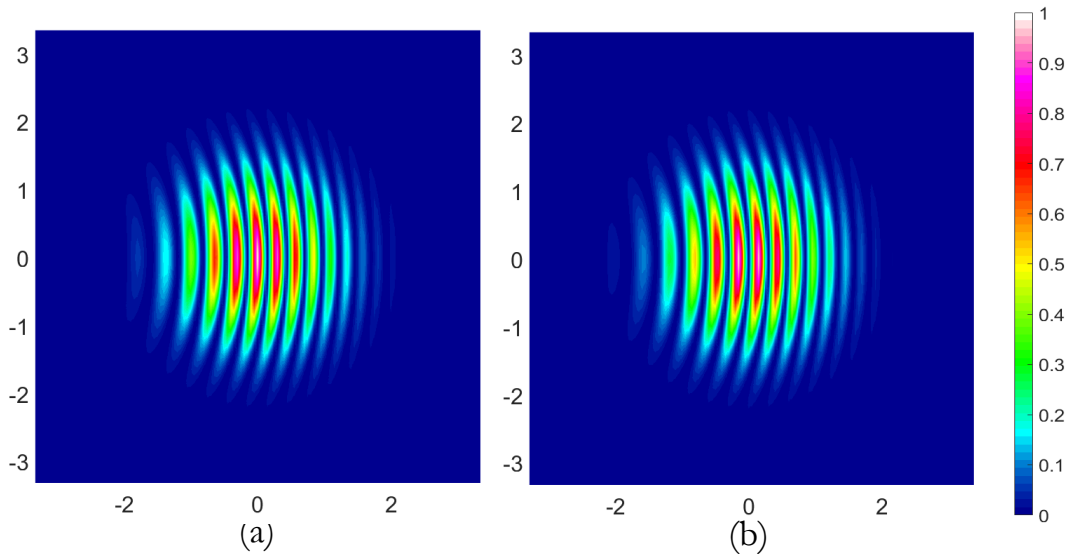


Figure 7.10: The effect of two beams, which were spatially overlapped at the point of observation, arriving at an angle with mismatched wavefronts. (a) A P-polarised beam. (b) An S-polarised beam.

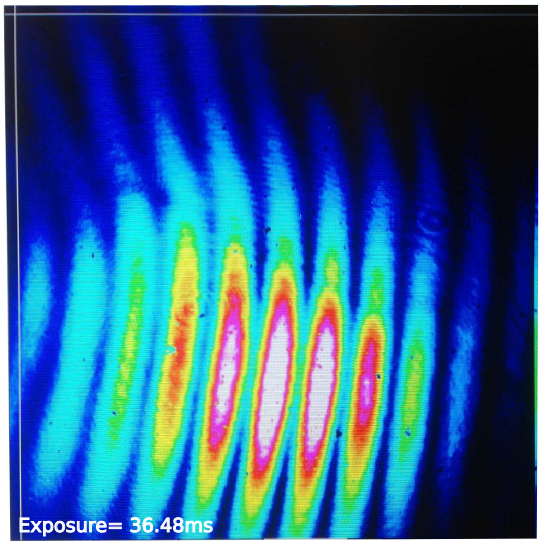


Figure 7.11: Misaligned combined beams observed on a camera through a PBS, when they were temporally overlapped, but not co-propagating. The fringes indicate that the pulses were interfering. Note the curvature appearing towards the edges.

A comparison between Figure 7.11 and 7.10 showed good agreement between the shape and intensity distribution of the fringes. The likely starting point for aligning a system would look like Figure 7.11, with mismatched wavefronts and non-co-propagating beams. As the angular mismatch between the beams was reduced, the fringes began to exhibit greater curvature as the wavefront mismatch effect began to dominate. The degree of wavefront mismatch determines the amount of curvature seen. As the angular mismatch tended to zero, giving co-propagating beams, the wavefront mismatch became the dominating factor

and interference rings were observed as seen in Figures 7.5-7.6. These interference fringes and rings are undesirable, as they reduce the quality of the combined beam. However, the interference fringes and their curvature could be used as an alignment tool, to help bring both beams into spatial overlap, and to assess if any wavefront mismatch is present. Consequently, I used this alignment tool for all of the combination experiments found in this thesis.

7.2 Beam divergence matching

The mismatch between the wavefronts of the two combining beams became a problem when trying to deliver combined beams with high degrees of linear polarisation and fractions of available beam power, to further stages of combination. I was able to address the wavefront mismatch problem by using a corrective telescope, calculated using beam propagation models written in MATLAB. The model was tested using the unamplified beams to gain confidence in the setup. For the experimental setup in this thesis, each beam exited their respective fibre with a different beam diameter, and was collimated by an aspheric lens. The imperfect collimation of the beams, and mismatched distances to the point of combination, contributed to the combining beam having different beam divergences. For review, each PCF had an aspheric lens to couple the pump into the inner cladding of the PCF. This same lens was used also to collimate the amplified seed beam exiting the core, as shown in Figure 7.12. Both PCFs had the same aspheric lens to ensure that they were treated the same, as would be seen in a large array of 1000s of PCFs. In addition, the aspheric lens used for the 70 μm PCF was moved closer to the PCF, to improve the beam profile as seen in Section 2.4.1. The combination of all of the factors mentioned above, meant that the PCF's beam diameters and divergences at the point of combination were far from matched.

Because of the dual functionality of the lens, I needed to optimise its position, both to collimate the amplified seed beam exiting the core, and also to give sufficient coupling of the pump to the inner cladding. The optimum position for collimating the amplified seed was not necessarily the optimum position for coupling the most amount of pump power to the inner cladding. For example, the lens position which best collimated the amplified seed beam,

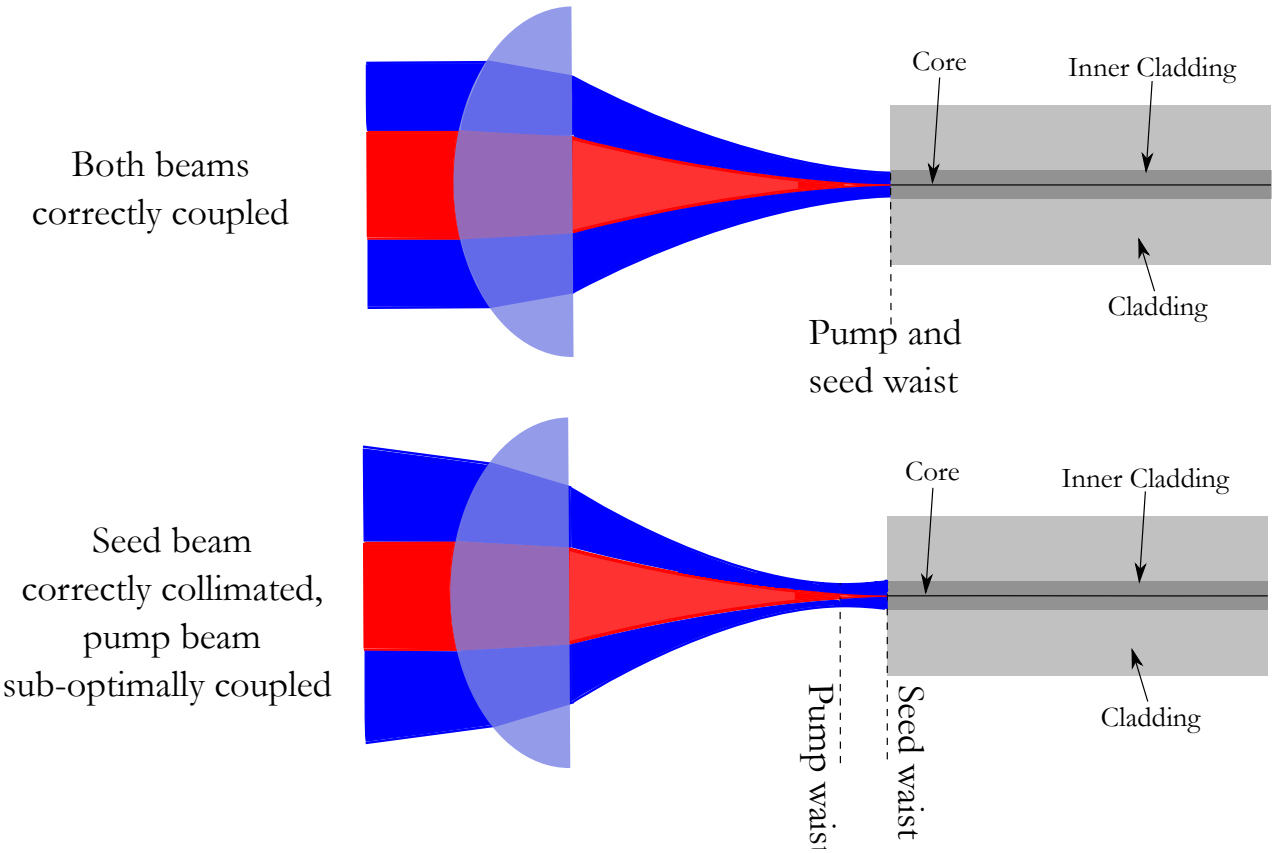


Figure 7.12: The lens used to couple the pump to the inner cladding, was also the lens used to collimate the amplified seed beam exiting the core. Given that the pump beam and amplified seed beams did not have matching beam sizes or divergences, the waists of the beams were in different longitudinal positions, making it difficult to optimise both the pump coupling and the amplified seed collimation.

may couple only 30 % of the pump power into the inner cladding. This made it difficult to choose a lens to fulfil both roles simultaneously. Therefore, I decided that collimating the amplified seed beam was the priority, to ensure good combination at the combining element. This meant that sub-optimal pump coupling was accepted, as it simply reduced the amount of pump absorbed for a given pump power. However, using the beam divergence matching techniques discussed later in this section, a telescope was added to adjust the pump beam divergence. I found that the quantity of pump power delivered to the inner cladding was increased, but there remained scope for improvement. Therefore, the issue of improving the pump coupling was seen as future work, and was not investigated further. To compensate for the differences in beam diameter and divergence of the amplified seed, a telescope was added to one of the arms of the interferometer. A telescope, with appropriately chosen focal length lenses, and positioned correctly would be able to match the beam sizes at the point of combination, as shown in Figure 7.13.

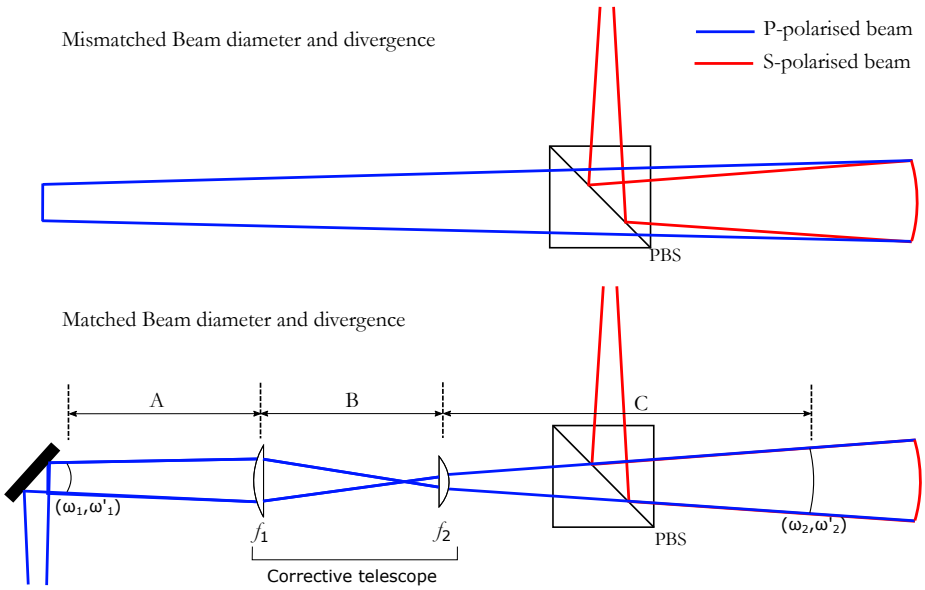


Figure 7.13: A diagram showing the principle of matching two beams with orthogonal polarisations. Initially, the S- and P-polarised beams had unmatched beam diameters and divergences at the point of combination. A telescope was added to the beam line of the P-polarised beam, to correct the beam diameter and adjust the beam divergence to match that of the other beam line. 'A' is the length from the first beam measurement point to the first telescope lens. 'B' is the length of the telescope. 'C' is the length from the end of the telescope to the second beam measurement point. (ω_1, ω_1') is the beam diameter and divergence prior to the telescope. (ω_2, ω_2') is the beam diameter and divergence measured after the PBS. The aim is to match (ω_2, ω_2') with the other beam's $(\omega_{70}, \omega_{70}')$.

In Figure 7.13, a telescope was used to change the beam diameter from one size to another, and to make the beam divergence difference as small as possible. However, this would not have guaranteed that the beam divergences would be matched, so the telescope length had to be adjusted to change the outgoing beam's divergence. Because this would change the beam diameter at the point of combination, I was required to calculate the telescope length, position in the beam line, and lens-focal-lengths, in order to match the combining beams. To achieve this, I wrote a beam propagation model in MATLAB, which used the measured beam diameters and divergences of the 70 µm and 85 µm PCFs, to optimise the telescope. The 70 µm PCF beam was measured at a set distance after the combining PBS, whilst the 85 µm PCF beam was measured as close as possible to the aspheric lens. Using ABCD matrices, the 85 µm PCF beam was propagated along the measured beam path, through the combining PBS, and to where the 70 µm PCF beam was measured, as shown in Figure 7.13. Here the diameters and divergences of the two beams were compared. Changes to the two telescope lenses, telescope position, and telescope length were made to minimise the residual between the two beam's diameters and divergences. With four free parameters to choose

from (two lenses, telescope length (B) and telescope position from the first mirror (A)), it was decided to fix the focal lengths of the lenses to values which were readily available. This reduced the free parameters to two. Various telescopes were trialled, including 1:1, 1.5:1 and 2:1 combinations. The virtual telescope was adjusted in length and position with different telescope lens combinations, until the beam diameter and divergence of the two combining beams was found to be similar. Once the telescope design and its position had been established, it was assembled on the optical table, as shown in Figure 7.14.

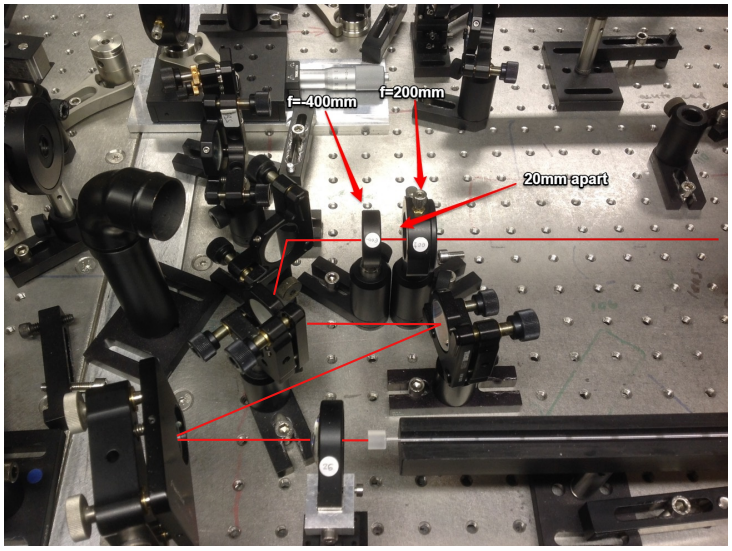


Figure 7.14: The telescope positioned in the beam line of the 85 µm PCF before the PBS. The lenses used were f=-400 mm and f=200 mm. Note that a Galilean telescope had been used.

With the telescope in place, the beams were spatially and temporally overlapped. The combined beam was viewed through a HWP and polariser by a beam profiling camera, as was done before. This was compared to the previous observation without a telescope, as shown in Figure 7.15.

The combined beam, from the bottom row of Figure 7.15, was found to be a large improvement over the combined beam, seen in the top row. Much more of the beam power was in the P-polarised beam, achieving a DOLP of $\sim 68\%$, compared to the previous $\sim 30\%$ DOLP measured for the top row. The wavefronts, which had been considered perfectly spherical up until now showed some structure in the S-polarised beam (Figure 7.15, bottom right). Subsequent adjustments to the pump coupling optics to optimise the amplified seed beam profile, forced the beam diameters to be remeasured and the wavefront correction optics

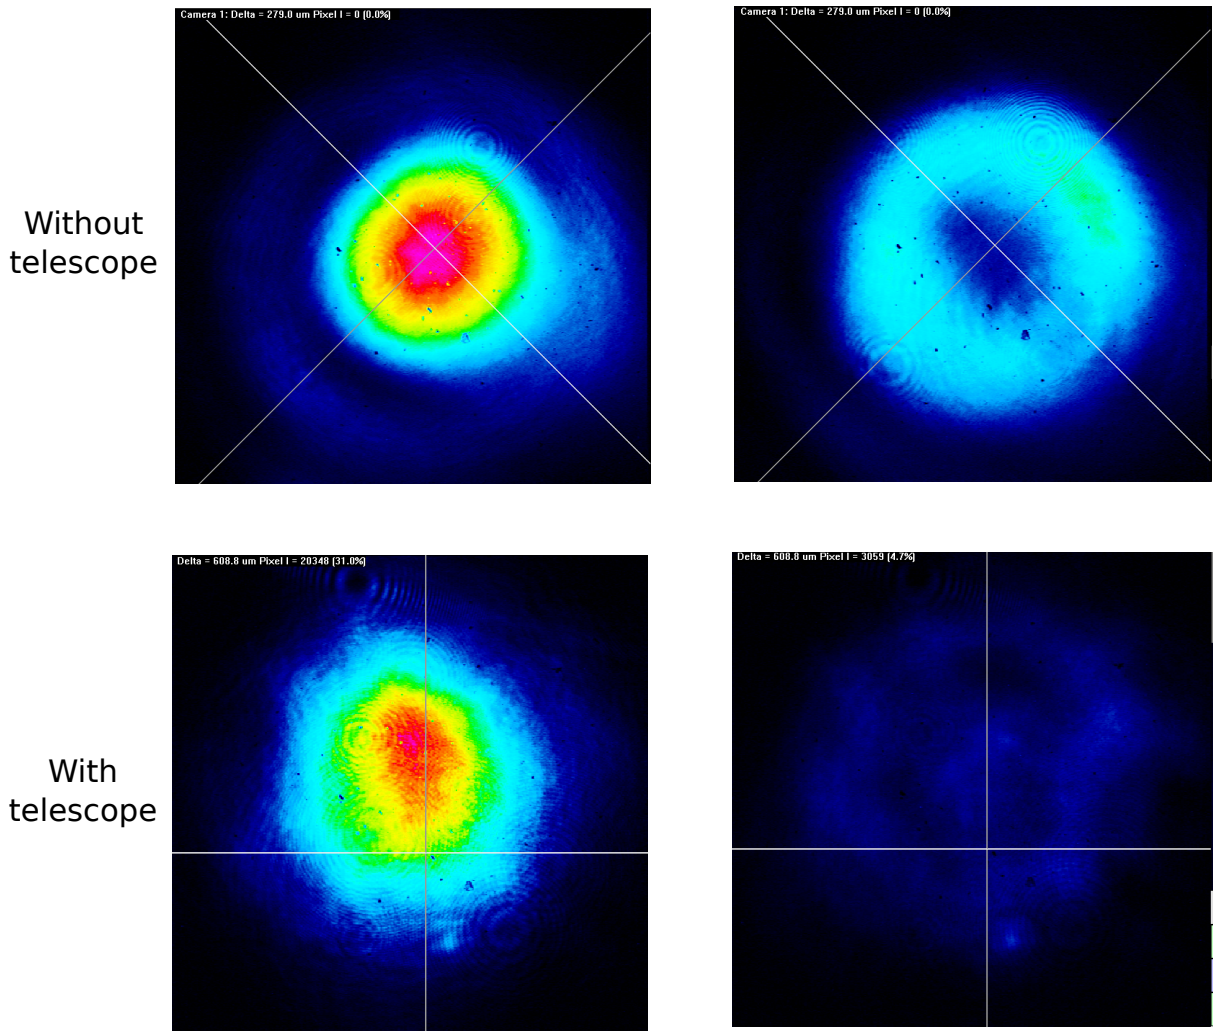


Figure 7.15: A comparison between the polarised beams transmitted, and rejected from the PBS. Top: The combined beam, when no corrections were made to the beam diameter or divergence. Bottom: The combined beam, resulting from the matching of the beam diameter and divergence of the individual beams and, using a specifically chosen telescope. Left: P-polarised beams. Right: S-polarised beams. The addition of a telescope improved the contrast between the desired and rejected beams. Ideally, the rejected beam image would be completely dark, meaning that all of the beam power was being directed into the combined beam.

to be recalculated. In this case, a single, long focal length was trialled, to ascertain if it was possible to use fewer components, and still produce a good match between the two beams. Fewer lenses meant a simpler setup compared to that using a telescope, as there were fewer components to position and align. Experimenting with overlapping the beams, I was able to increase the combined DOLP figure to $\sim 74\%$, as shown in Figure 7.16(c) and (d). The single lens approach created a simpler system, but I realised in hindsight that the telescope offered more flexibility in matching the beam sizes and divergences, at the point of combination. As such, subsequent wavefront correcting optics used a telescope. However, between the

two attempts I had made to match the beams, by inserting wavefront correcting optics into one of the beam lines, the combined beam DOLP figure had increased substantially. Given the improvement already seen in the unpumped regime, I decided that wavefront matching optimisation could continue to be tackled and refined while the system was pumped. This was due in part to the change in the beam diameter and divergence when the PCF became an amplifier, but was also due to the emergence of thermal effects. Thermal effects will be covered in the next section.

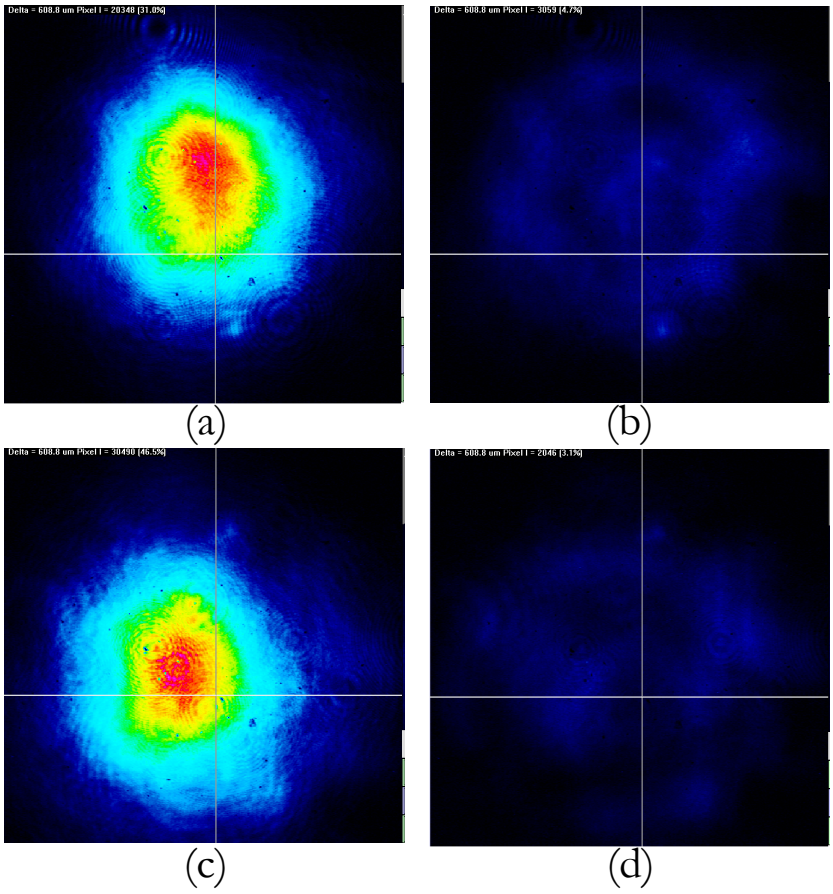


Figure 7.16: The combined beam before, and after the optimisation of the wavefront matching optics. (a) The combined beam in the desired polarisation state, using a telescope. (b) The combined rejected beam, using a telescope. (c) The combined beam in the desired polarisation state, using a single lens. (d) The combined rejected beam, using a single lens. The DOLP of the combined beam for the telescope and single lens setup was 68 % and 74 % respectively. The small increase in the DOLP was due to less beam power being lost to the rejected beam.

7.3 Thermal effects

The first major objective of the experiment was to ascertain if the two mismatched PCFs could be coherently combined. In an unamplified state, the beams had been coherently combined to give promising figures for the combined DOLP. The experiment had been run in an unamplified state, while I worked out the logistics of combining the different PCFs. This included building two models; one to ascertain if the material dispersion and non-linear phase picked up by each pulse, would create a barrier to pulse combination (as seen in Section 6.2); and a wavefront matching model, to ascertain why the combination was giving poor DOLP figures. Also, a new compressor and autocorrelator was built, to allow better control over pulse compression and measurement than the previous equipment offered.

Initially, the same setup that had been used in the unamplified experiment, I also used for the pumped experiment. The long focal length lens previously used for wavefront matching was not changed. To investigate the system under pumped conditions, I launched 15 W of pump power into each fibre, to produce two amplified seed beams. After some further adjustment to the pump power of the 70 μm PCF to match the PCF output powers, the combined output power was measured ∼4.5 W. However, the DOLP for this combination, was found to be poor, ∼ 50%. Inspection of the combined output showed that a faint ring structure had formed in the rejected beam, caused by the pumped fibres producing beams with different diameters and divergences, compared to the unpumped case. To diagnose the problem, I measured the beam diameter of the 70 μm PCF on a beam profiling camera at a distance of ∼ 1.9 m, and under different pump powers. The results of this are shown in Figure 7.17.

The method used to combine the beams, despite the mismatch in the wavefronts, was to match the beam diameter and divergence at the same point after the PBS. Figure 7.17 suggests that, if the pump power levels were changed after the system had stabilised, then the beam diameter and divergence at the point of combination, would also change. This effect has been reported before [122], where a change in the mode field diameter was witnessed as the pump power increased. The result here of both PCFs changing their beam diameters,

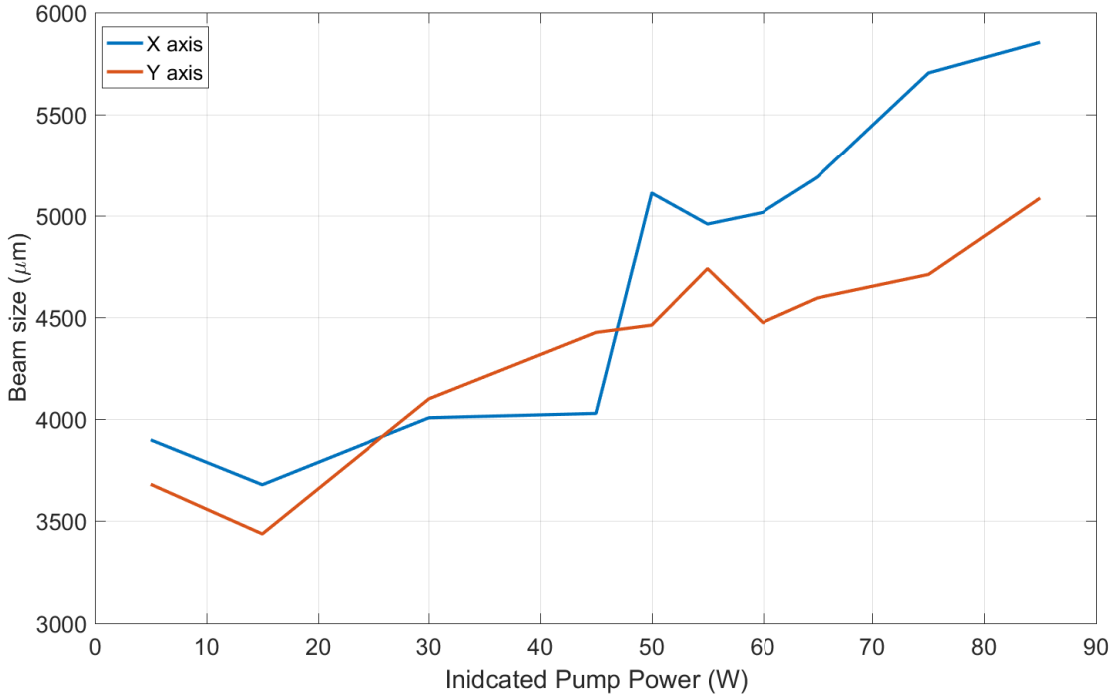


Figure 7.17: The change in beam diameter at 1.9 m from the 70 μm PCF, when seeded with 250 mW, and with increased pump power.

was a drop in the DOLP of the combined beam, as the mismatched wavefronts produced interference rings. Rather than establish the experiment at low power and then ramp up the pump power, I established the experiment at a set power and left it to thermalise for an hour. The effects of thermalisation of the PCFs were recognised earlier in previous single PCF pumping experiments, where I found that the beam pointing would wander after the PCF was pumped. This effect is shown in Figure 7.18. Once the PCFs had been pumped at a chosen power and left to thermalise, I remeasured the beam diameter and divergences for each PCF, and the wavefront correcting telescope was recalculated to match the two beams.

The heat produced from the non-zero quantum defect between the absorption and emission wavelengths [124], was responsible for the beam wandering as seen in Figure 7.18. The heat that was transferred from the core into the inner and outer cladding of the PCF, was transferred to the aluminium V grooves, in which they sat. The heating up and expansion of the silica fibre and metal V groove caused the beam to wander, as shown in Figures 7.18 and 7.19. The V grooves would heat up after 1/2 hour of operation, depending on the level of launched pump power. Once the system had thermalised at a given pump power, the PCF's output beam remained stable. As the system operated with passive cooling, it needed to be tailored to operate with static pump powers, as opposed to dynamically changing pump

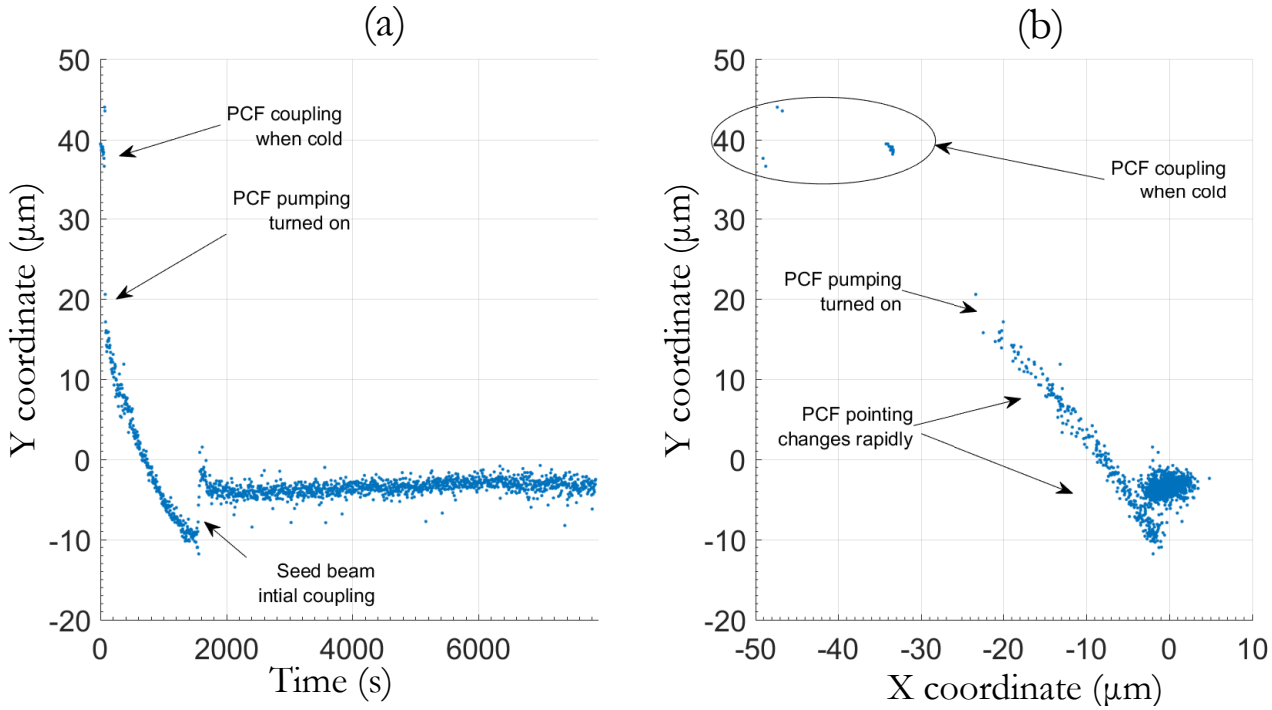


Figure 7.18: Plots of the relative coordinates of the beam centroid. (a) The movement of the beam centroid's Y coordinate as a function of time, as a result of heating due to pumping. (b) The position of the beam centroid over time in both X and Y coordinates. When the pump was turned on initially, there was a jump of the centroid position from the region of (-40,40) to the region of (-25,20); this is related to the location of gain in the PCF. As time evolved, the beam centroid position moved whilst its intensity profile degraded from a Gaussian-like mode. This was due to the seed beam becoming slightly decoupled from the core, as the PCF heated up. The discontinuity, shown at $\sim$ 1800 seconds, was due to a realignment of the seed beam to the core, which improved the beam profile after the thermal drift, to give a Gaussian-like mode. Once the system had thermalised, the centroid position remained stable from 1800 seconds onwards. The centroid is defined as the intensity weighted, arithmetic mean position of all pixel intensities above 13.5% of the peak intensity [123]. Using this method to determine the centroid, the beams position could be interpolated to an accuracy of 0.1 µm, despite the pixel size on the camera being 4.65 µm.

powers. The seed coupling can be optimised only when the system has thermalised to a given pump power. Changing the pump power level could make the system re-thermalise, so that the seed coupling to the PCF became sub-optimal.

Hence, I made decision to work in two different power regimes: a low power regime where the pump power was sufficient to get over 10x amplification of the seed, and to confirm that the two mismatched PCFs could be combined; and a high power regime, where the mismatched PCFs were pumped much harder and nearer to the limit of the available pump power. When the power regime was changed, the system was allowed to re-thermalise, before the seed coupling was optimised. The beam diameters and divergences were then remeasured and the coherent combination performed.

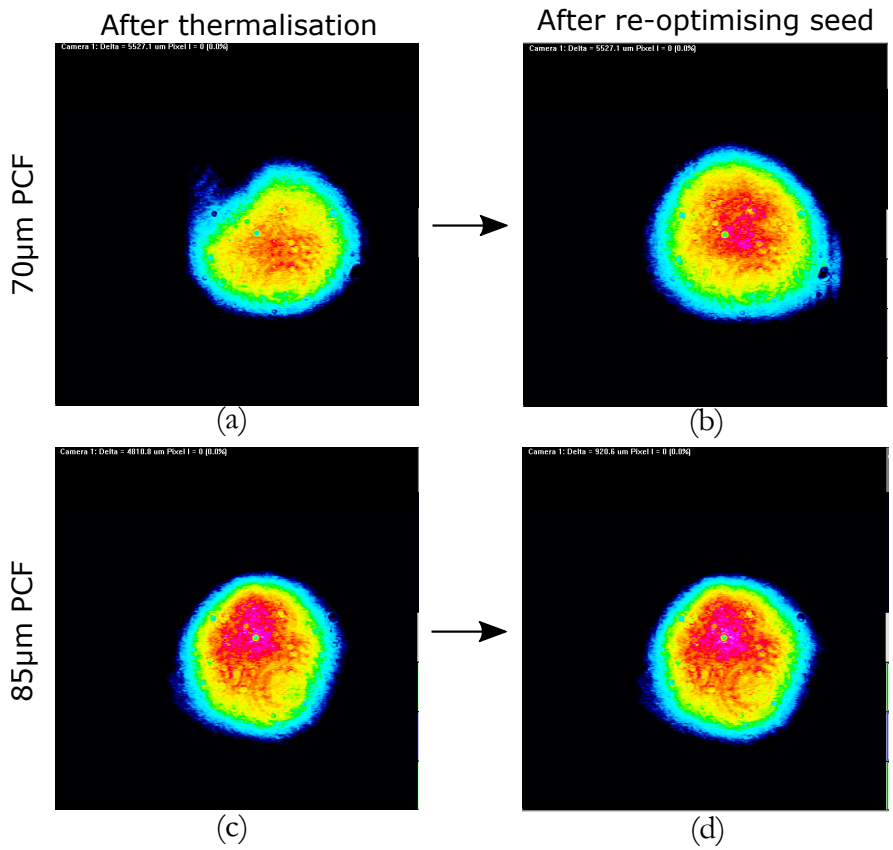


Figure 7.19: (a): The beam of the 70 µm PCF after thermalisation. (b): The beam of the 70 µm PCF after thermalisation and small corrections to the seed coupling optics. (c): The beam of the 85 µm PCF after thermalisation. (d): The beam of the 85 µm PCF after thermalisation and after a scan was performed to identify the best beam profile. The PCFs after being pumped for an hour at pump powers of ∼ 70 W were subject to thermal effects which affected the seed coupling to the PCF. It can be seen that the 70 µm PCF was more sensitive to this, with its Gaussian-like beam profile deformed due to the seed beam becoming slightly decoupled from the core as seen in (a). Small corrections to the seed beam optics could correct for this deformation, producing a uniform Gaussian-like beam seen in (b). In (b) and (c), the 85 µm amplified beam showed no obvious beam degradation or shifting.

The seed power used for this experiment was chosen to give the best contrast between the beam coupled to the core and that lost into the inner cladding. It was preferable to use the highest seed power possible, to extract the most energy from the PCF. However, at seed powers of ∼2.5 W, I noticed there was a significant amount of power leaking into the inner cladding of the 70 µm PCF, as shown in Figure 7.20.

This leakage affected the PCF's DOLP and also contributed to power loss at the point of combination. Reducing the input seed power to ∼ 0.250 W improved the contrast between the core and inner cladding, and this contrast was maintained during pumping. When the pumping was increased with 0.25 W of seed power, it was found that even at high pump

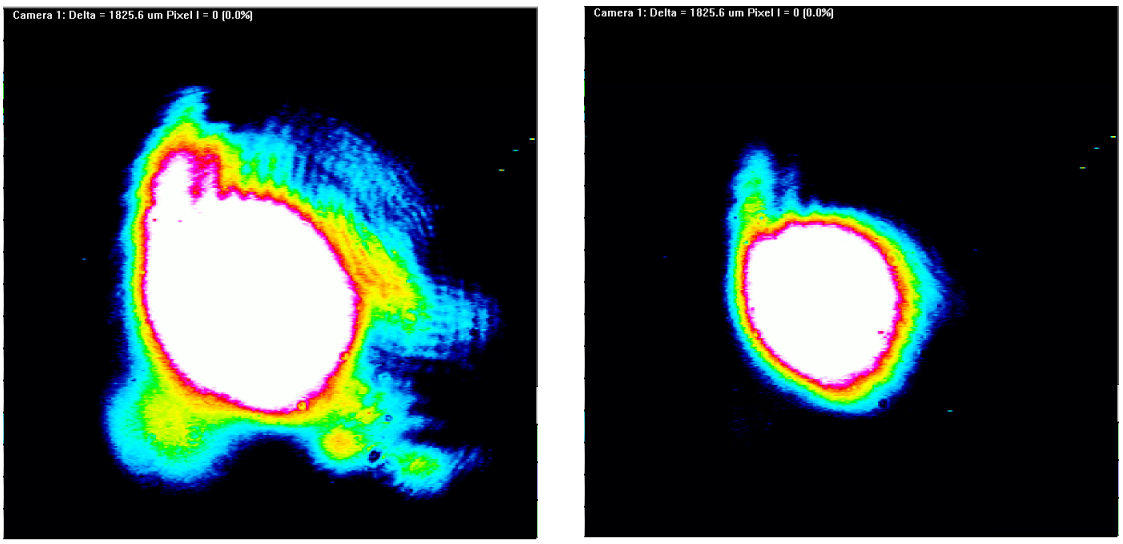


Figure 7.20: An overexposed image of the beam, highlighting the power lost around the core. Reducing the seed power, for the same level of pumping, from (a) ∼2.5 W to (b) ∼0.25 W, reduced the amount of power leaking into the inner cladding, which gave a much cleaner beam profile.

powers, the 70 µm PCF could produce a good quality beam, as shown in Figure 8.6 in Chapter 8. One thought, was that the seed beam was not perfectly coupled to the core of the PCF as this was not trivial to do. Therefore, some of the beam power that leaked into the inner cladding, propagated down the fibre, passed through the core region, and was amplified in the process. By reducing the initial seed power, the amount that escaped into the cladding could be reduced, therefore, improving the contrast as seen in Figure 7.20. With the seed power set to $\sim$ 0.250 W for the 70 µm PCF, the 85 µm PCF was also seeded with $\sim$ 0.250 W, to ensure parity of treatment between the two PCFs.

Initially, it was decided to deliver the same amount of pump power to each PCF. This parity of treatment was to replicate the PCFs in a large array of PCFs being seeded with the same seed power and pumped at the same pump power by identical pump lasers, a scenario that might be found in a coherent combination experiment of N PCFs. The PCFs were allowed to thermalise over an hour, before any amplified seed power measurements were taken. When the amplified seed powers were measured, a large difference was observed between the output powers of each PCF ($\sim$ 1.4 W and $\sim$ 3.4 W for the 70 µm and 85 µm PCF respectively), which grew when measured after the PBS (∼0.9 W and $\sim$ 2.7 W for the 70 µm and 85 µm PCF respectively). Some of this could be accounted for by the difference

in volume of active media of the two PCFs, with the ratio of their core volumes being a factor of ∼2. Other losses were due to: waveplates which weren't anti-reflection coated for the correct wavelength; and amplified unpolarised light from the 70 µm PCF being extracted by the PBS. With a three times difference in average power at the point of combination, I decided to better match the output powers of the PCFs. This would allow the contribution from each PCF to the final combined beam to be roughly equal, preventing the majority of the combined beam power from originating from a single PCF. The pump power launched into the 70 µm PCF was doubled, and varied to give an increased PCF output power of ∼3.39 W. This was very close to that of the 85 µm PCF. The pump laser for the 70 µm PCF was digitally controlled, meaning it was internally limited to finite steps in power due to the use of digital-to-analog converter. The current driver feeding the pump laser for the 85 µm PCF was an analog system, so a knob could be turned to vary the current driving the pump diode. Once the 70 µm PCF pump was set, the pump power of the 85 µm PCF was adjusted to match the same amplified seed output power.

Once the PCF output powers had been matched, I considered the difference in non-linear phase between the PCFs. The non-linear phase for each PCF was estimated to be 0.0355 and 0.0326 for the 70 µm and 85 µm PCFs respectively. The non-linear phase difference between the PCFs, $\Delta B = 0.003$, was deemed to have a negligible effect on the coherent combination. For the wavefront matching, the beam diameter and divergence for each beam line was measured for both PCFs, and the corrective telescope requirements were calculated. A 1:4 telescope was constructed, using a 50 mm and 200 mm lens placed 250 mm apart in the beamline for the 85 µm PCF. After running the wavefront matching model to find a single lens solution, it was found that a telescope with two lenses gave an extra parameter to vary in terms of the telescope length. This extra degree of freedom made it easier to find workable solutions. Effort was invested into the spatial and temporal overlap of the beams, to try to improve the DOLP of the combined beam. Micrometre to millimetre scale adjustments made to the telescope length allowed the beam divergence of the 85 µm PCF to be varied slightly. Making a small change to the telescope length, and then optimising the spatial and temporal overlap, allowed the combined beam's DOLP to be improved over the previous record of

$\sim$ 73% to 83.3%. The results of the combined PCFs, operating in the low power regime, were recorded from this point onwards and are presented in Chapter 8.

7.4 Summary

In this chapter, the issue preventing the combined beam from producing high degrees of linear polarisation, was identified as being due to wavefront mismatch. A mathematical model was produced to simulate the wavefront mismatch between two beams, and to estimate the impact of the wavefront mismatch on the degree of linear polarisation. A telescope was identified as a suitable method of matching the wavefronts at the point of combination and another mathematical model was produced. This model used the beam sizes of the beams from both PCFs at different places in their beam lines, in order to estimate the beam divergence. With the beam size and divergence for each beam, the model then calculated what telescope was needed in the 85 μm PCF's beam line, to ensure the beam size and divergence of both beams were matched at the point on combination. The successful telescope design was built and the effectiveness was analysed by measuring the degree of linear polarisation of the combined beam before and after the addition of the telescope. The degree of linear polarisation was found to improve significantly, increasing from $\sim$ 30% to $\sim$ 74 %. This increase signalled that the most dominant effect reducing the degree of linear polarisation of the combined beam was the mismatch between the wavefronts.

The thermal effects which could affect the performance of the photonic crystal fibres were also considered. The thermal issues identified were; the change in the beam pointing and spatial profile as the PCFs thermalised; and the change in beam size as more pump power was launched into the PCF. The beam pointing and spatial profile measurements dictated that the PCFs should be given time to thermalise before any work was done with the beams. Else, beams aligned to have good spatial overlap initially, could later see that overlap deteriorate. The change in beam size due to the pump power required that the pump power was chosen a-priori and left fixed during the experiment. This led to the creation of two power regimes; a low power regime and a high power regime. The low power regime was chosen to demonstrate coherent combination of the mismatched PCFs at lower power levels with 10x amplification. The high power regime was chosen to demonstrate coherent combination of mismatched PCFs whilst they were pumped at the maximum possible pump power.

Chapter 8

Results

In this chapter, I present the results for the coherent combination of two mismatched photonic crystal fibres (PCFs). I ran the experiment using amplification in two regimes: a low power regime, and a high power regime. These choices are related to the observed and documented behaviour [122] of the PCFs, while being pumped at different levels of pump power, as discussed in Section 7.2. In the low power regime, the coherent combination of mismatched PCFs demonstrated ten times amplification of the amplified seed over the input seed. In the high power regime, the coherent combination of mismatched PCFs was demonstrated with larger pump powers. The pump lasers were set to the highest levels of pump power feasible for the experimental conditions, taking into account the thermal effects on the PCFs [122].

8.1 Low power combination Results

The coherent combination of mismatched PCFs was demonstrated successfully. A combined average output power of 4.86 ± 0.01 W was achieved which, at a repetition rate of 6.49 MHz, gave a combined pulse energy of 749 ± 1.5 nJ. Using the amplified compressed pulse duration of 370 ± 30 fs, the peak power was estimated to be 1.5 ± 0.2 MW. The combined beam had a system efficiency of $85.8 \pm 0.3\%$ and a degree of linear polarisation of $83.3 \pm 0.3\%$. The linearly polarised output of the system was measured to be 3.79 ± 0.01 W, which gave a fraction of available beam power of $91.6 \pm 0.3\%$. A more detailed analysis of the results from the low power regime are presented in this section.

8.1.1 Power

The two PCFs were coherently combined and locked at a total system output power of 4.86 ± 0.01 W. Power measurements, as shown in Table 8.1, were taken from different points in the experiment, as shown in Figure 5.1, and used to determine the PCF gain (dB), degree of linear polarisation, system efficiency and fraction of available beam power. The annotated experimental setup diagram, in Figure 5.1a, Section 5.8, aids the identification of the measurement points shown in Table 8.1. Cross-referencing Table 8.1 to Figure 5.1, 'cube input' refers to the beams entering the combining polarising beam splitter (PBS) at Points (4) and (5). 'cube input (Combined)' refers to beams combined in the PBS, shown at Points (3) and (4). These measurements take into account the power loss of the beam from the 70 μm PCF, on its first pass through the combining PBS. Combining only occurs on the beam's second pass through the polarising beam splitter, so the power available for combining was 2.45 W, not 2.68 W. LP output is the maximum linearly polarised beam power found in the desired polarisation state measured at Point (1A). The rejected beam was also measured at (1A) using the Half Wave Plate (HWP), adjusted to give the minimum power throughput.

The losses quoted in Table 8.1 are defined as the power lost from the available system output power at Point (1), which is available to further stages of combination. It is a subtraction of the LP Output and rejected beam from the system output, (A-(B+C)) from Table 8.1. The losses in the output from the combining PBS are due to: back reflections of the HWP; back reflections from the diagnostic PBS after the Hänsch-Couillaud-detector pickoff wedge; and the power lost into a separate rejected beam from the diagnostic PBS. A number of factors were responsible for the losses experienced by the 70 μm PCF beam, as it propagated along the beam line. One of these factors was previously mentioned in Section 7.3, where the increase of the seed beam power increased the proportion of power leaking into the inner cladding. Other factors include the PCF's birefringence leading to a small quantity of elliptically polarised light and unpolarised light being produced from scattering effects inside the core. Both of these effects could cause amplification within the core of the fibre, resulting in beam power being lost at the combining PBS.

	70μm PCF	**80μm PCF**	**Combined**
Seed input	0.234 ± 0.002 (10)	0.225 ± 0.001 (11)	-
PCF Unamp. Output	0.0821 ± 0.004 (6)	0.090 ± 0.004 (7)	-
PCF Amp. Output	3.39 ± 0.01 (6)	3.39 ± 0.02 (7)	-
Cube input	2.68 ± 0.007 (5)	3.2 ± 0.05 (4)	-
Cube input (Combination)	2.45 ± 0.008 (3)	3.2 ± 0.05 (4)	-
Dark Port	0.187 ± 0.002 (4, 7B)	0.0962 ± 0.002 (5, 6B)	-
Cube output (unlocked)	2.06 ± 0.01 (2, 4B)	2.8 ± 0.05 (2, 5B)	4.85 ± 0.05 (2)
(A) System output (locked)	-	-	4.86 ± 0.01 (1A)
(B) LP Output	-	-	3.79 ± 0.01 (1A)
(C) Rejected Beam	-	-	0.345 ± 0.005 (1)
Losses (A-(B+C))	-	-	0.735 ± 0.01
Pump Powers	26.2 ± 0.06	15.3 ± 0.04	-
Amp. Factor (Over input)	11.6 dB	11.8 dB	-
Amp. Factor (Over emitted seed)	16.2 dB	15.76 dB	-

Table 8.1: The average powers (W) of the low power coherent combination experiment. All numbers in brackets refer to the positions indicated on Figure 5.1. The numbers postfixed with a 'B' indicate that that position was blocked during the measurement. The seed powers shown above are slightly lower than the 0.250 mW initially used for both PCFs, due to long term thermal effects in the seed laser. The LP output and rejected beam measurements were taken after sending the system output through a HWP and PBS. The loss was defined as the power lost from the available system output power at Point (1), which was available to further stages of combination. It was a subtraction of the LP Output and rejected beam from the system output, (A-(B+C). These losses are due to: back reflections of the HWP and diagnostic PBS after the Hänsch-Couillaud-detector pickoff wedge; the small amount of power directed to the HC detector from the pickoff wedge for locking; and power lost into a separate rejected beam from the diagnostic PBS.

Given the beam power which was available for combination, a pulse energy of $\sim 748 \pm 1.5$ nJ was calculated by dividing the system output power, by the repetition rate of 6.49 MHz. An autocorrelation width of 520 ± 45 fs was measured for the combined pulse, which was deconvolved with a Gaussian to recover the pulse duration of 370 ± 30 fs. From the autocorrelations, it was found that $\sim 75\%$ of the energy existed within the central peak. Therefore, the peak power of the system was calculated as 1.6 ± 0.2 MW.

8.1.2 Degree of linear polarisation and efficiencies

The degree of linear polarisation of the output beam gave an indication of the purity of polarisation of the combined beam. Furthermore, it also determined how much of the combined beam power could be recovered (fraction of available beam power) if it were to encounter polarisation sensitive optics. For polarisation based coherent combination, the fraction of

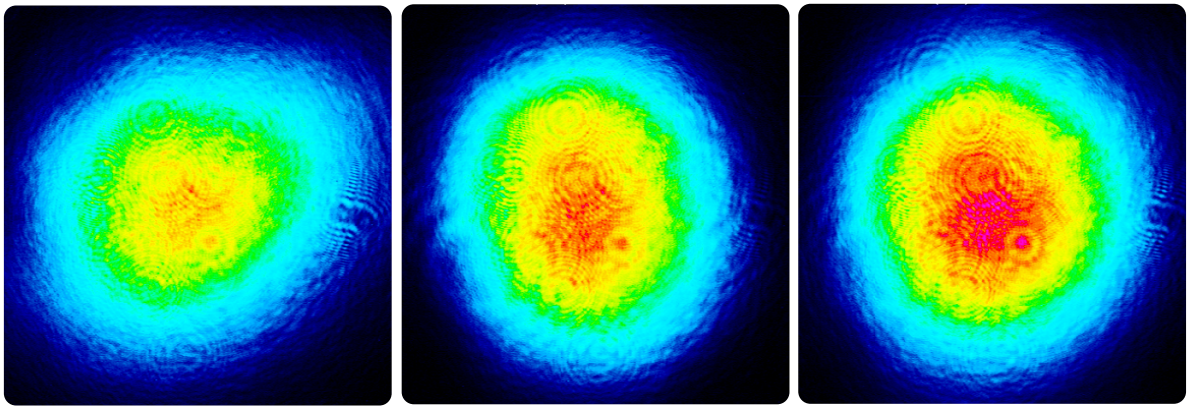
available beam power is the power which would contribute to further stages of coherent combination. Given that this experiment aimed to demonstrate that two mismatched lasers could contribute to a larger array of N coherently combined lasers, the degree of linear polarisation and fraction of available beam power were important figures. Using the LP output power (B) and rejected power (C) from Table 8.1, the degree of linear polarisation of the combined beam was calculated to be $83.3 \pm 0.3\%$. The 70 μm and 85 μm PCF had a degree of linear polarisation of $\sim 86\%$ and ~ 93 % respectively. However, because the combining PBS was polarised, the linear polarisation of the individual beams leaving the PBS was improved to $\sim 99\%$ each. The fraction of available beam power was calculated to be $91.6 \pm 0.3\%$, indicating that despite the mismatch between the PCFs, there was a high power fraction available to subsequent stages of combination.

The system efficiency compared the combined beam power with the summed powers of the individual beams at the combining PBS's inputs. This figure was calculated to be $85.8 \pm 0.3\%$, using the input powers to the PBS, at the point of combination. The powers are those in Table 8.1, from row 'cube input (Comb)' taken from Points (3), (4) and (2). For further stages of combination, the figure of interest is the fraction of available beam power, as seen in Equation (1.10), Chapter 1. This figure is dependent upon the degree of linear polarisation of the beam, and so for the combination of N lasers, a combined beam should have as high a degree of linear polarisation as can be achieved.

8.1.3 Beam profiles

The beam profiles for the 70 μm and 85 μm PCF respectively are shown in Figure 8.1(a) and 8.1(b).

The beam profiles show Gaussian-like characteristics with only a small asymmetry present in the beam profile from the 70 μm PCF. The combined beam shown in Figure 8.1(c) also shows a clean Gaussian-like beam profile, that is more uniform than either one of the individual beams. A number of factors prevented the actual measurement of M^2 for the combined beam, but it can be inferred that the LP output beam, which will propagate into further stages of combination, must be single mode. This is because only the fundamental mode



(a) (b) (c)

Figure 8.1: The beam profiles of the (a) 70 µm PCF, (b) 85 µm PCF and (c) combined beam. These beam profiles are viewed through a HWP and PBS in series.

from one beam will combine with the fundamental mode of another. The likelihood of both PCF's having the same higher order modes present, and in the same orientations and magnitudes is remote. The higher order modes will not combine efficiently, and therefore will be largely filtered out by any polarisation sensitive optics, as a rejected beam. The improvement of the combined beam profile over that of the individual beams, was also observed for four PCFs [38]. Each PCF's beam profile showed good single mode characteristics, with the combined beam's profile also being single mode and looking more uniform than any one of the individual beams. Looking specifically at the combined beam from the low power regime experiment, Figure 8.2 shows the beam transmitted through the PBS (desired polarisation state) and the rejected beam reflected from the PBS.

In Figure 8.2, there is good contrast between the desired and rejected beams, with beam powers of 3.79 W and 0.345 W for combined and rejected respectively. This was an improvement over the unamplified average powers seen in Section 7.2 during wavefront matching, where beam powers of 1.66 W and 0.317 W were measured for the combined and rejected beams respectively. Additionally, the combined beam profile from Figure 8.2(a) appears to be much more uniform, than that seen in the bottom row of Figure 7.15.

8.1.4 Autocorrelations

To show that the mismatched PCF amplifiers were not having a significant adverse effect on the compressed pulse, as predicted by the model, the autocorrelation of the compressed seed

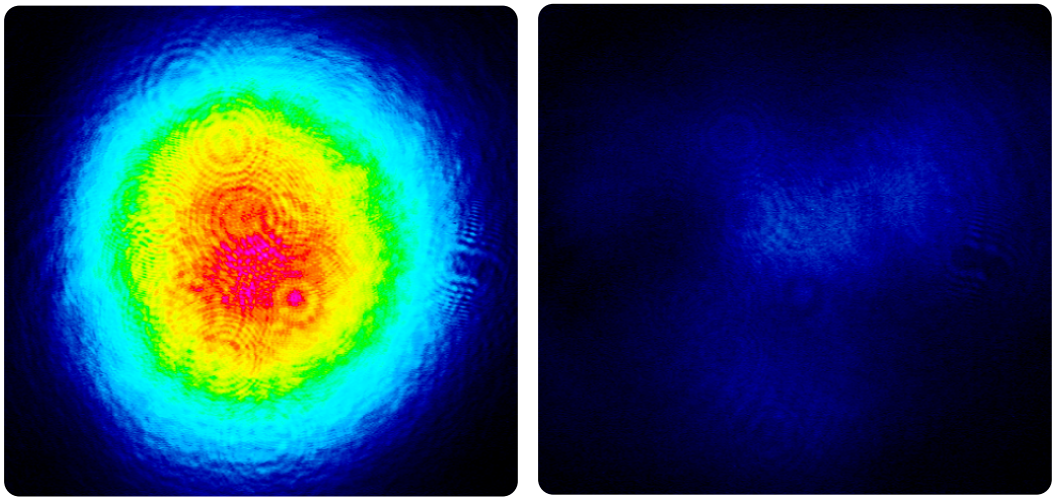


Figure 8.2: The comparison between (a) the combined beam with the desired polarisation state and (b) the rejected beam. The combined beam had an average power of 3.79 W, and the rejected beam had an average power of 0.345 W. Both beams were captured at the same exposure setting for comparison.

and the combined pulse were measured. The autocorrelation duration and pulse duration are both defined as the full width at half maximum of the peak, with the pulse duration having a deconvolution factor of $\sqrt{2}$ over the autocorrelation. This deconvolution factor depends upon the shape of the laser pulse and can vary depending on what fitting function fits the pulse shape best. A Gaussian was found to give a good fit to the temporal profile, as shown in Section 5.8. So, for all further measurements in this section, a deconvolution factor of $\sqrt{2}$ is implied, when it is stated that a pulse is deconvolved with a Gaussian. Figure 8.3 shows both the autocorrelation of the seed and amplified combined output, overlaid for comparison.

The seed pulse, shown in Figure 8.3, was optimally compressed to a pulse duration of 340 ± 30 fs, which was within the uncertainty of the ~ 315 fs transform limit predicted by the spectrum available to the pulse. The introduction of a large amount of non-linear phase in the main amplifier of the seed laser, ~ 7 rads, meant there was some residual phase that could not be compensated by the compressor. The split step Fourier method model predicted that this would cause the pulse duration to grow to ~ 366 fs, which was just within the uncertainty of the measured seed autocorrelation. The combined pulse had a deconvolved pulse duration of 370 ± 30 fs, which was in good agreement with the ~ 366 fs simulated value, and remains within the uncertainty of the measured seed autocorrelation. The two autocorrelations were taken at different times, with slightly different compressor settings. This was due to having

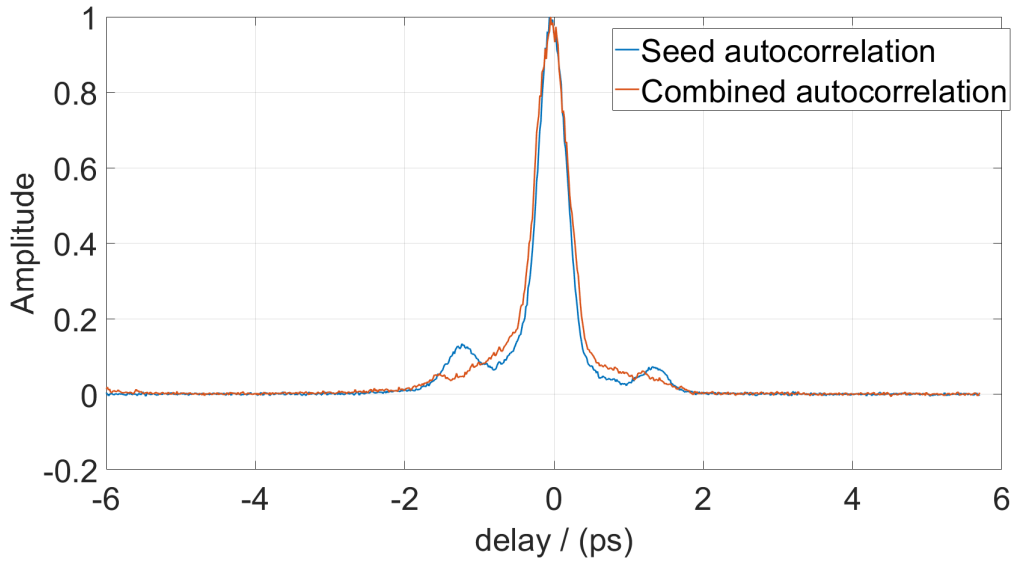


Figure 8.3: The autocorrelation taken of the seed pulse incident to the PCFs and the combined pulse. The seed and combined pulse have autocorrelation durations 480fs $\pm$ 40 fs and 520fs $\pm$ 45 fs respectively with the deconvolved pulse durations being 340 $\pm$ 30 fs and 370 $\pm$ 30 fs respectively when deconvolved with a Gaussian. Both are within uncertainty of one another.

to realign the beam into the compressor from scratch for each measurement. The small features seen either side of the seed pulse were likely smeared out of the combined pulse autocorrelation due to the additional phases acquired by the PCFs and then subsequently removed by the compressor.

8.1.5 Summary of low power coherent combination

The low power regime has demonstrated the coherent combination of two PCFs of vastly different lengths, core sizes, and structures. The wavefront matching was the largest hurdle to achieving this result, which when corrected, boosted the degree of linear polarisation value from the lower values of 40% - 50%, to the much higher value of $\sim 83.3 \pm 0.3\%$. The fraction of available beam power was measured to be $91.6 \pm 0.3\%$, indicating that a high percentage of the power would be available to further stages of combination. The comparison between the seed and amplified pulse autocorrelations in Figure 8.3, showed that the mismatch between the PCFs was not preventing compression back to the $\sim$200 fs transform limit of the oscillator. All of the structure that existed in the pulse, and prevented compression back to the oscillator's pulse duration, was observed during the seed laser characterisation in Chapter 6. Both the seed and combined amplified pulse durations were within the uncertainty of each other, and the pulse duration predicted by the model. The two PCF's were able to

produce a combined beam profile, which looked more uniform than either of the individual combining beams, and which showed what looked to be single mode characteristics. In the absence of an M^2 measurement due to equipment shortages, the linearly polarised combined beam observed after passing through a PBS, could be implied to be single mode because of the physics behind combining beams. For efficient coherent combination, matched wavefronts and intensity profiles are required. The PCFs were designed to operate in effective single mode operation, and the only common mode between the PCFs was the fundamental mode. The higher order modes were mismatched between the PCFs and contained only small amounts of beam power. The high fraction of available beam power signifies that only a small percentage of the total beam power was due to unpolarised light or higher order modes, which had not combined. Subsequent stages of combination would further reinforce this mode cleaning effect, so one could expect the combined beam from N lasers to be single mode, $M^2 = 1$. In terms of general power loss, the system efficiency was measured to be $85.8 \pm 0.3\%$. On a separate experimental run, the system was tested to determine the length of time that the system could be stabilised. After the usual system warmup time of an hour, the system was locked and monitored over a period of 6 hours 45 minutes, after which the system was shut down. Once the system had thermalised and the beams had been combined, the system could be repeatably locked and stabilised.

8.2 High power combination results

The coherent combination of mismatched PCFs was demonstrated in the high power regime. A combined average output power of 20.3 ± 0.05 W was demonstrated which, at a repetition rate of 6.49 MHz, gave a pulse energy of $3.1 \pm 0.01\,\mu$J. Using the amplified compressed pulse duration of 380 ± 60 fs, the peak power was estimated to be 5.8 ± 1.3 MW. The linearly polarised output of the system was measured to be 15.3 ± 0.05 W, with a system efficiency of $88.4 \pm 0.4\%$. The degree of linear polarisation of the system was measured to be $85.9 \pm 0.5\%$, which produced a measured fraction of available beam power of $93.0 \pm 0.5\%$. A more detailed analysis of the results from the high power regime is presented in this section.

8.2.1 Experimental modification

The aim of the high power coherent combination experiment was to push the system to higher output powers, whilst trying to improve upon the degree of linear polarisation and the efficiency of the system. Initially, the objective was to see how much power the system was capable of putting into a combined beam, which still possessed the high efficiencies and degree of linear polarisation from the low power experiment. The 85 μm PCF was able to operate with a large range of output powers for the given $\sim$ 250 mW input seed, and was able to reach average output powers well over 20 W. The 70 μm PCF, given the seed power of $\sim$ 250 mW, was only able to reach amplified seed average powers of $\sim$8.4 W, with mode degradation becoming a barrier to further amplification. For equal pump powers, the 85 μm PCF was producing approximately twice the amplified seed power of the 70 μm PCF. This was due, in part, to the 70 μm PCF having a smaller volume core, so had less gain medium to hold energy. Another factor that might have affected the power difference was the actual doping level within the 70 μm core, but this knowledge was propriety so this could not be scrutinised further. With the intention of getting large, combined output powers, without one PCF supplying the majority of the output power to the combined beam, it was decided to increase the seed power of the 70 μm PCF. This resulted in the 70 μm PCF output power increasing from 8.4 W to 13.3 W. However, this power increase did increase the quantity of power being leaked into the inner cladding, as was observed in Section 7.3. This power leakage would increase the amount of loss around the combining PBS. However, the aim was not to minimise the loss, but to get as much power as could be achieved into a combined beam. So for the high power regime, some loss was accepted in order to increase the beam power of the 70 μm PCF that would actually contribute to a combined beam.

Given that the longer 85 μm PCF, working at higher power was going to acquire more non-linear phase than the shorter PCF, it was necessary to control the non-linear phase difference between the two PCFs. One method was to alter the pump power delivered to the 85 μm PCF, thereby, affecting the amplified pulse energy. However, as discussed earlier in Section 7.3, the pump power had been set to allow the PCF to thermalise, so that the output beam's divergence did not change. Another method of reducing the non-linear phase

was to change the polarisation state of the incident pulse. By propagating a pulse down a PCF with a circular polarisation state, rather than a linear polarisation state, the non-linear phase acquired by the pulse could be reduced to 2/3. This effect occurs due to the non-linear refractive index having a smaller value in the presence of circularly polarised light, than when in the presence of linearly polarised light. Additionally, it has been demonstrated [125] that, for strongly chirped pulses, non-linear phase acts as an anomalous chirp-contribution, which can be compensated for by the positive chirp-contribution of normal dispersion. In summary, for chirped pulse amplification, the non-linear phase acquired by stretched pulses can be compensated for, to an extent, by adjusting the quantity of normal dispersion acquired by the pulse.

In Figure 4.6 in Section 4.3, it can be seen that if there is a length difference between the PCFs, then a non-linear phase difference can be used to compensate for this extra material dispersion, to help ensure good combination. Looking at the effect of adding quarter wave plates, before and after the 85 μm PCF, the non-linear phase, given the measured input and output powers from the high power regime, was calculated to have reduced from $\phi_{\mathrm{NL}} = 0.1629$ rads to $\phi_{\mathrm{NL}} = 0.1086$ rads. For the 70 μm PCF, given an increase of the seed power from 250 mW to 750 mW, the non-linear phase acquired in the 70 μm PCF increased from $\phi_{\mathrm{NL}} = 0.0699$ rads to $\phi_{\mathrm{NL}} = 0.1267$ rads. This left a non-linear phase difference between the PCFs of $\Delta\phi_{\mathrm{NL}} = -0.0181$, which is a reduction of approximately five times from the $\Delta\phi_{\mathrm{NL}} = 0.093$ that existed before the modification to the experimental conditions. The degree of linear polarisation predicted by the split step Fourier method (SSFM) model from Section 6.2, was 99.99%, with the fraction of available beam power also predicted to be 99.99%. A high figure of merit (FOM), equation 4.35, of 99.6% can also be seen in Figure 4.6, when a non-linear phase difference of $\Delta\phi_{\mathrm{NL}} = 0.093$ is observed between the two combining beams from mismatched PCFs with $\Delta L = 300$ mm. Hence, the agreement between the SSFM model and the FOM estimation both suggested that the coherent combination of mismatched PCF's in the high power regime should work.

However, using the degree of linear polarisation measurement from the low power regime, the difference between the model predicted value of 99.99% and the previously measured

value of $\sim 83\%$ is fairly significant. Both the approximate FOM method and the SSFM model had predicted that the combination of two mismatched PCFs, at the increased power levels, should combine to give high 99% degree of linear polarisation values. Hence, the small difference in non-linear phase cannot be the dominant factor in the 17% drop in the degree of linear polarisation. Given the non-linear phase was small, and the model predicted good combination, the quarter wave plates used to reduce the non-linear phase were left in place, as they were still making a net improvement on $\Delta\phi_{\mathrm{NL}}$ overall. Another factor was that the pump power delivered to the 70 µm PCF led to degradation of the single mode output. The 85 µm PCF was a newer design, and was quoted as being capable of operating at amplified seed average powers of $\leq$ 100 W. The 70 µm PCF was an older design, and its maximum power capabilities were not quoted in the datasheet. So a test was performed to see how far the 70 µm PCF could be pumped, given a seed power of 0.750 W. At a level above $\sim$ 75 W, the PCF output beam profile was found to degrade, as shown in Figure 8.4, and was a repeatable threshold-like behaviour. This threshold-like behaviour has been observed before [126] and relates to a phenomena known as 'avoided-crossings' of modes. Typically, this effect limits the maximum power that can be deposited into a linearly polarised pulse before it degrades. Hence, it is usually avoided.

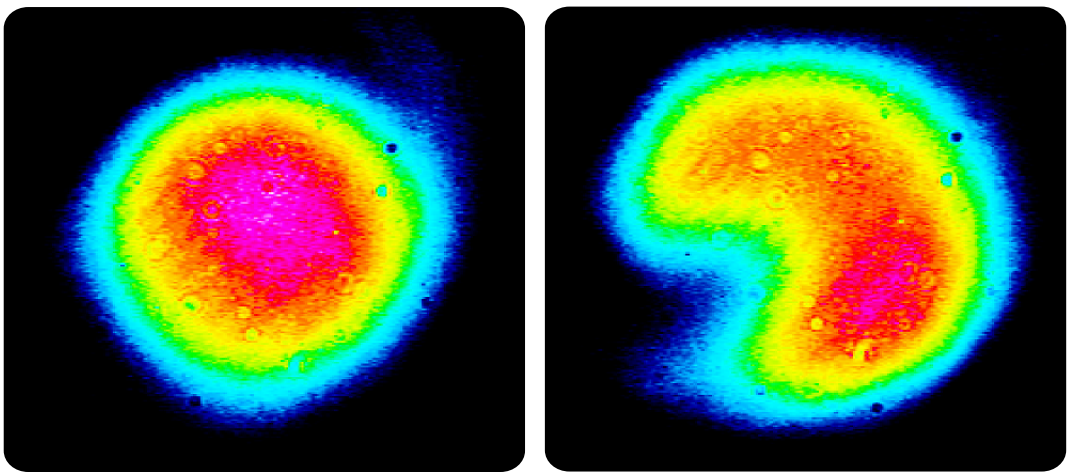


Figure 8.4: An example of the threshold-like behaviour observed with the 70 µm PCF, when the pump power was increased passed the threshold of $\sim$ 80 W. (a) The 70 µm PCF, when pumped with 75 W of pump power (b) The 70 µm PCF, when the pump power was increased to $\sim$ 80 W, and it was not allowed time to thermalise.

When the system had been left to thermalise for hour, it was found that the PCF given its current alignment, would begin to show mode degradation at a pump power beyond $\sim$ 75 W. So, to remain clear of this effect and to avoid having to realign the system to work at these higher powers, it was decided to set the maximum pump power available to the 70 µm PCF to $\sim$ 70 W. The PCF could be cooled and re-thermalised to this pump power repeatably, with only minor adjustments to the input seed beam alignment required from day to day, to maintain or improve the PCF's Gaussian-like output beam profile. The 85 µm PCF was able to surpass this pump power, and still maintain stable Gaussian-like beam characteristics. However, the pump power for the 85 µm PCF was limited, so that its power contribution to the combined beam did not dominate. Both PCFs were pumped with roughly equal amounts of pump power, to represent the case where the two PCFs from this experiment are part of a large array, and are all treated the same.

Given that the seed power for the 70 µm PCF had increased from 0.250 W to 0.750 W, a check was performed to confirm if the beam expanding behaviour of the PCF, given an increasing pump power, was still present. This measurement was taken on a beam profiling camera fixed at a distance of $\sim$ 1.9 m from the PCF output, as shown in Figure 8.5. The same beam expanding behaviour was still observed to be present, ruling out the idea that the beam expansion effect is related to the input seed beam power. This also suggested that the wavefront matching telescope should be recalculated, when operating in the higher power regime. The PCFs were pumped at $\sim$ 70 W each, and left to thermalise in the high power regime. Once the beam diameter and divergence had been measured, several lens combinations were trialled, based upon what was available in the lab. An approximate 1:2.6 telescope using $f = 15.29$ mm and $f = 40$ mm lenses separated by $\sim$ 53 mm was found to provide suitable matching in beam diameter and divergence at the point of combination. After setting up the telescope in the 85 µm PCF beam line, and adjusting the telescope distance, the system was coherently combined and locked.

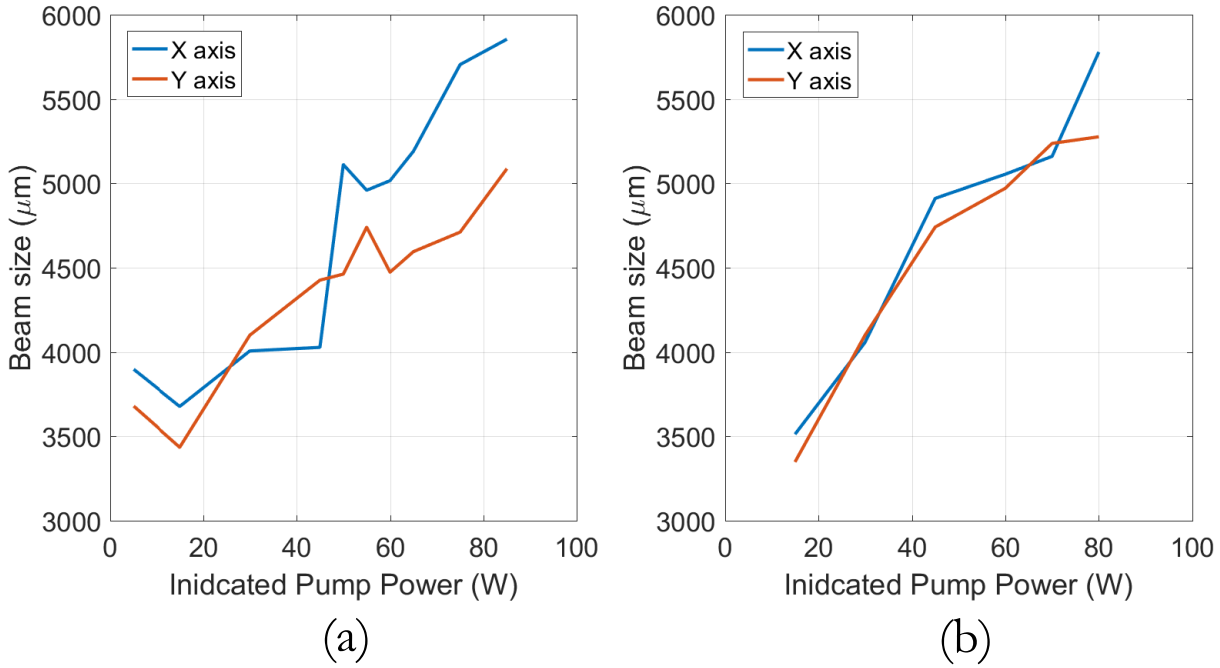


Figure 8.5: (b) The change in beam diameter, as the 70 µm PCF was seeded with 750 mW and pumped with increasing amounts of pump power. The beam diameter for the 250 mW seed input, (a), has been included for comparison, and both show an increase in the beam diameter.

8.2.2 Power

The measurement of the beam power was done, as previously in the low power setup. The results can be found in Table 8.2. A pulse energy of $\sim 3.1 \pm 0.005\,\mu$J was measured, which using the recovered Gaussian pulse duration of $\sim 380 \pm 60$ fs deconvolved from an autocorrelation of $\sim 540 \pm 80$ fs, gave a peak power of 5.8 ± 1.3 MW. This assumed a repetition rate of 6.49 MHz, and 70% of the total pulse energy existing under the central peak. Losses from the system were calculated to be 3.94 W, taken as a subtraction of the LP Output and rejected beam from the system output, (A-(B+C)). These losses are due to: a back reflection from the HWP (which was 1.85 W alone) and diagnostic PBS after the Hänsch-Couillaud-detector pickoff wedge; and power lost into a separate rejected beam from the diagnostic PBS.

8.2.3 Degree of linear polarisation

The degree of linear polarisation of the combined beam was measured to be $85.9 \pm 0.5\%$. This figure, measured for the high power regime, showed that time invested into ensuring that the combining beams were properly overlapped spatially and temporally, the previous

	$70\mu m$ PCF	$80\mu m$ PCF	Combined
Seed input	0.719 ± 0.003 (10)	0.236 ± 0.001 (11)	-
PCF Unamp. Output	0.395 ± 0.004 (6)	0.154 ± 0.004 (7)	-
PCF Amp. Output	13.3 ± 0.05 (6)	17.0 ± 0.1 (7)	-
Cube input	12.7 ± 0.05 (5)	15.3 ± 0.05 (4)	-
Cube input (combined)	8.91 ± 0.03 (3)	15.3 ± 0.05 (4)	-
Dark Port	2.48 ± 0.02 (4, 7B)	0.166 ± 0.007 (5, 6B)	-
Cube output (unlocked)	7.69 ± 0.03 (2, 4B)	14.2 ± 0.05 (2, 5B)	21.4 ± 0.1 (2)
(A) System output (locked)	-	-	20.3 ± 0.05 (1)
(B) LP Output	-	-	15.3 ± 0.05 (1)
(C) Rejected Beam	-	-	1.16 ± 0.04 (1)
Losses (A-(B+C))	-	-	3.94 ± 0.05
Amp. Factor (Over input)	12.67dB	18.58dB	-
Amp. Factor (Over emitted seed)	15.27dB	20.43dB	-

Table 8.2: The average powers (W) of the high power coherent combination experiment. All numbers in brackets refer to the positions indicated on Figure 5.1a. The numbers postfixed with a 'B' indicated that that position was blocked during the measurement. The seed power deficits from the target figures of 750 mW and 250 mW, were due to long term thermal effects in the seed laser. The LP output and rejected beam were taken by sending the system output through a HWP and diagnostic PBS. The losses are taken as a subtraction of the LP Output and rejected beam from the system output, (A-(B+C). These losses are due to: back reflections of the HWP and diagnostic polarising beam splitter after the Hänsch-Couillaud-detector pickoff wedge; and power lost into a separate rejected beam from the diagnostic polarising beam splitter.

best degree of linear polarisation of $83.3 \pm 0.3\%$ could be improved upon. The degree of linear polarisation of the 70 µm PCF does appear to be much worse than was measured previously, at 56.5%. The degree of linear polarisation of the 70 µm PCF was calculated using the input power (combination) (8.91 W) to the PBS and the dark port power (2.48 W). Although this is a low degree of linear polarisation, because the beam from the 70 µm PCF makes another pass through the PBS, the power that was lost into the dark port at the point of combination (1.22 W) was the difference between the 'cube input' (combination) (8.91W) and the unlocked cube output (7.69 W). After the first pass, the degree of linear polarisation was improved to 72.6%. After the combining PBS, both beams had degree of linear polarisation values at the +99% as both could be extinguished completely at the degree of linear polarisation diagnostic-measurement. One explanation for the drop in the degree of linear polarisation of the 70 µm PCF, was the amount of beam power leaking into the inner cladding of the PCF. Around the central single mode spot, some structure could be seen, highlighted when a card was placed into the beam. This structure was much harder to see on a camera, due to its

limited dynamic range. Increasing the exposure level to observe the faint structure, caused the camera to become saturated from the intense central beam of the core. Given that the PCF design using stress rods has been discontinued by the manufacturer, this might suggest that the design might be partially responsible for this degree of linear polarisation behaviour. The 85 μm PCF had a much better degree of linear polarisation from the outset, sitting at 97.7%.

8.2.4 Efficiencies

For two PCFs, the system efficiency takes into account the losses due to absorption, polarisation, and combination, when comparing the input to output. The system efficiency was measured to be $88.4 \pm 0.4\%$. This metric indicated that the vast majority of beam power, which entered the combining PBS, was making it to the point of combination. The degree of linear polarisation of $85.9 \pm 0.5\%$ suggested that a high proportion of the beam power, which was available for combining, minus losses, was combining coherently. The fraction of available beam power confirmed this, being measured as $93.0 \pm 0.5\%$, giving only a $\sim 7\%$ projected power loss at this stage of combination. These figures show that both the degree of linear polarisation, and the fraction of available beam power, have made small improvements over those measured in the low power regime. This was due to better alignment and matching of the beam diameter and divergences of the two combining beams.

8.2.5 Beam profiles

With the system optimally coupled, the beam profiles seen in Figure 8.6, show Gaussian-like beam characteristics. As raised earlier in section 8.1.3, there were a number of obstacles to measuring the M^2 of the beams, which would have given a more quantitative analysis of the beam quality, and whether the combined beam was single mode. However, due to the physics of coherent combination, the combined LP polarised beam should be largely single, mode as only two beams with the same wavefronts and intensity distributions can efficiently combine. Higher order modes, which do not combine efficiently, contribute to power losses in the system. Figure 8.7 shows the comparison between the combined beam and the rejected beam.

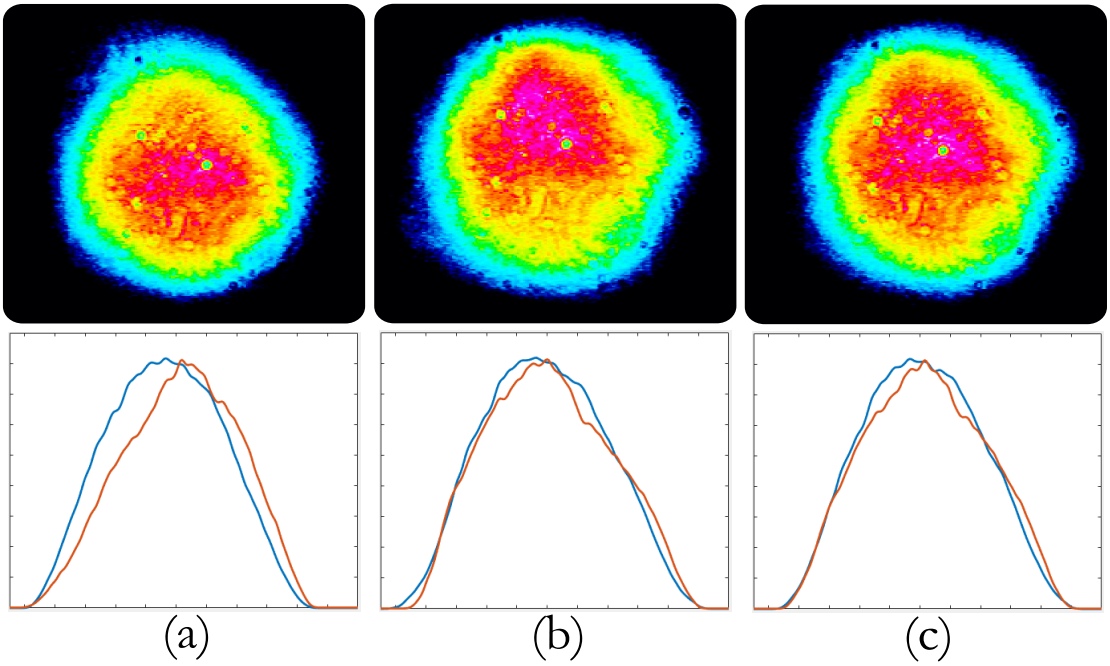


Figure 8.6: The beam profiles recorded, while operating in the high power regime: (a) The 70 µm PCF. (b) The 85 µm PCF. (c) The combined beam. In the absence of the M^2 measurement, the horizontal (blue) and vertical (red) effective slit profiles of each beam are shown below, to demonstrate that the combined beam has a slightly more uniform beam profile that either the 70 µm or 85 µm beam.

From figure 8.7, the contrast seen between the desired and rejected beam was more significant than it was for the low power regime. The extra image in figure 8.7(c) has had the exposure increased to show the presence of the rejected beam. The rejected beam showed a ring type structure, which could be evidence to suggest that a small wavefront mismatch remained between the combining beams. Based upon the contrast between figures 8.7(a) and 8.7(b), the mismatch could be expected to be small. However, the rejected power measurement of 1.16 W showed that it still remained significant.

8.2.6 Autocorrelations

The autocorrelation of the combined beam was taken and found to have a good match to the seed autocorrelation. The FWHM pulse duration was 380 ± 60 fs, as seen in Figure 8.8.

The large uncertainty accompanying this measurement was due, in part, to the pixel size of the camera and the size of the autocorrelation on the CCD. The high power autocorrelation used a telescope to magnify the image, in order to collect enough second harmonic light. This had the knock-on effect of reducing the resolution of the autocorrelation profile on the camera. In addition, because the camera was set to sum up 2x2 groups of pixels into one large macro

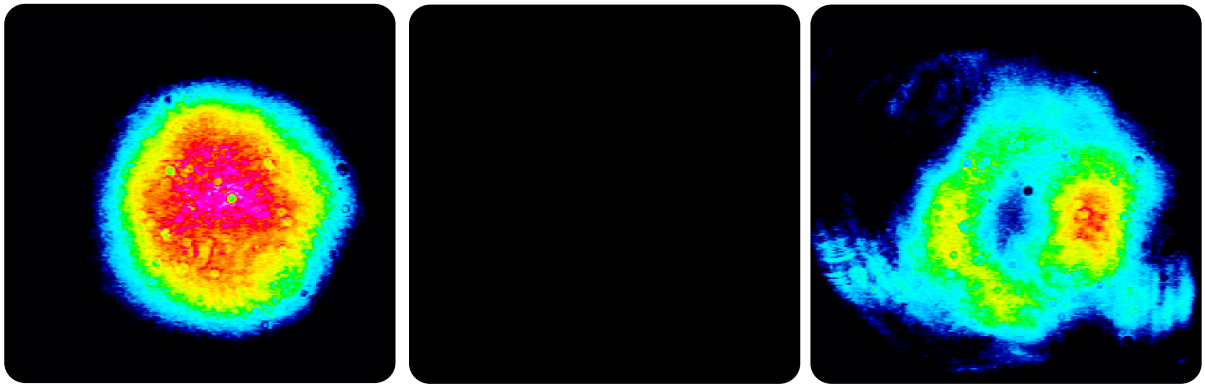


Figure 8.7: The beam profiles of the combined beam observed at point (1A) from Figure 5.1, Section 5.8. (a) The combined beam with the desired polarisation state. (b) The rejected beam. (c) The rejected beam observed with an increased exposure level. Even though the rejected beam cannot be clearly seen in Figure (b), increasing the exposure shows that it is still present.

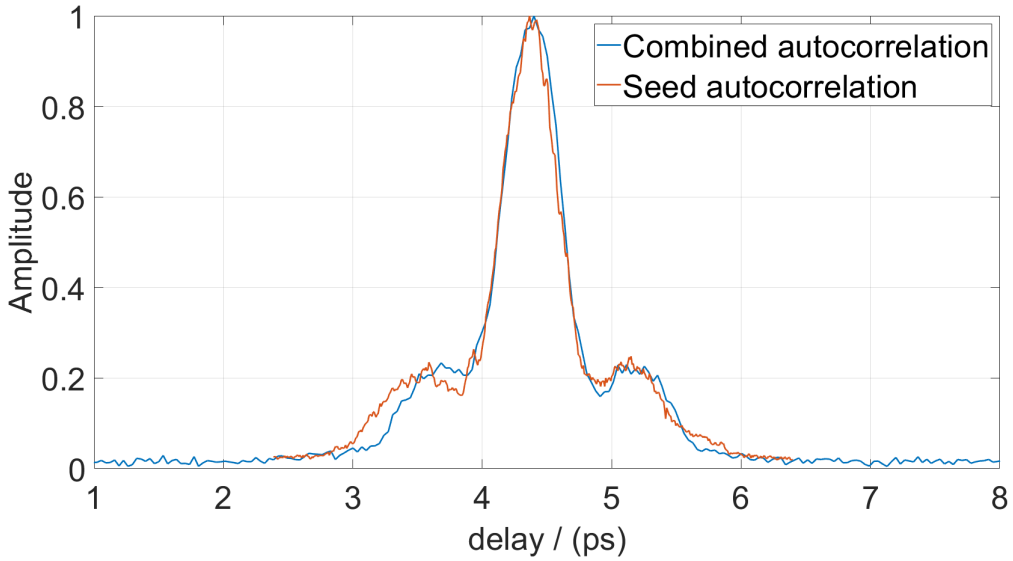


Figure 8.8: The autocorrelation of the seed laser and high power combined beam. The seed and combined beams had autocorrelation widths of 530 ± 20 fs and 540 ± 80 fs respectively, with the pulse durations being 370 ± 10 fs and 380 ± 60 fs respectively, when deconvolved with a Gaussian.

pixel, the resolution of the autocorrelation was further reduced. This was not desired, but despite this, it can be seen that the autocorrelation remained similar in shape to that of the seed pulse (which was retaken shortly afterwards with better resolution) and within the uncertainty of its pulse duration. This implied that the mismatched PCFs, at least at these power levels, were having minimal effects on the pulse envelope and duration, with the pulse's observed temporal structure being due to the seed laser itself. A comparison of the simulated autocorrelation and the measured autocorrelation showed good agreement also in terms of their pulse durations, as shown in Figure 8.9.

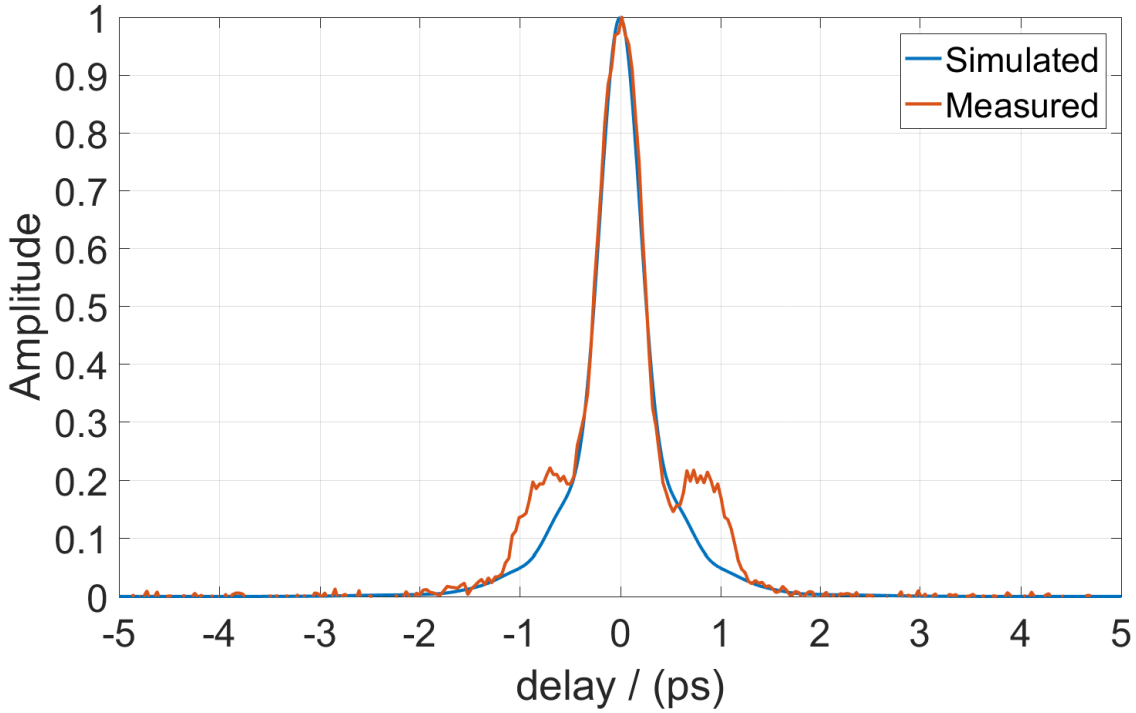


Figure 8.9: The comparison between the high power autocorrelation from the model, and the measured combined high power autocorrelation. The measured pulse duration of (380 ± 60 fs) agreed with the simulated value (366 fs) to within the uncertainty. The wings of the autocorrelation did show some disparity, which could have been due to the pre-pulse, identified previously in Figure 6.13, being amplified more than was seen in the model.

From Figure 8.9, simulations and measurements show good agreement, as the measured pulse was within the 60 fs measured uncertainty of the simulated $\sim$ 366 fs pulse. The simulated value was similar to the simulated pulse, taken at lower power, for two reasons. Firstly, the compressor settings were adjusted for optimal compression each time the model was run, and secondly, the amount of NL Phase acquired by both pulses was relatively small, as shown in Section 8.2.1. Even though the high power experiment acquired more non-linear phase in total, it was compensated for by adjusting the compressor. In the side lobes, there was an obvious difference between the simulated and measured pulse, and this could have been due to the pre-pulse, identified by the model in Figure 6.13, Section 6.1.7, being subject to greater amplification than was simulated in the model. To confirm that, a Frequency Resolved Optical Gating (FROG) or Spectral Phase Interferometry for Direct Electric-field Reconstruction (SPIDER), which was not available, would have been required to reconstruct the pulse. Without that, one can only make a best guess, using the reconstructed pulse from Figure 6.13 as a guide.

8.2.7 Lock stability

To characterise the stability of the system, the beam power was monitored through a PBC onto a photodiode for the system when locked, and when unlocked. The system was stabilised using two layers of locking. The first was a PID controller driving a piezo-mounted mirror, to keep the arm lengths of the interferometer matched. The second was a PID controller, driving a small translation stage, which monitored the output signal sent to the piezo. The aim of this secondary locking loop was to keep the piezo mounted mirror near the centre of its range. By monitoring the piezo output, when the signal was within 10% of the boundaries, the stepper motor would move the piezo mounted mirror back towards the centre of its range. The benefit was that the system remained locked over large changes in path length. The drawback was that this could affect the locking stability, by creating an extra noise source. Figure 8.10 shows how the beam power changed over time.

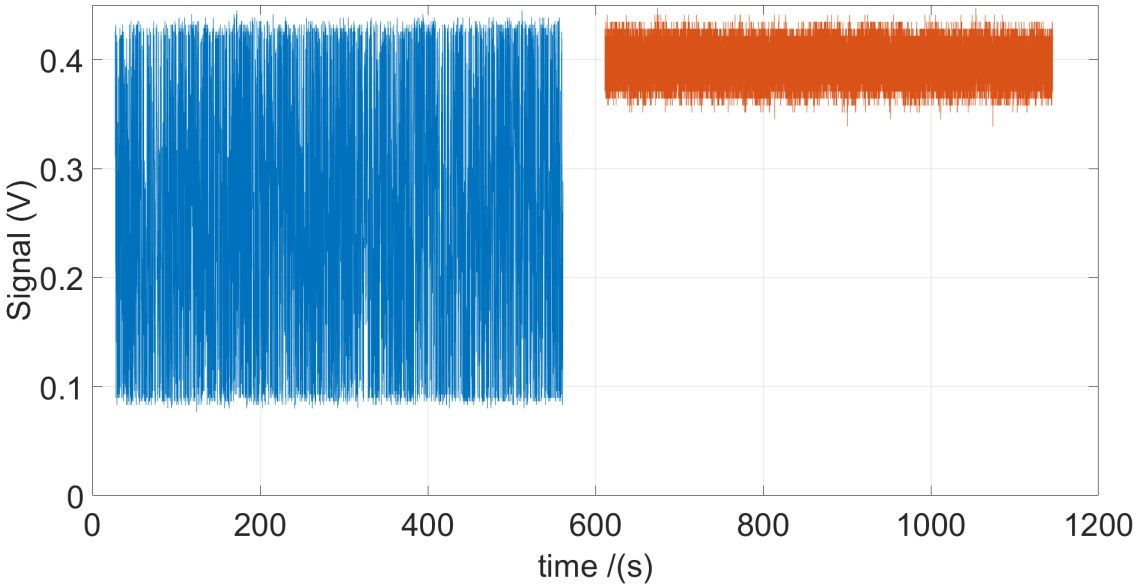


Figure 8.10: The temporal evolution of the combined beams average power, measured as a photodiode voltage, in the unlocked (blue) and locked (red) state. Both were taken at different times, so a gap was left between data sets. The power variations in the unlocked signal caused the voltage to vary from ~ 0.09 V to ~ 0.44 V. The locked signal showed variations in the range of ~ 0.36 V to ~ 0.44 V

Here the stability of the system had been locked to better than an RMS of $\sim \lambda/18$. To calculate the stability of the system, the photodiode voltage needed to be converted to a length scale. The process used is demonstrated in Figure 8.11. The locked data series, from Figure 8.10, was then binned into a histogram to allow the RMS deviation to be calculated,

as shown in Figure 8.12. Compared to other experiments, stabilities of $< \lambda/10$ [81, 82] and $< \lambda/25$ [127] have been reported when using tiled aperture arrangements, with $\lambda/40$ [85] being reported for the coherent combination of two PCFs using a fibre coupled phase modulator as opposed to a piezo-mounted mirror. The power spectral density of the locked and unlocked beams from Figure 8.10 were compared to show the noise suppression effect of the PID lock, as shown in Figure 8.13. In Figure 8.13, it can be seen that the lock had a large impact in stabilising the interferometer. In addition, several sources of noise from 50-500 Hz were suppressed, at the cost of a small increase in the background noise.

The bandwidth of the lock was 1 kHz, which was what determined the upper-frequency limit of 500 Hz, as shown in Figure 8.13. For comparison to the other phase shifting options available, spatial light modulators have a frame rate of approximately 30-60 Hz, piezo-electric driven mirrors have a bandwidth of a few kHz, and electro-optical phase-modulators have a bandwidth of a few MHz. Spatial light modulators are not suitable for suppressing large sources of noise due to their slow update rates. Electro-optical phase-modulators are very good at suppressing noise sources, but their range of phase correction is generally limited to the range of $\pm\pi$. Piezo-electric driven mirrors have a smaller bandwidth, but make up for this with their larger range of phase compensation and lower cost.

$$V \approx R_\lambda P R_L$$
$$V = V_0 \left(\frac{\cos(\Phi_D) + 1}{2} \right)$$
$$\Phi_D = \text{acos} \left(\frac{2V}{V_0} - 1 \right)$$
$$L_d = \lambda \frac{\Phi_D}{2\pi}$$

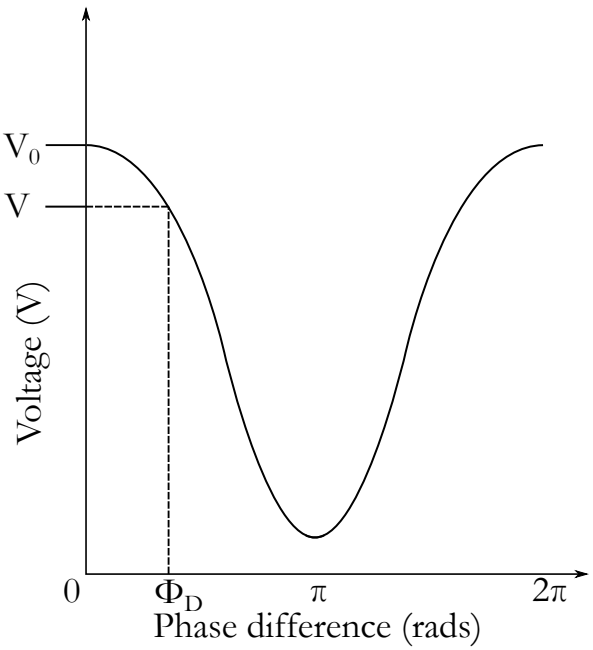


Figure 8.11: V is the photodiode voltage, V_0 is the peak photodiode voltage, R_λ is the responsivity of the photodiode at a wavelength of λ, P is the beam power striking the photodiode, R_L is the resistance of the photodiode circuit, Φ_D is the phase difference between the two pulses, and L_d is the path difference in terms of fractions of a wavelength. The first equation shows the approximate relationship between the beam power striking the photodiode and photodiode voltage. The second and third equation show how the change in voltage can be converted into a phase.

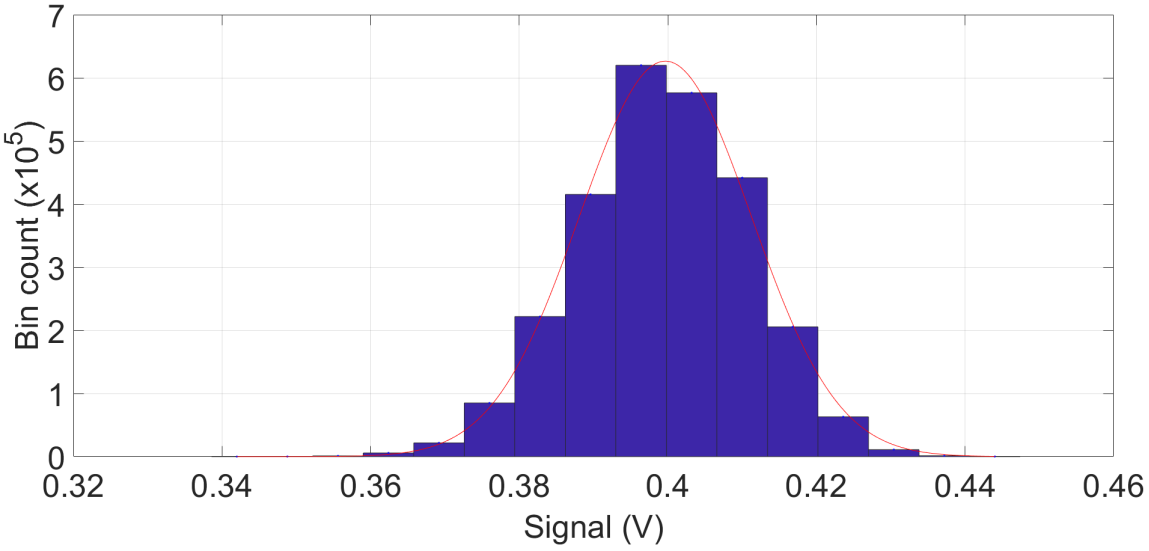


Figure 8.12: The locked photodiode voltage data binned into a histogram (blue), with a Gaussian fitted to the distribution (red line). The RMS of the data gives a stability of $\sim \lambda/18$, taken over a time span of $\sim$ 18 minutes.

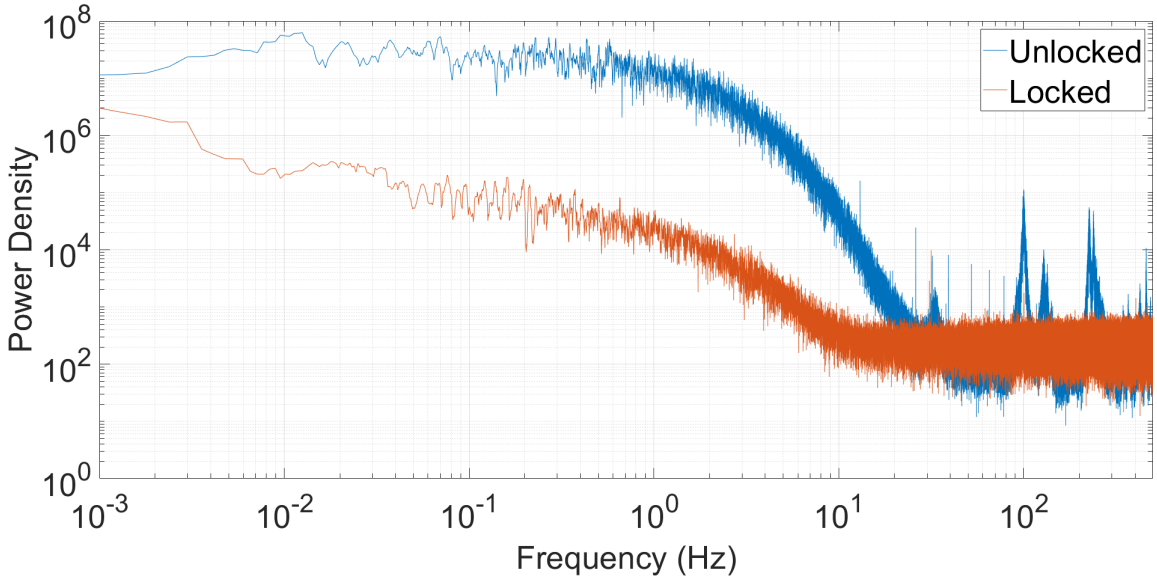


Figure 8.13: The power spectral density of the combined beam power in the locked and unlocked state, sent through a PBC and observed on a photodiode. The lock had a significant impact on the stability of the system.

8.2.8 High power Experiment summary

The high power experiment was successful in combining two mismatched PCFs, which were pushed to the readily achievable limits of the system. The combined average output power was ~ 40 times larger than the input seed powers, sitting at 20.3 ± 0.03 W. Compared to the low power experiment, there was a four times increase in the average combined output power. The measured pulse energy was $3.13 \pm 0.01\,\mu$J, with $\sim 70\%$ of the pulse energy measured to be within the AC's central peak. The peak power of the laser pulse was calculated to be 5.84 ± 1.3 MW, with a compressed pulse duration of $\sim 380 \pm 60$ fs. The calculated peak power for the high power regime was an improvement of 3.8 times over that of the low power experiment, where the peak power was calculated to be 1.55 ± 0.19 MW. The degree of linear polarisation and fraction of available beam power was measured to be $85.9 \pm 0.5\%$ and $93.0 \pm 0.5\%$ respectively.

The beam profile for the 70 μm PCF was found to be sensitive to the pump power launched into it. The beam profile quality deteriorated in a threshold-like fashion, once 80-90 W of pump power had been reached. Although the deterioration was not as extreme as some of the ∼kHz rate modal instabilities simulated and observed by other groups when pushing their PCFs to the limit [73, 74], it did have the potential to ruin the combination between the two PCFs. This led to the decision to back off from the maximum pump power, until the

PCF was seen to be operating reliably and repeatably. The system's stability was judged to have an RMS piston phase uncertainty (discussed in Chapter 7) better than $\sim \lambda/18$, which is on par with other combination experiments seen in the literature. A floating optical bench helped to mitigate these noise effects, but noise sources on the table, such as cooling pipes from chillers and air flow from air conditioning units, all contributed to the background noise sources. The combined degree of linear polarisation figure increased slightly over that measured in the low power regime, and that was attributed to the increased effort made to achieve optimum spatial and temporal overlap. However, something was preventing the degree of linear polarisation from reaching the high 90% range. Some of this was due to wavefront mismatch error, as the simple model used to calculate the wavefront correcting telescope would not produce perfect wavefront matching. The simple model did not take into account the wavefront perturbations that could occur due to imperfectly flat surfaces, scratches and pits in the optical components. To identify why this loss was occurring, the wavefronts of the beams from both PCFs, were subject to further investigation. This analysis is found in Chapter 9.

8.3 Summary

The results were presented for the coherent combination of two mismatched PCFs, covering a low and high power regime. The low power regime was a proof of principle, successfully demonstrating that, despite the very large difference between the PCFs in terms of length and core diameter, the beams could be combined and give a degree of linear polarisation of $83.3 \pm 0.3\%$, with a fraction of available beam power of $91.6 \pm 1.5\%$. The pulse energy of 749 ± 1.5 nJ produced a peak power of 1.5 ± 0.2 MW, when a pulse duration of 370 ± 30 fs was used, with 75 % of the pulse energy residing under the central peak. The autocorrelations for the seed input and amplified output pulses were the same to within uncertainty, which showed that the coherent combination of mismatched PCF's in the low power regime had a negligible impact on the pulse duration.

The high power regime aimed to maximise the combined output power. Limitations of the 70 μm PCF due to mode instabilities at high pump powers prevented the PCF from using all the available pump power. With these limitations in mind, sensible conditions for

the 85 μm PCF were chosen, and the high power output beams of both PCFs were combined. This produced a combined output power of 20.3 ± 0.03 W which, after losses were removed, gave 15.3 ± 0.05 W of linearly polarised light. The degree of linear polarisation of the LP beam power was measured to be $85.9 \pm 0.5\%$, with the fraction of available beam power being $93.0 \pm 0.5\%$. This was found to be a small improvement over the low power regime. The pulse energy was measured to be 3.1 ± 0.01 μJ , which gave a peak power of 5.8 ± 1.3 MW when a pulse duration of 375 ± 10 fs was used, with 75 % of the pulse energy residing within the central peak.

The autocorrelations of the seed and the amplified output were found to be the same to within uncertainty, showing that the difference in length between the PCFs was not having a detrimental impact on the pulse compression. In addition, the pulse duration was compared to the split step Fourier method model, introduced earlier in Section 6.2, and it was found to be a good match in terms of the pulse duration, with the simulated and measured pulse duration being the same to within uncertainty. One disparity between the simulated and measured autocorrelation was the larger wings seen on the measured pulse. During the execution of the model, the simulated intensity profile was calculated and found to have a small pre-pulse present. It was possible that the pre-pulse had a larger energy experimentally, and contributed to the wings seen on the measured pulse. However, without a FROG or a SPIDER, it was not possible to verify this.

Chapter 9

Analysis of Combination

After conducting two experiments in the low and high power regime, I performed analysis to identify where the differences between the model and experimentally observed degree of linear polarisations (DOLPs) were originating. I constructed computer models of Zernike polynomials and spherical wavefronts to estimate the degradation of the DOLP, and the fraction of available beam power due to the wavefront differences between the co-propagating beams. The model using the Zernike polynomials used libraries found online, and credit for these goes to Paul Fricker [128]. In this chapter, I used these libraries as a tool to help analyse the wavefront and identify the source which could explain the experimentally observed DOLP.

9.1 Estimate of degree-of-linear-polarisation loss

As could be seen from the DOLP measurement of both the low and high power experiment, a small amount of power was rejected at the polarising beam splitter after the system pickoff wedge. The theoretical model from Section 6.2, predicted a DOLP of 99.6%, whereas power measurements produced DOLP values for the low and high power regime experiments of $\sim 83.3 \pm 0.3\%$ and $\sim 85.9 \pm 0.5\%$ respectively. Because of the disparity between the measured and modelled DOLP values, it was thought that the phase differences acquired between the pulses were not to blame and that the wavefront might be the largest contributor. Although the telescope had gone some way to improving the DOLP by matching the wavefronts, it was necessary to identify the source of the remaining loss to assess whether this value could be improved further.

The wavefronts of the PCFs were measured using a Thorlabs wavefront sensor (WFS150-5C), possessing a 1280×1024 pixel CCD, with a 36×28 array of micro lenses attached to the front. A beam incident on the lenslet array would be focused down to 36×28 spots, which vary their position on the CCD based upon the curvature of the wavefront striking the lenses. However, the wavefront sensor must first be calibrated with a laser beam known to possess a spherical wavefront. This wavefront would act as the 'reference' wavefront. In the case of this experiment, the reference wavefront was calibrated using a Fianium ultrafast laser known to have excellent beam characteristics and therefore, assumed to have a perfectly spherical wavefront. All subsequent wavefront measurements taken, with respect to the reference wavefront, were relative measurements. The wavefronts of the 70 μm PCF and 85 μm PCF were measured using a custom beam line with the wavefront sensor placed in the far field, not shown in Figure 5.1. This beam line consisted of the pump-coupling/seed-decoupling lens, two turning mirrors, and an additional lens. Because the lenslet array reduced the effective resolution from 1280×1024 to 36×28, the additional lens was used to ensure the beam in question filled as much of the wavefront sensor as possible. The wavefront sensor was also prone to damage, so a large beam was chosen to mitigate this effect. The reason for the custom beamline was to remove the optics of the coherent combination experiment from the wavefront measurements of each PCF.

The fine structure on the wavefronts was hard to discern in surface plots in Figures 9.1(a) and 9.1(b). However, a small difference could be seen between the ROC in plots (a) and (b), by noting the minimum values of the figure corners. This difference was due to the different lenses used to expand and collimate the beam to fill the wavefront sensor, so different radii of curvature were expected. A contour plot highlighted the small wavefront perturbations that were present on the wavefronts. Taking the difference between these two wavefronts and converting into a phase, gave a residual which made the wavefront differences more apparent as shown in Figure 9.2.

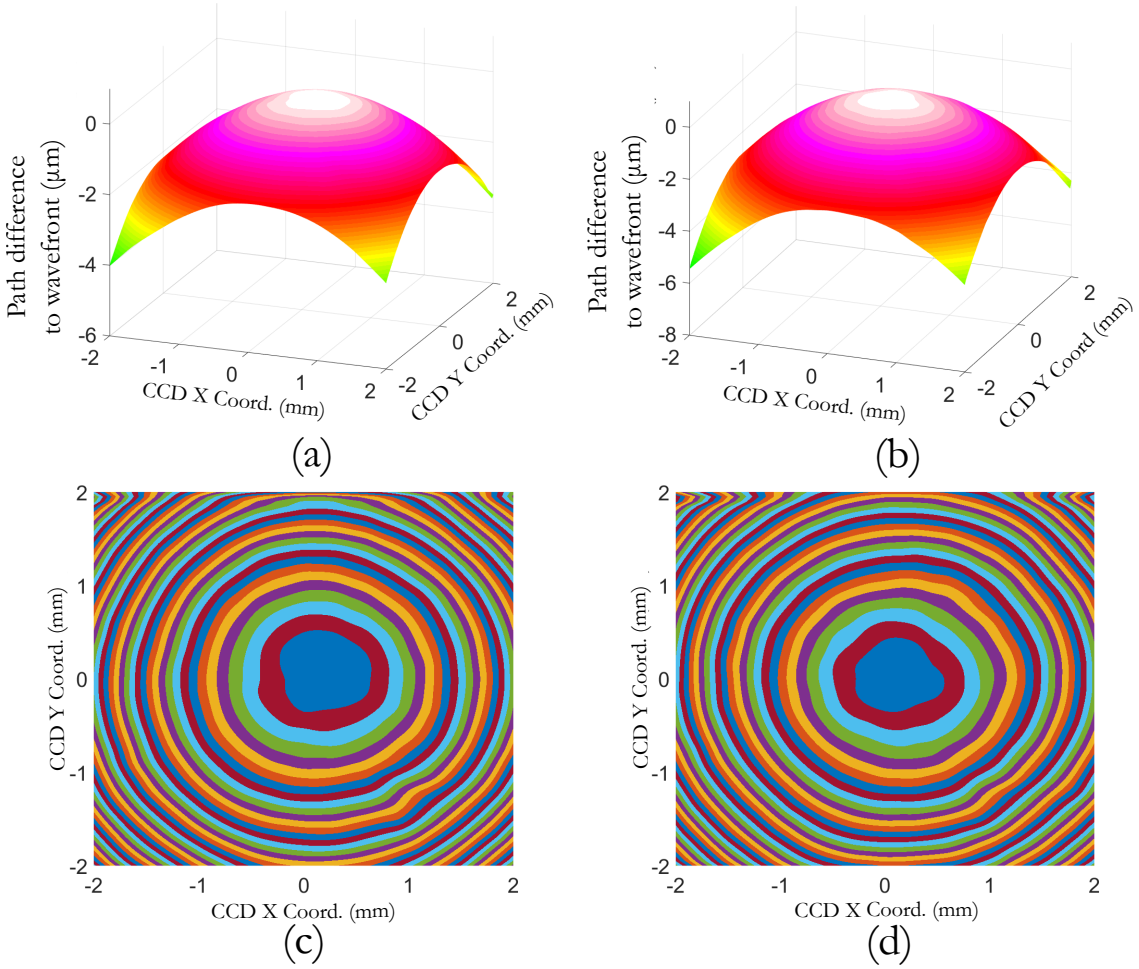


Figure 9.1: The wavefronts of the 70 µm (a and c) and 85 µm PCF (b and d) taken by the Thorlabs wavefront sensor (WFS150-5C), and with another laser's spherical wavefront used as a reference. (a) and (b) show the wavefronts of the two fibres to highlight the difference in wavefront curvature, which was due to the lenses used to collimate the beams onto the wavefront sensor. (c) and (d) use a periodic colourmap on bands of equal path difference to highlight the small perturbations in the wavefront, which were not obvious in (a) and (b).

For two perfectly spherical wavefronts of different radii of curvature, the residual between the two would have the coloured rings marking $\pi/4$, $\pi/2$, π and $3\pi/2$ following perfectly circular contours. Figure 9.2 shows wavefront perturbations in both beams, leading to phase variations across the beam and creating non-circular contours. The non-circular contours result in a spatially varying polarisation state made up of linear and varying degrees of elliptical polarisation. Propagating a combined beam with a spatially varying polarisation state through a polarising beam splitter (PBS) will produce leakage into a rejected beam. To quantify the impact of the wavefront differences and distortions, wavefronts can be decomposed into a linear combination of orthogonal modes called Zernike polynomials. These are a sequence of orthogonal polynomials defined upon a unit disk, as shown in Figure 9.3. An observed wavefront can be decomposed into a series of terms via a Zernike transform

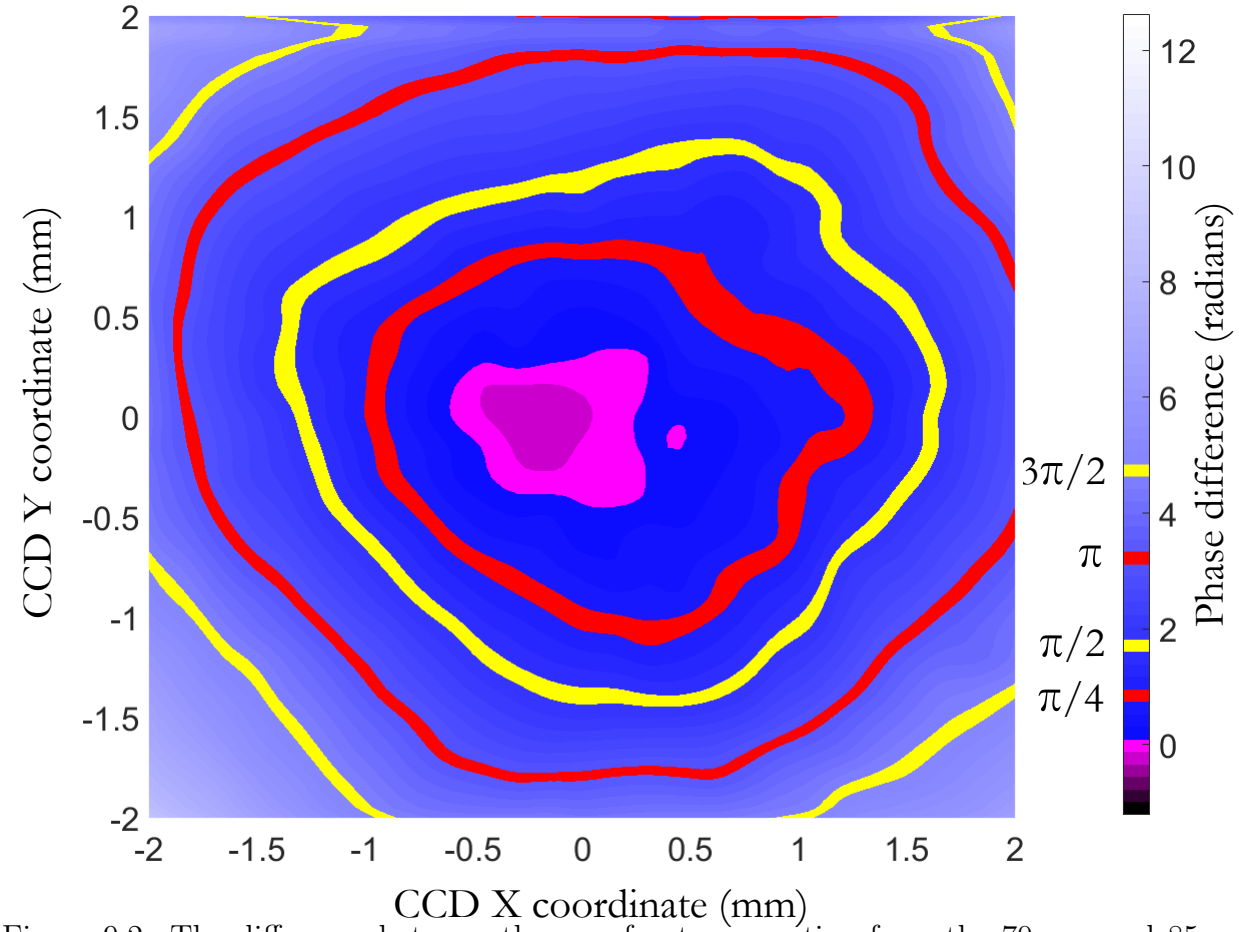


Figure 9.2: The difference between the wavefronts emanating from the 70 µm and 85 µm PCF converted into a phase using Equation (7.5). Coloured contours have been added, which represent the point at which the two beams combine with a phase difference of $\pi/4$, $\pi/2$, π and $3\pi/2$. The non-circular contours result from the perturbations acquired by the wavefronts. If the beams had perfectly spherical wavefronts, then the plot would show circular contours.

[129, 130]. These terms are described by the Zernike polynomials

$$Z_n^m(\rho,\phi) = R_n^m(\rho)\cos(m\phi), \tag{9.1}$$

$$Z_n^{-m}(\rho,\phi) = R_n^m(\rho)\sin(m\phi), \tag{9.2}$$

where there are both even ($Z_n^m(\rho,\phi)$) and odd ($Z_n^{-m}(\rho,\phi)$) polynomials. Both of which are distinguished by their cosine and sine terms respectively. The radial polynomials are described by

$$R_n^m(\rho) = \sum_{k=0}^{\frac{m-n}{2}} \frac{(-1)^k(n-k)!}{k!(\frac{n+m}{2}-k)!(\frac{n-m}{2}-k)!}\rho^{n-2k}, \tag{9.3}$$

where n and m are integers used to select modes with the condition $n \geq m$ and ρ is the normalised radial distance $0 \leq \rho \leq 1$. Zernike polynomials are limited to the range of $|Z_n^m(\rho,\phi)| \leq 1$.

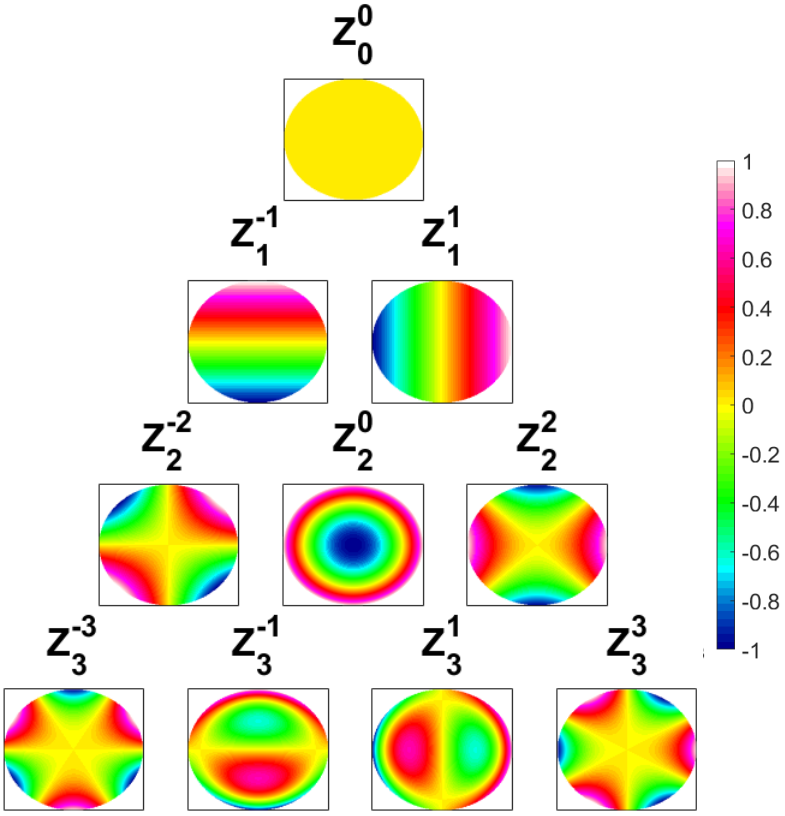


Figure 9.3: The first 10 Zernike modes. The lowest order mode is the piston phase and it is a flat plane. Z_1^{-1} and Z_1^1 are the tip / tilt errors that are seen on the wavefront. Z_2^{-2} and Z_2^2 are astigmatic contributions, while Z_2^0 is the defocusing error, i.e. the curvature of the wavefront. Higher order terms give coma and higher orders of astigmatism. The code used to generate these Zernike polynomials is available online [128], and written by Paul Fricker.

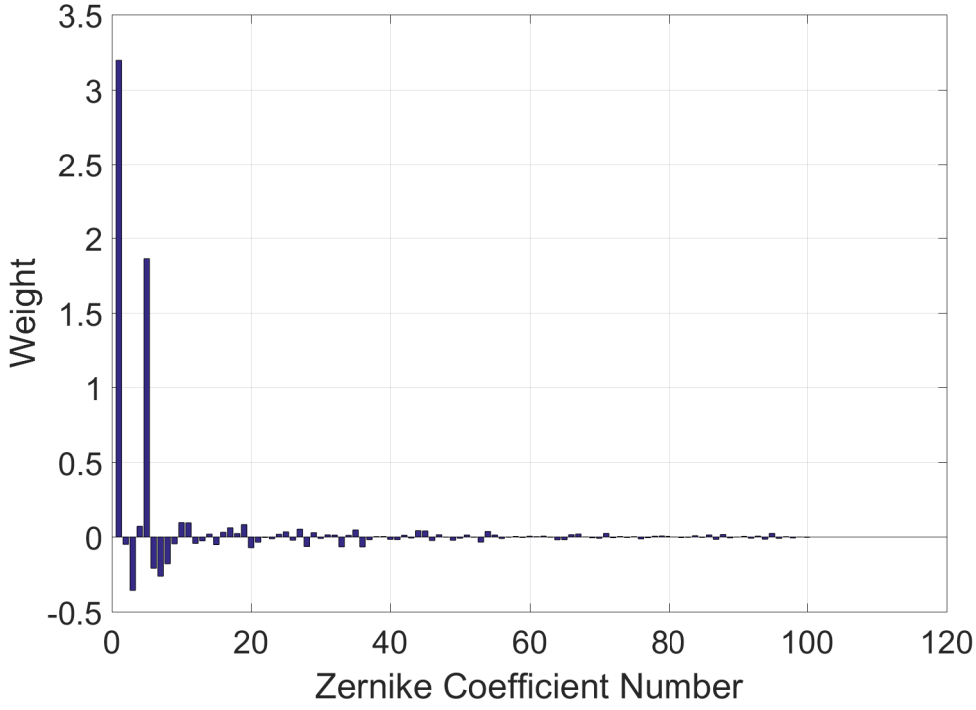
To represent any real wavefront, WF, a linear combination of these terms can be taken

$$\mathrm{WF}(\rho,\phi) = \sum_{\mathrm{m,n}} \left[\mathrm{a_{m,n}} \mathrm{Z_n^m}(\rho,\phi) + \mathrm{b_{m,n}} \mathrm{Z_n^{-m}}(\rho,\phi) \right], \tag{9.4}$$

where $\mathrm{a_{m,n}}$ and $\mathrm{b_{m,n}}$ are the coefficients of the even and odd terms respectively. To generate the coefficients, one calculates the overlap integral of each Zernike mode with the wavefront

$$\mathrm{a_{m,n}} = \int_0^1 \int_0^{2\pi} \mathrm{WF}(\rho,\phi) \mathrm{Z_n^m}(\rho) \rho \mathrm{d}\rho \mathrm{d}\phi. \tag{9.5}$$

Using Equation (9.5), the phase difference map in Figure 9.2 was decomposed into a linear combination of Zernike polynomials, using a MATLAB library which was available online [128] and written by Paul Fricker. This MATLAB script calculates Equation (9.5). The coefficients for this decomposition are shown in Figure 9.4. Consequently, it was possible to ascertain the type of perturbations the wavefront possessed and determine the impact of these perturbations on the degree of linear polarisation (DOLP) and fraction of available



m coefficient

n coefficient	-13	-12	-11	-10	-9	-8	-7	-6	-5	-4	-3	-2	-1	0	1	2	3	4	5	6	7	8	9	10	11	12	13
0														1													
1													2		3												
2												4		5		6											
3											7		8		9		10										
4										11		12		13		14		15									
5									16		17		18		19		20		21								
6								22		23		24		25		26		27		28							
7							29		30		31		32		33		34		35		36						
8						37		38		39		40		41		42		43		44		45					
9					46		47		48		49		50		51		52		53		54		55				
10				56		57		58		59		60		61		62		63		64		65		66			
11			67		68		69		70		71		72		73		74		75		76		77		78		
12		79		80		81		82		83		84		85		86		87		88		89		90		91	
13	92		93		94		95		96		97		98		99		100		101		102		103		104		105

Figure 9.4: The magnitude of the Zernike coefficients ($a_{n,m}$), when the measured phase difference wavefront was decomposed into a linear combination of 105 Zernike polynomials, limited to 105 by the algorithm used. The simplest way to represent the coefficients was in a 2D bar graph, as a 3D plot made it difficult to properly visualise the coefficients and their weights. The translation of the Zernike coefficient numbers, from the graph to their m and n indices, is shown in the table.

beam power. Figure 9.4 shows that the dominant contributions came from lower order terms, these being Piston a_1, tilt a_2 and a_3, defocus a_5 and astigmatism a_6.

The large defocus term seen in Figure 9.4 was almost certainly due to the different lenses used by each beam to fill the wavefront sensor. The tilt (a_2,a_3) terms could also be down to the alignment of the beam with respect to the wavefront sensor, therefore these three contributions were ignored. The remaining Zernike polynomial terms, except for the piston term, were assumed to be due to the difference in the wavefronts of the two PCFs. The piston

term seen in Figure 9.3, was present due to how the code functions. Zernike polynomial modes are all normalised to ±1. In order to decompose a wavefront, it too should be normalised to the range of ±1. The phase map for the wavefront residual seen in Figure 9.2 sat between ∼ 0 and 8, meaning it would normalise to between 0 and 1. The piston phase a_1 contribution seen in Figure 9.4 was a correction factor used by the algorithm to shift the wavefront residual seen in Figure 9.2 to between ∼ ±4, to allow it to be normalised to ±1. As a sanity check, the Zernike polynomials and their calculated coefficients should be able to reconstruct the phase difference map using Equation (9.4). Figure 9.5 shows the comparison between the measured phase difference map and the reconstructed phase difference map.

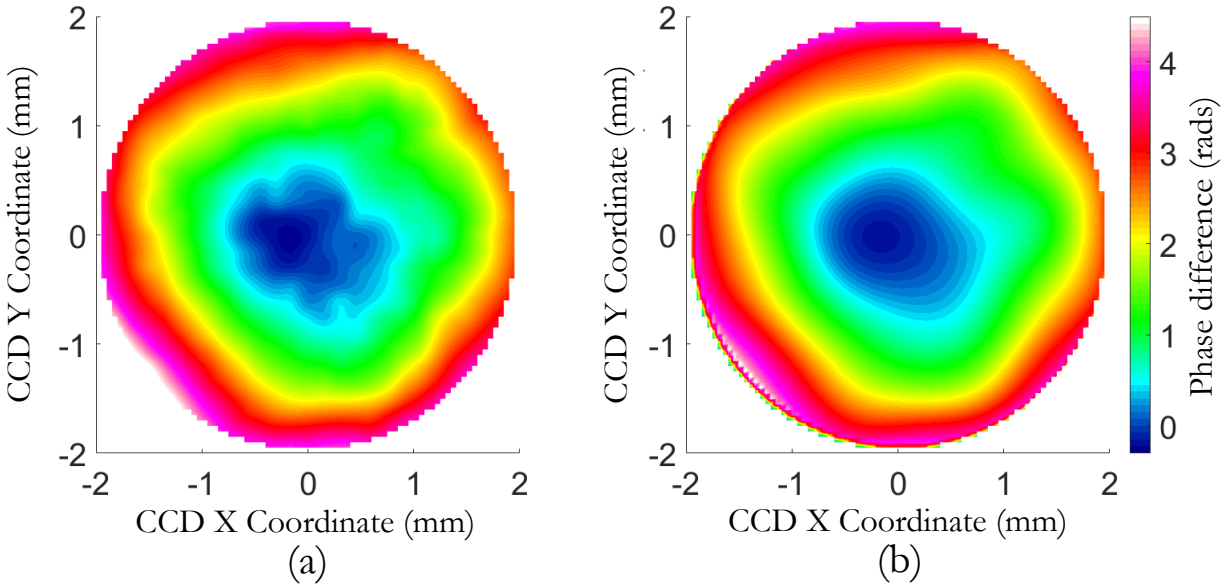


Figure 9.5: A comparison between (a) the measured phase difference colour map and (b) the reconstructed phase difference map, generated from the summation of the 105 Zernike polynomial modes, using the coefficients shown in Figure 9.4. The limitation to 105 Zernike polynomial modes is a limitation of the software library used in MATLAB.

Figure 9.5, shows that the Zernike polynomials can reconstruct the phase difference map, with good agreement to the measured phase difference map. However, the measured map does contain some fine structure that could not be reproduced in the reconstructed map, due to the limited number of Zernike polynomials used. The difference between these two maps produces a residual, that highlights this fine structure, as shown in Figure 9.6.

Using the definition of the fraction of available beam power from Equation (1.10), the fraction of available beam power could be estimated by taking the phase difference map and calculating the power that would be transmitted through, and rejected from, a polarising

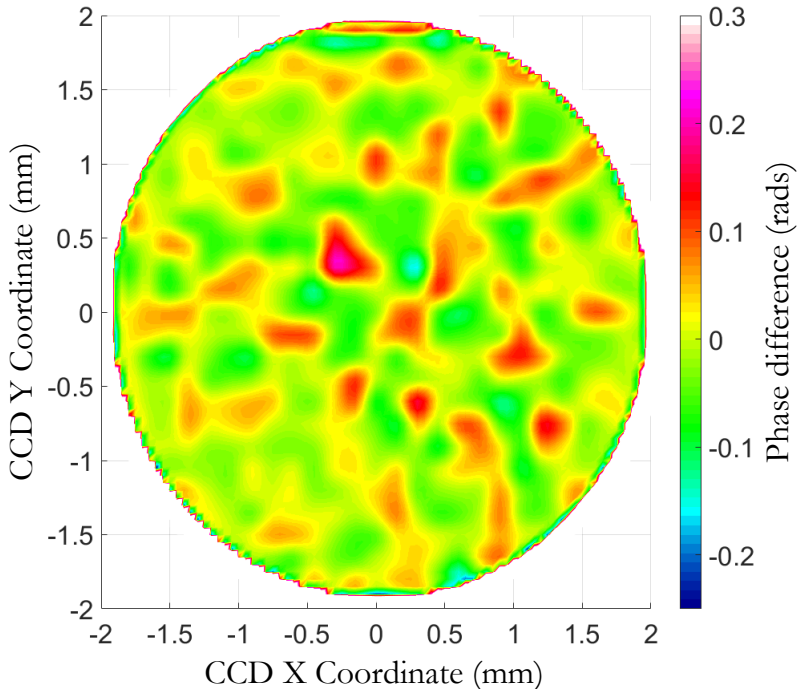


Figure 9.6: The phase difference residual that was produced when the phase difference maps seen in Figure 9.5, are subtracted from one another. This was the fine structure that could not be reproduced by the 105 Zernike polynomials available.

beam splitter. This was done by taking the phase difference map of 9.5a, leaving the phase as it was for the transmitted beam, and then adding a phase of π at all locations for the rejected beam. The effect on the fraction of available beam power from the fine structure, as seen in Figure 9.6, was calculated to ascertain the upper limit on the fraction of available beam power. The fraction of available beam power was calculated to be 99.995%, with the associated DOLP calculated to be 99.9%. Given the negligible impact these perturbations have had on the fraction of available beam power, any major power loss must be due to the contributions of the Zernike modes, which are seen in Figure 9.4. The next task was analysing the impact of the Zernike polynomial modes, starting with the most dominant modes in the Zernike series expansion (omitting piston phase for reasons given earlier), as seen in Figure 9.4.

To get a realistic measure of the fraction of available beam power, the reconstructed phase difference map of figure 9.5b was weighted with a 4 mm diameter ($1/e^2$) Gaussian beam (to fill the unit circle) positioned with its centre at (0,0). This weighting ensured that the areas with high intensities and larger phase differences, affected the fraction of available beam power more than those areas where there were small phase differences and lower intensities.

Using the reconstructed phase difference map, it was possible to remove terms from the series expansion and observe the effect they would have on the fraction of available beam power. Starting with the full expansion of terms, modes were subtracted sequentially and cumulatively from the total simulated phase map, starting from Mode 2 and working up to Mode 20. After each mode removal, the fraction of available beam power and the DOLP were calculated and plotted on a graph, as shown in Figure 9.7.

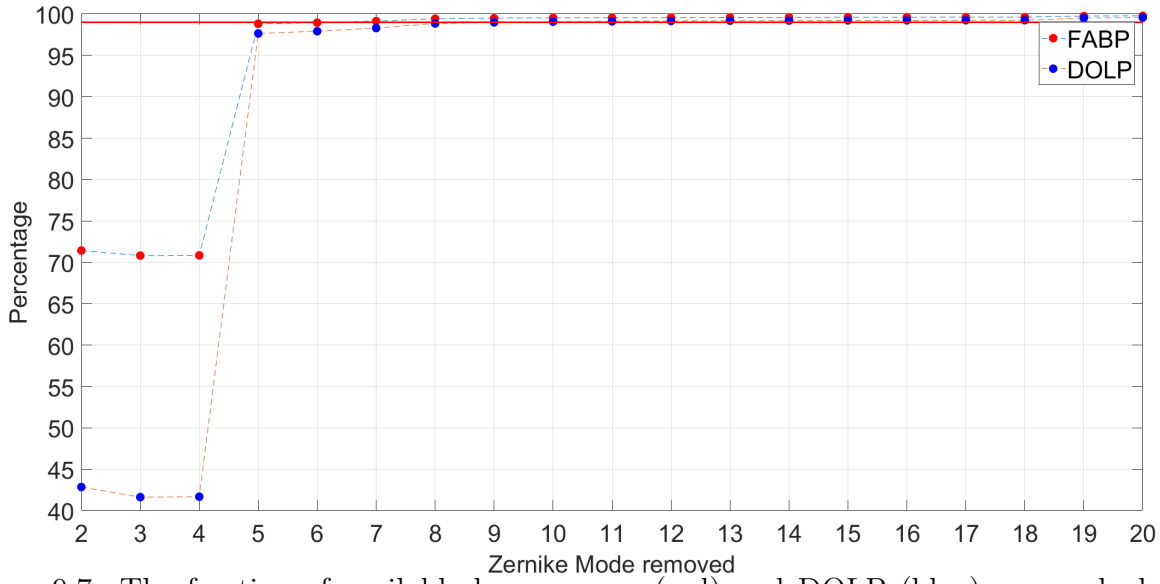


Figure 9.7: The fraction of available beam power (red) and DOLP (blue) were calculated after the removal of each mode from the simulated phase map. The red solid red line indicates the 99% level. The removal of modes was sequential from mode a_2 (tip / tilt) upwards, and their effects were cumulative as more modes were removed. As can be seen, after mode a_5 (defocus term) was removed, there was a large jump in the FABP up to 98.6%.

From Figure 9.7, it was concluded that the removal of the first ten modes was sufficient to increase the fraction of available beam power and DOLP to over 99 %. The largest contribution came from the defocus term a_5, which improved the fraction of available beam power from 74.8 % to 98.6 %. Looking solely at the effect on the fraction of available beam power and DOLP, in the absence of mode a_5 (defocus), Figure 9.8 shows the effects of sequential mode removal.

What can be seen from Figure 9.8 is that the next largest contributors to the reduction of the fraction of available beam power were modes a_6 (astigmatism), a_7 (higher order astigmatism) and a_8 (coma). The effects of astigmatism and coma were difficult to eliminate in the absence of carefully designed lenses to meet individual beam requirements. If no corrections were possible, these terms would set an upper limit to the DOLP and the fraction of available

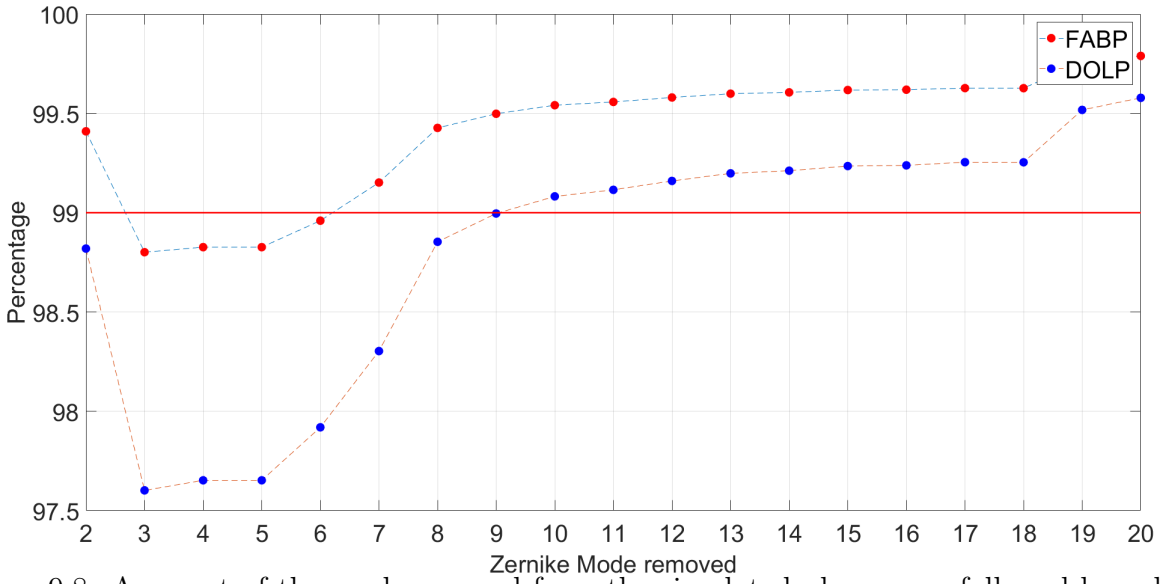


Figure 9.8: A repeat of the mode removal from the simulated phase map followed by calculating the fraction of available beam power (red) and DOLP (blue). The mode removal was sequential and cumulative, with the contribution from mode a_5 (defocus) removed to allow better resolution of the effects caused by the other terms. The next largest contributors to the FABP were modes a_6 (astigmatism), a_7 (higher order astigmatism) and a_8 (coma).

beam power possible from a coherent combination system. Modes a_2 and a_3 (Tip and tilt) did not have a large impact on the simulated fraction of available beam power. However, much effort was expended to ensure the beams were properly aligned to the wavefront sensor, and the on-screen software package provided feedback that assisted in removing a large quantity of tip and tilt. The wavefront sensor was not used for the coherent combination of beams, to avoid causing unnecessary damage. Instead, the alignment method as discussed in Section 7.1 was used to make the beams co-propagating. Even though the effect of tip and tilt was minimised through careful alignment of the two combining beams, near-perfect alignment was difficult to achieve using this alignment method. It was suspected that in these experiments, there might have been a non-negligible tip/tilt contribution leading to a drop in the fraction of available beam power in the high power experiment. The method used to align the beams to be co-propagating depended upon observing the lateral motion of the interference pattern to ascertain which direction to make corrections to the beam pointing. If there was a lack of observable lateral motion on the interferogram (the combined beam looked like a pulsating Gaussian spot), then no further corrections could be made, creating uncertainty in the tip/tilt contributions.

Under further scrutiny, it was considered that the alignment method discussed in Section 7.1 had two limitations: the precision to which the two Gaussian beams could be overlapped in the far field; and the ability to observe the subtle transverse shifts of the interference pattern on the camera that indicated a tilt between the beams. The precision to which the beams could be overlapped on a camera was a combination of two factors: the pixel size on the camera; and the user's interpretation of what constituted the best spatial overlap. For the pixel size, the beams could not be overlapped to better than one pixel on the camera. For the latter, the two Gaussian beams may have had intensity profiles which were different, perhaps slightly elliptical, and with the point of highest intensity slightly off their geometrical centre (the centre of mass/centroid of the beam's spatial profile). The user could choose to overlap the beam profiles for geometric overlap with the beam's maxima being slightly offset, or overlap the points of highest intensity. For this thesis, it was assumed that the geometrical centre of the beam was also the centre of the wavefront. So, the best overlap of the beam intensity profiles would correspond to the overlap of the geometric centres.

Another source of error was the user's perceived location of the centre of the beam. The cross-hair from the beam profiling software was used to mark the centre of the reference beam before that beam was blocked. The cross-hair was left fixed at that position on the screen, and the alignment beam moved and centred on the cross-hair by eye. Given that the beam profiles had no fixed points on them, which would allow the two beams to be overlapped repeatably to within one pixel, the user had to estimate the centre of the beam. Figure 9.9 shows that this was not easy to do on a beam profiling system, where the cross-hairs must be placed manually to prevent the software package re-optimising the centroid when the beam disappears. Assuming a beam centroid offset-uncertainty of $\sim$50 µm at the camera and the distance from the combination cube to the beam profiling camera being 0.875 m, the angular uncertainty between two co-propagating beams could be calculated. The angular uncertainty between the two co-propagating beams also defined the tilt between the two matched wavefronts. Using geometry, the difference in path length between the two wavefronts, at all points across the beams spatial profile, could be calculated and converted into a phase using Equation (7.5), see Figure 9.10.

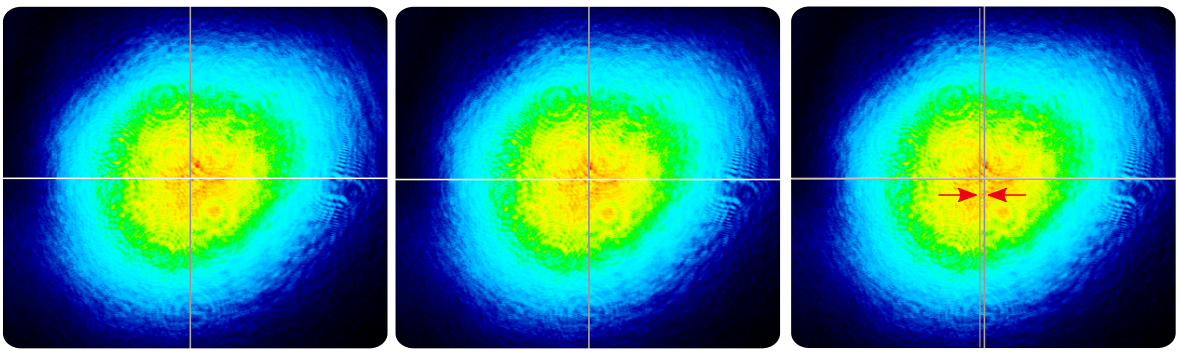


Figure 9.9: The 70 μm beam used in the low power regime experiment, with two cross-hairs aligned by eye over the beam's geometric centre. Even though the cross-hairs appeared to be central, it can be seen that a small difference exists between their placement. This small mismatch was approximately 4 pixels ($\sim$ 13 μm pixel size on the camera used), which equated to a total offset of $\sim$ 52 μm.

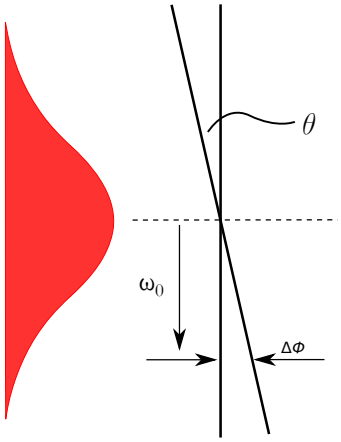


Figure 9.10: A representation of the beam size, with a simple example showing two wavefronts tilted at an angle of θ shown right. The path difference can be calculated at all points and converted into a phase difference $\Delta\phi$, as a function of position on the beam. In this 1D example, the position is represented as the beam radius from the centre of the beam, ω_0.

Using the $\sim$ 5 mm beam diameter, as measured by the high power data run, the phase difference calculated at one beam radius, was found to be $\Delta\phi \approx 0.91$ radians. With a known phase difference at one beam radius away from the centre of the beam, the tip/tilt mode a_5 could be used to reconstruct a phase difference map, where the same phase difference could be seen at the same point. a_2 was scanned from 0-2 and the phase difference at one beam radius was calculated and plotted into a graph, as seen in Figure 9.11. The point where the phase difference was equivalent to $\Delta\phi \approx 0.91$ at one beam radius was when the coefficient was $a_2 = 1.05$. Using the calculated tip/tilt a_5, the impact on the fraction of available beam

power and DOLP could be calculated. A Gaussian beam profile with a 5 mm beam diameter was used to model the distribution of the beam power and weight the phase differences at all points. This was done using using Equation (9.6) [95]

$$\mathrm{P(x,y)} = \mathrm{P_0(x,y)}(1 + \cos(\Delta\phi(\mathrm{x,y}))) \tag{9.6}$$

where P(x, y) is the power of the beam observed through a PBC at point (x, y), $\mathrm{P_0(x,y)}$ is the input beam power at point (x, y), and $\Delta\phi$(x, y) is the phase difference between the combining wavefronts at point (x, y). The power in the desired polarisation state depended upon the phase difference between the two wavefronts of the beam, $\Delta\phi(x, y)$. The measured or reconstructed phase difference map could be substituted into this equation along with an array describing $P_0(x, y)$ to allow calculation of P(x, y). The difference in power between the desired and rejected beams could be compared by using a phase difference of $(\Delta\phi)$ for a piston phase of zero and $(\Delta\phi + \pi)$ for a piston phase difference of π, and substituting both into Equation (9.6). This gave $\mathrm{P_{max}}$ and $\mathrm{P_{min}}$ respectively, which could be used to calculate the fraction of available beam power and DOLP. Using $\mathrm{a_2} = 1.05$ to calculate the phase differences at all points, the fraction of available beam power and DOLP were calculated to be $\sim 93\%$ and $\sim 86\%$ respectively.

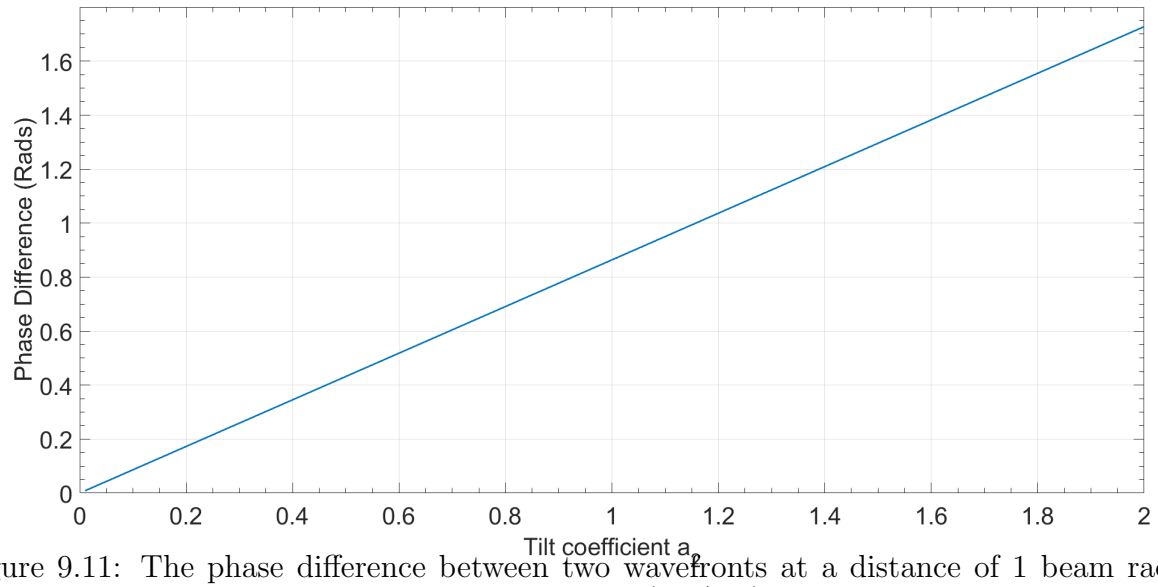


Figure 9.11: The phase difference between two wavefronts at a distance of 1 beam radius from the centre of the beam, when either $\mathrm{a_2}$ or $\mathrm{a_3}$ (tip/tilt) are scanned from 0-2.

The calculated numbers for the fraction of available beam power and DOLP, given a tip/tilt coefficient of $a_2 = 1.05$, were very similar to that observed experimentally for the high power experiment, and only a few percent different for the low power experiment. To see if this amount of tilt was feasible, the previous wavefront model, seen in Section 7.1, was used to simulate two beams interfering with a crossing angle equivalent ~ 50 μm geometrical centre offset at the camera. The amount of tilt between the two beams leads to an interference pattern, which moves laterally whilst its intensity changes. As previously mentioned, one metric of how well the two beams were aligned, was to observe the interference pattern as the piezo mounted mirror scanned, and look for the lateral movement of the combined beam's interference pattern. During alignment, it was found that the two beams could be aligned using the camera, to the point where none of this lateral motion could be easily observed, only the global intensity variation of the beam. The model showed that, with 50 μm of offset between the beams, and by plotting the intensity profiles to screen, small lateral motions of the combined beam interferogram were visible. As this motion could be seen in simulations but was absent from experimental observations, it meant that although the beams could have an uncertainty of 50 μm on the camera due to non-parallel beams, the beams were likely aligned to better than 50 μm.

It was found that a 20 μm offset provided a barely noticeable amount of lateral motion of the simulated intensity profile on screen. The phase difference at one beam radius from the centre due to this 20 μm error was $\Delta\phi = 0.35$, corresponding to $a_2 = 0.4$ which was 2.5 times less than the $a_2 = 1.05$ from before. The fraction of available beam power and DOLP resulting from an $a_2 = 0.4$, was calculated to be $\sim 99\%$ and $\sim 98\%$ respectively. These figures suggested that the alignment method by eye, using a camera and cross-hair to overlap the beam's spatial profiles, was not the cause for the observed fraction of available beam power and DOLP value seen in the low and high power regime experiments. Therefore, the alignment method used throughout this thesis was deemed adequate for use in the coherent combination of beams.

A residual defocus a_5, caused by a mismatch between the wavefront radii of curvature, was the next likely culprit for the observed values of the fraction of available beam power and

the DOLP seen in the low and high power experiments. If the defocus term a_5 was assumed to be the main driver behind beam power loss, by performing a numerical parameter scan of the defocus coefficient a_5, the fraction of available beam power and DOLP could be calculated for every possible value of a_5. Once the value of a_5 was found, which produced the fraction of available beam power and DOLP as seen in the high power regime experiment, the radius of curvature difference between the wavefronts of the combining beams could be calculated. Using the same 5 mm beam size, the defocus coefficient was scanned from 0-2, as shown in Figure 9.12, and the fraction of available beam power and DOLP were calculated using Equation (9.6).

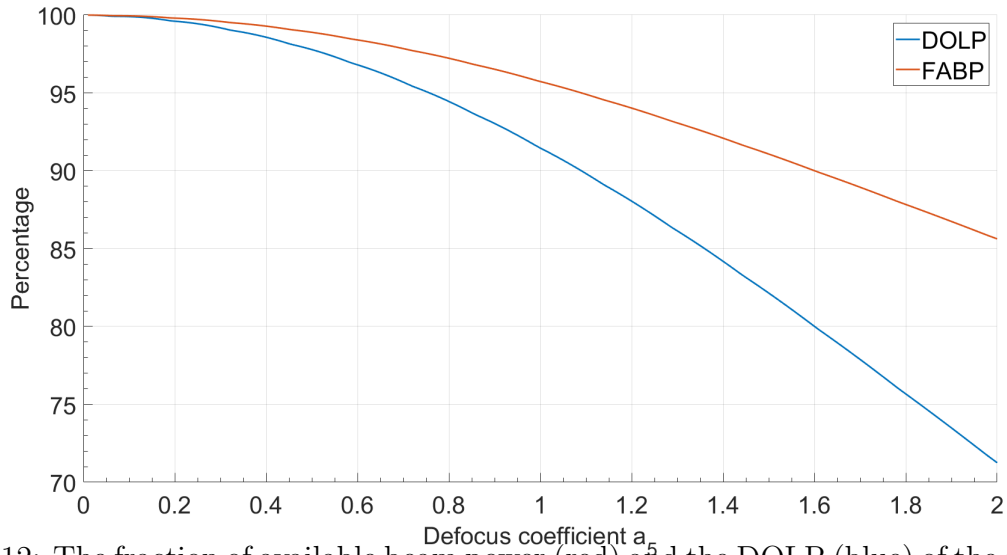


Figure 9.12: The fraction of available beam power (red) and the DOLP (blue) of the simulated combined beam when a defocus term a_5 was introduced with varying magnitude. To match the observed values of the fraction of available beam power and DOLP as seen in the high power regime experiment (93% and 86% respectively), a defocus coefficient of (a_5 1.3) would be required.

Figure 9.12, showed that the defocus term a_5 could have a large impact on the fraction of available beam power and DOLP, as the coefficient grew. Also, the system was able to tolerate a defocus term of $a_5 \approx 0.47$ and only suffer a 1% reduction in the fraction of available beam power. However, as the defocus term a_5 increased, the rate of degradation of the DOLP and fraction of available beam power increased also. A defocus coefficient of $a_5 \approx 1.3$ was found to be necessary to produce a fraction of available beam power and a DOLP of 93% and 86% respectively. Using $a_5 \approx 1.3$, the phase difference seen between two wavefronts at one beam radius away from the centre of the beam is shown in Figure 9.13. Using Equation (7.5),

the phase difference between two wavefronts with different radii of curvature was determined to be

$$\Delta\phi(\mathrm{x},\mathrm{y}) = \frac{2\pi}{\lambda} * \left(\mathrm{R}_1 - \mathrm{R}_2 - \sqrt{\mathrm{R}_1^2 - \mathrm{x}^2 - \mathrm{y}^2} + \sqrt{\mathrm{R}_2^2 - \mathrm{x}^2 - \mathrm{y}^2}\right), \tag{9.7}$$

where R_1 and R_2 are the radii of curvature of the two wavefronts and x and y are the transverse coordinates along the wavefronts. Equation (9.7) could then be rearranged to solve for R_2 in terms of R_1 and $\Delta\phi(\mathrm{x},\mathrm{y})$. Figure 9.14 was produced by scanning over R_1 and inserting the phase difference produced by a defocus term of $\mathrm{a}_5 = 1.3$ to find R_2.

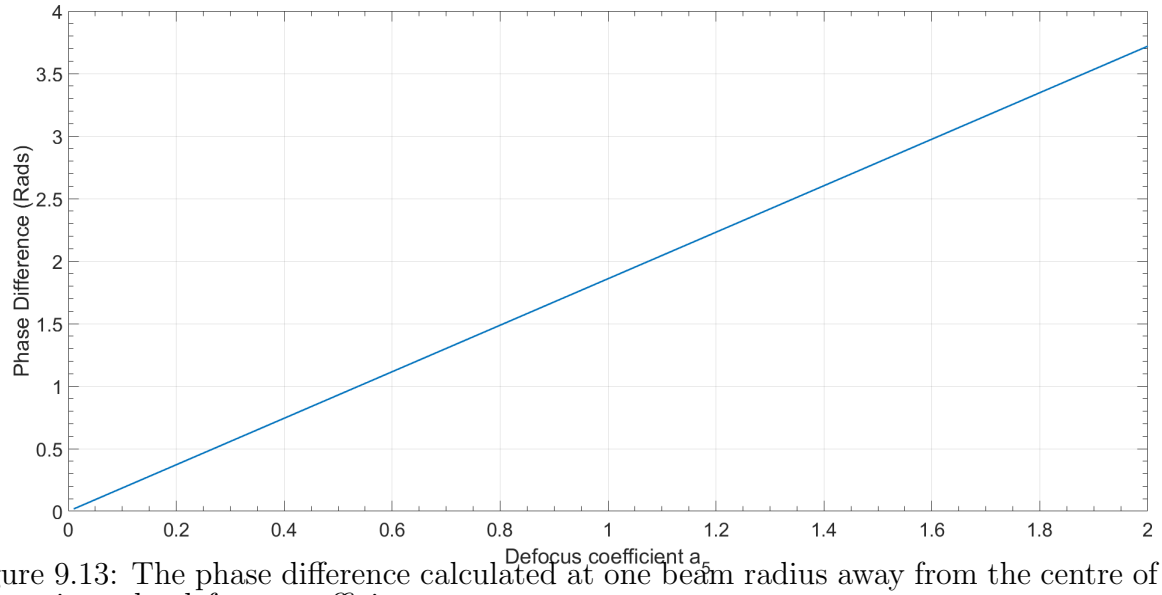


Figure 9.13: The phase difference calculated at one beam radius away from the centre of the beam given the defocus coefficient a_5.

It can be seen, from Figure 9.14 that a defocus term of $\mathrm{a}_5 = 1.3$ could have a large impact on the difference between the wavefront's radii of curvature. Even though the simple telescope model had gone some way to correcting the differences between the wavefronts, if a small residual remained, which was similar to that seen in Figure 9.14, it was possible that beam power was being lost due to an uncompensated defocus term a_5. The wavefront matching model used measurements of the beam size at several locations to infer the beam divergence, and to allow the model to estimate the position and length of a wavefront matching telescope, given two preselected lenses. However, it was possible that the telescope used was not the right length or in the right position (longitudinally). If the length of the wavefront-matching telescope was responsible for the beam power loss, then its effect on the wavefront of a beam could be investigated by scanning the telescope length over a short range around the

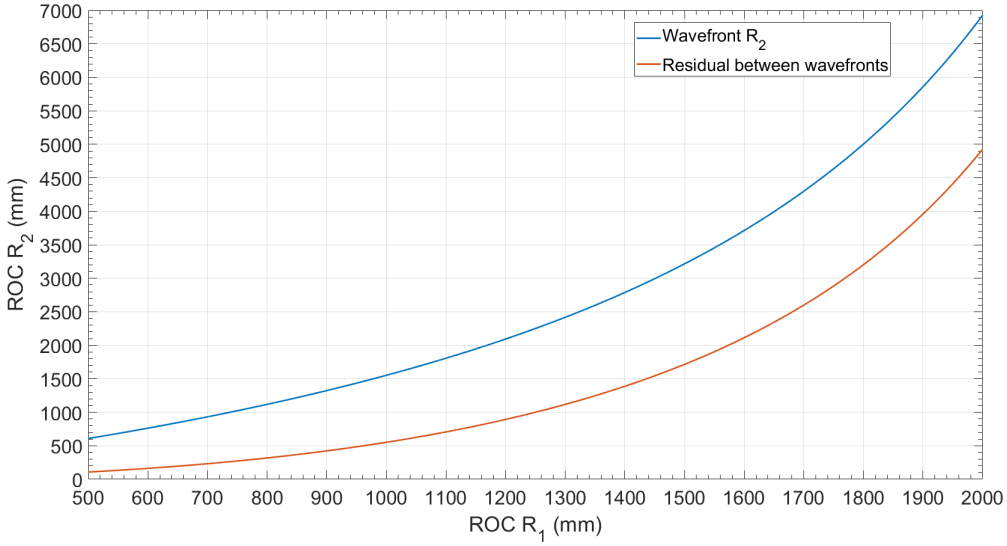


Figure 9.14: The radius of curvature R_2 (blue), given the radius of curvature R_1, and a defocus term of $a_5 = 1.3$, required to generate a fraction of available beam power and a DOLP of 93% and 86% respectively. Plotted in red is the residual between the two wavefronts for a given value of R_1 (using the same vertical scale).

calculated value. A millimetre range was chosen, as this was the smallest increment found on the tape measure used to measure the length of the telescope. The result of this can be seen in Figure 9.15.

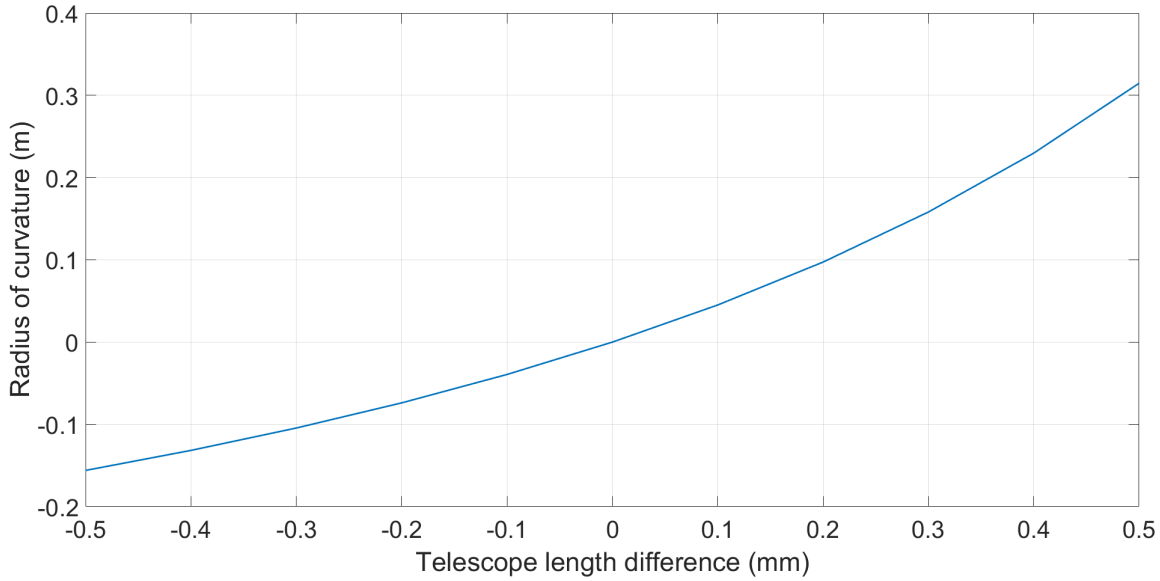


Figure 9.15: The length difference of the telescope scanned from -0.5 mm to 0.5 mm around the predicted correct position. If it was assumed that the telescope was in the ideal position to begin with, then even a deviation from the ideal location by 300 µm, could lead to an increase in the radius of curvature of $\sim$ 160 mm.

Figure 9.15 shows that a small error in the position of the second lens can have a large impact on the wavefront's radius of curvature. The telescope length was measured using a

tape measure, with the intention of fine-tuning the telescope length using a translation stage. The fine tuning involved observing the interference pattern on a beam profiling camera and judging how well the wavefronts were matched. Better matching was seen when less or no interference rings could be observed when the piezo was set to scan over its full range of motion. This objective approach to wavefront matching was sufficient for getting the majority of the beam power. However, for better results more rigorous techniques would be required. It was possible that the telescope not being exactly the right length, created a defocus term sufficient to reduce the fraction of available beam power and the DOLP to 93% and 86% respectively.

9.2 Endcap and path length analysis

For the experiments performed in this thesis, the extra path length used to balance the arms of the interferometer was added after the PCFs, which therefore required the telescope to help correct the divergence mismatch. In hindsight, an alternative method to fixing the divergence mismatch problem might be to remove the telescope and place the extra path length before the PCFs. This would have allowed the path lengths, post-PCF, to be accurately matched from the outset. If the combining pair of PCFs both had endcaps and the same coupling lens, then the differences between the radii of curvature of the wavefronts could be minimised, even though the beam sizes would still be slightly different. This would reduce the amount of wavefront difference that needed to be corrected at the point of combination. A simple ABCD beam propagation model was constructed which assumed that the beam emanating from the core of a PCF was at a waist. Then, taking the mode field diameters of the 70 μm and 85 μm PCFs, four scenarios were considered: both PCFs with endcaps and with the same path length to the combining PBS; both PCFs with endcaps but with a small 10 mm error in the path length from each PCF to the combining PBS; the 70 μm PCF without an endcap, but with a matched path length to the combining PBS; and finally, the 70 μm PCF without an endcap, and a small 10 mm error in the path length from each PCF to the combining PBS. The beam parameters seen at the point of combination for these four scenarios are shown in Table 9.1.

	70 µm EC 85 µm EC PL matched	70 µm EC 85 µm EC 10 mm error	70 µm NEC 85 µm EC PL matched	70 µm NEC 85 µm EC 10 mm error
70µm PCF ROC	1.858m	1.908m	1.671m	1.688m
70µm PCF W_D	5.7 mm	5.1 mm	2.8 mm	2.8 mm
85 µm PCF ROC	1.807m	1.807m	1.807m	1.807m
85 µm PCF W_D	5.0 mm	5.0 mm	5.0 mm	5.0 mm
Calculated DOLP	97.2%	95.9%	71.2%	71.6%

Table 9.1: The beam parameters seen at the point of combination for scenarios (1)-(4). EC - end cap, NEC - No end cap, PL - path length

From this model, it can be seen that the presence of the end cap makes a large difference to the degree of linear polarisation at the point of combination due to the mismatched wavefronts. When only the core diameter differs between the two PCFs, the estimated degree of linear polarisation remains high, sitting at around 96%. However, this high degree of polarisation does assume that the endcaps are the same length, and have the same refractive index. Differences to end cap length and path length travelled to the point of combination could account for a 10% drop in the degree of linear polarisation. With one PCF having no end cap, the degree of linear polarisation can reduce to the $\sim 70\%$ level. With mismatched path lengths to the point of combination, the resulting difference in wavefronts explains why early work realised degree of linear polarisations of $\leq 40\%$.

The next simulation was run, with the 70 µm PCF without an endcap, the 85 µm PCF with an endcap, the 10 mm error in path length, but the coupling lens's position at the exit of the 70 µm PCF given freedom to be adjusted. This adjustment allowed for better matching of the beam parameters at the point of combination, with the simulation predicting a 99% degree of linear polarisation. A core diameter reduction of the 70 µm PCF, to give a mode field area reduction of 30% from $2250 \times 10^{-12}\,\mu m^2$ to $1578 \times 10^{-12}\,\mu m^2$ gave a calculated degree of linear polarisation of $\sim 92\%$. The 85 µm PCF (mode field area of $3300 \times 10^{-12}\,\mu m^2$) could be combined with a $1578 \times 10^{-12}\,\mu m^2$ mode area PCF, with a 10 mm path length mismatch, provided the coupling lens had the freedom to be adjusted. This suggests that the combination of two PCFs, regardless of small core diameter and length differences, can be achieved with high levels of degree of linear polarisation, by matching the path lengths post-PCF.

The lack of an endcap on the shorter 70 μm PCF, meant that eventually, it would be power limited due to breakdown at the end facet from high intensities. This PCF was afflicted by a few other issues, which made working with it difficult. Being older and a completely different design of PCF, it was affected more by thermal effects than the 85 μm PCF. This highlights the need to specify to manufacturers, the power levels at which the PCFs are required to operate, so they can design PCFs to operate stably in these conditions. It would be pointless to have matched PCFs from different manufacturers, if a batch from one of them experienced thermal effects and mode degradation near the desired operating power.

9.3 Summary

The use of Zernike polynomials allowed a greater insight into the effects causing the degradation of the fraction of available beam power and DOLP. A combination of defocus a_5, tip/tilt a_2, a_3 and higher order terms, all contributed to reducing the simulated DOLP of the two measured wavefronts to $\sim 40\%$, as seen in Figure 9.7. Gradual removal of the modes from the reconstructed wavefront showed that the largest impact on the DOLP and fraction of available beam power occurred when the defocus term a_5 was removed. This led to an increase of over 50% for the DOLP and 30% for the fraction of available beam power. The large defocus term a_5 was caused by the difference of the wavefronts striking the wavefront sensor. Each beam used a separate lens to ensure the beam filled the CCD of the sensor. However, the measured defocus term a_5 shown in Figure 9.4, was not a direct indication of the amount of wavefront mismatch present at the point of combination in the coherent combination experiment. Looking at the reconstructed wavefront without a defocus term a_5, allowed scrutiny of the other terms that were assumed to be present at the point of combination. Further improvements to the fraction of available beam power and DOLP, to over 99%, could be achieved by removing the first 10 Zernike modes. However, some of these modes would be difficult to remove, unless custom built optics or a spatial light modulator or deformable mirror were available. The modes which could be controlled without these were the tip/tilt a_2, a_3 and the defocus a_5 terms.

From the modelling presented in this chapter, the contributions of tip and tilt to power losses were deemed to be negligible, due to the inferred accuracy (within 20μm) of overlapping

the beam spatial profiles. It was believed initially that the simple wavefront matching model was sufficient for removing the majority of the wavefront mismatch, due to the defocus term a_5. However, tuning the telescope length by observing the interference pattern of the two beams could lead to a residual error. With a telescope length error of up to 0.5 mm, the difference in wavefronts would be sufficient to cause significant reductions to the DOLP and fraction of available beam power. If a defocus term of $a_5 = 1.3$ still remained after the telescope length tuning, then this could have caused the experimentally observed values of the fraction of available beam power and DOLP of 93% and 86% respectively. The indications from this analysis, are that a residual defocus term a_5 might be present, despite the tuning of the wavefront matching telescope. In future iterations of this experiment, effort should be invested in ensuring that the residual defocus term a_5 is minimised. This minimisation could lead to the largest gains to the DOLP and fraction of available beam power.

An alternative approach was considered for correcting the wavefront mismatch due to the endcaps and mismatched path lengths, from the end of the photonic crystal fibre (PCF) to the point of combination. In simulation, the effect of matching the path lengths and adding end caps on each PCF was found to make a significant difference to improving the wavefront matching of two PCFs. In the extreme case of one PCF having double the mode field area of another, the degree of linear polarisation was found still to give values of over 90 %. Before wavefront matching optics are considered in future experiments, gains may be made more economically by ensuring the path lengths from the PCF to the point of combination are matched, and both PCFs have end caps.

Chapter 10

Conclusions and Future work

10.1 Conclusions

To the best of my knowledge, this is the first demonstration of two significantly mismatched photonic crystal fibres (PCFs) being coherently combined together. The thermal behaviour of the PCFs under different amounts of pumping meant that the experiment was operated in two different power regimes. The low power regime had the purpose of combining the PCFs at low enough powers to achieve gain, therefore demonstrating the principle of coherent combination of mismatched PCFs. This was achieved, with the system demonstrating a system efficiency of $85.8 \pm 0.5\%$, a degree of linear polarisation of $83.3 \pm 0.3\%$, a fraction of available beam power of $91.6 \pm 0.3\%$ and a peak power of 1.55 ± 0.2 MW, given a compressed pulse duration of 370 ± 30 fs and 75% of the energy under the central peak in the time domain. The high power regime was also successful, with the system demonstrating a system efficiency of 88.0%, a degree of linear polarisation of 86%, a fraction of available beam power of $93.0 \pm 0.5\%$ and a peak power of 5.84 ± 1.3 MW, given a compressed pulse duration of 380 ± 60 fs and 70% of the energy under the central peak in the time domain.

In related work demonstrated in the field, a combining efficiency of 88% was demonstrated for a 22 GW peak power pulse using 4 PCFs [38]. A system efficiency of 78% was demonstrated for a 12 mJ combined pulse using 8 PCFs and divided pulse amplification [42]. However, these systems utilised more PCFs and lower repetition rates to achieve higher pulse energies. The amount of non-linear phase acquired by Klenke's PCFs was approximately two

orders of magnitude larger than that seen in the high power regime from this thesis, though their experimental setup differed in many ways. Firstly, their system started with a 60 fs pulse (3 times shorter than in this thesis) which was stretched to a 2 ns pulse (6 times longer than in this thesis. Their PCFs also had core diameters of 135 μm , and were capable of producing pulse energies of 2.2 mJ. Their subsequent 8 PCF experiment demonstrated 12 mJ pulse energies from fibres that had similar core diameters to those used in this thesis. The similarities between the PCFs motivated the modelling of my mismatched PCFs under the same conditions as used previously (conditions as seen in Section 6.2 using the split step Fourier method), except for the PCF input and output pulse energies. These were set to 3 μJ and 600 μJ respectively for both the 70 μm and 85 μm PCFs.

Despite the large pulse energies and the poor combined-pulse autocorrelation, the coherent combination of the pulses still gave a predicted degree of linear polarisation of 92%. However, the simulated non-linear phase for the 70 μm and 85 μm PCFs, was 20.7 and 19.2 radians respectively, highlighting why the autocorrelation seen in Figure 6.15 was so poor. To bring the non-linear phase into line with the non-linear phase of 5 radians acquired by the PCFs in the 8 PCF coherent combination experiment [42], the output pulse energy was reduced. Trialling a few different energies, it was found that an amplified pulse with 100 μJ of energy gave simulated non-linear phases for the 70 μm and 85 μm PCFs of 5.2 and 4.8 radians respectively. The degree of linear polarisation and fraction of available beam power of the combined pulse at the reduced pulse energy, was calculated to be over 99%. The combined pulse autocorrelation seen in Figure 6.16, also showed significant improvement, with a deconvolved pulse duration of 400 fs. Assuming that a pulse of 100 μJ and 400 fs duration was achievable using the 70 μm and 85 μm PCFs, the peak power possible for a 200 μJ combined pulse with 75 % of the energy under the central peak, would be 375 MW. Assuming a larger number of fibres were used, with the same pulse energies, these numbers scale to 750 MW for 4 PCFs and 1.5 GW for 8 PCFs. However, in this thesis, the main limitation to further amplification was single mode degradation and, beyond this, the amount of available pump power. Given sufficient mode stability, as the pump power increased, the system efficiency measured would remain stable until larger non-linear phase differences between the PCFs emerged. These non-linear phase differences would need to be balanced by calculating

the non-linear phase accumulated by the pulse in each PCF. By carefully controlling the amplification and the pulse's polarisation state, the non-linear phase could be tuned to the desired value. Non-linear phase differences between many PCFs are not easy to control and may explain why the previous 4 [38] and 8 [42] PCF combination experiments saw system efficiencies lower than the high power regime experiment of this thesis.

Continuing with the comparison to the 4 and 8 PCF experiments [38, 42], the fraction of available beam power calculated in the high power regime can be used to predict the performance of 4 and 8 mismatched PCFs. The metric used by the above authors is the system efficiency [52], as described by Equation (1.9). However, what is not mentioned by the authors across several publications, is the degree of linear polarisation of the individual PCFs, or that of the combined beam. Therefore, it is difficult to make a meaningful comparison between their system efficiency metric, and the fraction of available beam power metric used in this thesis. It was assumed that their efficiency numbers related only to the polarised beam powers which actually contributed to the coherent combination. It was also assumed that the total power of each beam, prior to combination, was made up of a linearly polarised component, which was desired and an elliptical/orthogonal linearly polarised component, which formed a rejected beam. These assumptions then make Klenke's definition of system efficiency [38, 42] fit with the definition of the fraction of available beam power used in this thesis. Using Equation (4.37), for a fraction of available beam power of 93%, we see that for 4 and 8 PCFs respectively, the total fraction of beam power that makes it into the final beam is 86.5% and $\sim$ 80% respectively. In their work, the authors have reported for the 4 and 8 PCF experiments, efficiencies of $\sim$88% and $\sim$78%. These efficiencies are very similar to those calculated using the measured fraction of available beam power from the high power regime.

However, it is interesting that the system efficiency of the 8 PCF combination experiment was as low as it was, given the care that was taken to match the PCFs. Simulations seen in Section 6.2 using the general non-linear Schrödinger equation with the same input and output pulse energies as for the 8 PCF experiment, predicted that the degree of linear polarisation of the combined pulse would be over $\sim$ 99%. Therefore it is considered that the spectral

phase differences between the pulses were not responsible for the reduction in the fraction of available beam power. Another possibility was that the layout of the 8 PCF coherent combination experiment [42] unintentionally added wavefront mismatch between combining beams (assuming the PCFs and coupling lenses were also not contributing). Either the authors had overlooked the impact of mismatched beam paths to the point of combination or had no choice given the layout of the components. Beams combining with slightly mismatched wavefronts in a compound like fashion could almost certainly reduce their beam power fraction from 100 % down to the 80%-88% range. Given the complexity of the Klenke's experiments and the number of factors they needed to consider, it is possible that the wavefront mismatch, due to mismatched path lengths after the PCFs, contributed to a reduction in their system efficiency.

In this thesis, when analysing the source of the disparity between the measured degree of linear polarisation values of $83.3 \pm 0.3\%$ (low power regime) and $85.9 \pm 0.5\%$ (high power regime) and the simulated value of 99.6%, the wavefront mismatch seen in this thesis was concluded to be the reason. In terms of the low and high power regime experiments, the mismatch in PCF length was determined by simulations to have a negligible effect on the degree of linear polarisation of the combined pulse. This led to an investigation of the wavefronts of the combining beams using Zernike polynomials, which could deconstruct a wavefront into a linear combination of radial modes. This deconstruction of the wavefront allowed some quantification of the wavefront differences between the two PCFs, and it was found that small variations in the wavefront could have a large impact on the degree of linear polarisation. If it is assumed that there was a degree of linear polarisation reduction of 5% due to the curvature of the wavefronts, a few percent due to tip / tilt contributions, and the remainder due to the remaining HOM contributions, as seen in Figure 9.4; given equal beam sizes, the simulated degree of linear polarisation could be dropped from high 99% values down to $\sim$ 90%. The simple model, used to perform the wavefront matching calculations for the telescope, was an ABCD-matrix, beam propagation model. It was not expected that this model would calculate telescope parameters, which completely compensated for the differences in the wavefronts. However, as the addition of a wavefront matching telescope did make a significant improvement to the combined degree of linear polarisation, as seen

in Section 7.1. The initial degree of linear polarisation of $\sim$ 20% was improved up to the higher $\sim$ 80% range. Nevertheless, there is still scope for improvement and refinement in the wavefront matching process. Assuming a system with excellent temporal overlap and phase stabilisation, the next most significant contribution to increasing the performance was to address the mismatch of the wavefronts at the point of combination. The most important factor, when combining pulses, is to ensure that the beam sizes and divergences are matched.

An alternative approach to the use of a telescope for correcting the wavefront mismatch is to consider the endcaps and mismatched path lengths from the end of the PCF to the point of combination. In simulations, matching the path lengths and adding end caps to each PCF, was found to make a significant difference in improving the wavefront matching of the two PCFs. In the extreme case of one PCF having double the mode field area of another, the degree of linear polarisation was found still to give values over 90 %. Before wavefront matching optics are considered in future experiments, gains could be made more economically by ensuring that the path lengths from the PCF to the point of combination are matched, and both PCFs have similar end caps. Endcaps aside, the best way of ensuring that the PCFs have matched path lengths to the point of combination, is to insert the extra path length, prior to coupling to the PCFs. The lack of an endcap on the shorter 70 μm PCF eventually would present power limitations to the amplified pulses, due to the breakdown of the end facet from high intensities. The 70 μm PCF was afflicted by a few issues also, which made it difficult to work with. Being older and a completely different design of PCF, it was affected more by thermal effects than the 85 μm PCF . This highlights the need to specify to manufacturers the power levels at which the PCFs are required to operate so they can design PCFs to operate stably under these conditions. It would be pointless to have PCFs with matched properties (length, core diameter) from different manufacturers, if a batch from one of them experienced thermal effects and mode degradation near the desired operating power. It is strongly recommended that all PCFs should be constructed with endcaps, and all manufacturers should design their PCFs to operate stably at a specified average power.

The core diameter variation seen between PCFs is a result of the core diameters being difficult to control in manufacture, leading to a spread around the target diameter. This

thesis has shown that, even with core diameters differing by almost 20% from each other, the PCFs can be coherently combined using wavefront correction optics. This suggests that the requirement for core diameters to be precisely matched can be relaxed. Differences in core diameter, by as much as 10% of a target core diameter, would not be a barrier to the efficient combination of any two given PCFs. While manufacturers are able to meet tight tolerances of 5% when making small batches of PCFs, this research suggests that a relaxed tolerance in the core diameter would be acceptable. Despite manufacturers making many hundreds of PCFs to a deadline, they could achieve sufficient precision for use in a large scale coherent combination system. This thesis has shown that the coherent combination of PCFs with a 10% spread about a target core diameter is achievable. If manufacturers still produce waste due to PCFs falling outside of these relaxed tolerances, then another approach would be to embrace the fact that waste would be produced. Typically when technologies are produced on mass scale, the cost of buying the product will fall, enabling consumers with less buying power to purchase and make use of the technology. Manufacturers could use the waste to produce a PCF factory outlet, which would sell these non-ideal PCFs at a reduced price to universities and laboratories around the world. This would have two benefits: firstly, the PCFs made for a facility scale laser could demand smaller tolerances to enable them to extract as much energy into the final pulse as possible; secondly, PCF technology could see greater proliferation towards the private sector.

Earlier in the thesis, it was mentioned that the risk of decoupling the seed beam from the PCF was a justification for adding the path length at the combining polarising beam splitter (PBS). This is still true in the scope of this thesis, as the piezo mounted mirror was not immune to small knocks, which could cause the reflected beam to wander. The beam has to wander only a few microns across the end face of the PCF before it decouples from the core. So it was decided to remove the risk entirely by placing the piezo mounted mirror at the combining PBS. However, given a more robust mirror mounting solution, the piezo mounted mirror could be aligned to be retroreflecting from the splitting PBS, allowing the seed decoupling risk to be minimised. Additionally, the power loss due to the two passes through the combining PBS could be reduced as both beams would combine on their first pass through the PBS. If the variable path length is added at the combining PBS side of

the interferometer, then the result will be mismatched beams, unless corrected. However, even with a large initial degradation of the combined beam's degree of linear polarisation, high degrees of linear polarisation are still achievable given the addition of some wavefront correction optics, as demonstrated in this thesis. The group which reported the highest peak powers (22 GW [38], 35 GW[42]) and highest pulse energies (12 mJ [42]) used PCFs with closely matched properties in terms of length, core diameters and endcaps. Also, the extra path length used to balance the two arms of the interferometer, was added onto the splitting polarising beam splitter using the same retroreflection from a piezo-mounted mirror, as was used in this thesis. Therefore, the path lengths post-PCF should be well matched. However, they do not explicitly state this, and some diagrams reported in the literature also seemed to show mismatched beam paths to the point of combination. Therefore it again leads one to believe that some of the author's efficiency drop for these two experiments might be caused by some wavefront mismatch due to asymmetric path lengths to the point of combination.

10.2 Future work

For scenarios where the layout makes it difficult to allow for accurate path length matching post-PCF, wavefront correction optics can be used. However, a wavefront-correction-optics setup could lower the damage threshold, and add more cost and loss to the overall system, especially if multiple copies were required. The first port of call should always be to match the path lengths post-PCFs as well as possible. One of the drawbacks with the wavefront matching telescope used in this thesis, was that it was constructed using lenses that were available in the laboratory, making it hard to optimise the telescope. A telescope with variable focal length lenses would allow more control over the beam divergence leaving the telescope while keeping the beam size striking the second lens the same. Modelling such a telescope for a wavefront matching scenario would be straightforward, as a solver would be able to iterate over the telescope length, position and lens focal length until a solution was found.

Other work seen in the literature that showed promise involved the use of birefringent lenses [131, 132] and polarisation dependent focusing [133]. Birefringent lenses have a variable refractive index, based upon their orientation with respect to the polarisation vector of the

laser. A calcite lens with a varying focal length has been reported [131], which varied from $f = 88$ mm to $f = 109$ mm for S- and P-Polarised beams respectively. For a single beam of S-polarised light, it is trivial to rotate the lens to change its focal length (see Figure 10.1). However, the limitation of this is that the lens cannot be adjusted to give a continuum of focal lengths. Orienting the beam to be at 45° to both refractive index directions will produce a pulse splitting effect, as seen earlier in Section 3.3.2, with divided pulse amplification. To enable more use of this type of lens and to go beyond binary focal lengths, a series of lenses would be required. Combining a birefringent lens with a spatial light modulator would allow macro and fine adjustment, respectively, over the beam divergence, providing something which more closely resembles a variable focal length lens.

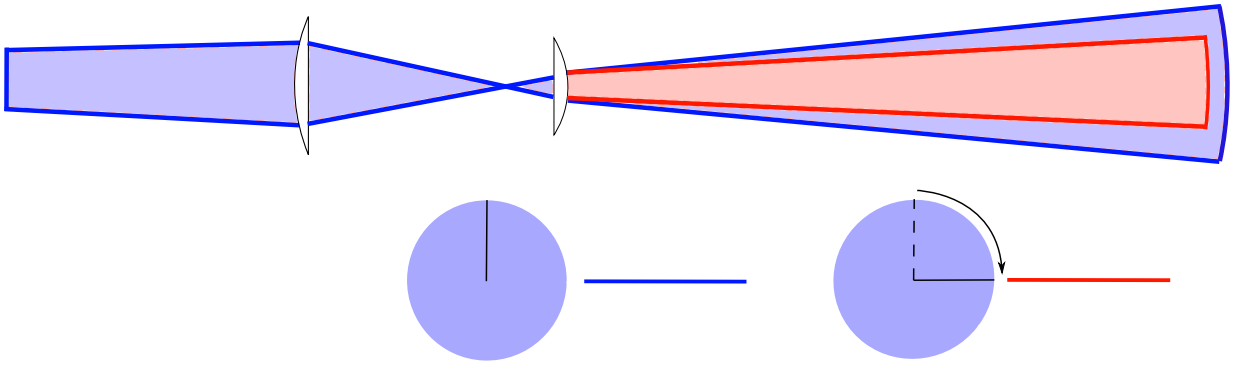


Figure 10.1: A wavefront correcting telescope with a birefringent lens as the second lens. When the telescope is subject to linearly polarised light (blue beam), the birefringent lens can be rotated to change its focal length. This has the effect of changing the beam divergence from the telescope (red beam), whilst the lens remains fixed in position. However, there are only two possible focal lengths, not a continuum.

A spatial light modulator is a device with an array of individually-addressable pixels made from liquid crystals, each of which allows a small, tuneable phase shift to be added to the beam at that point. They can be made to be transmissive or reflective, where the beam takes a double pass through the LCD, allowing twice as much phase shift to be applied. The typical use of a spatial light modulator, with regards to wavefront correction, would be to create a feedback loop between it and a wavefront sensor. This feedback loop would allow small wavefront distortions to be measured and compensated for at all points on the beam. This is shown in Figure 10.2.

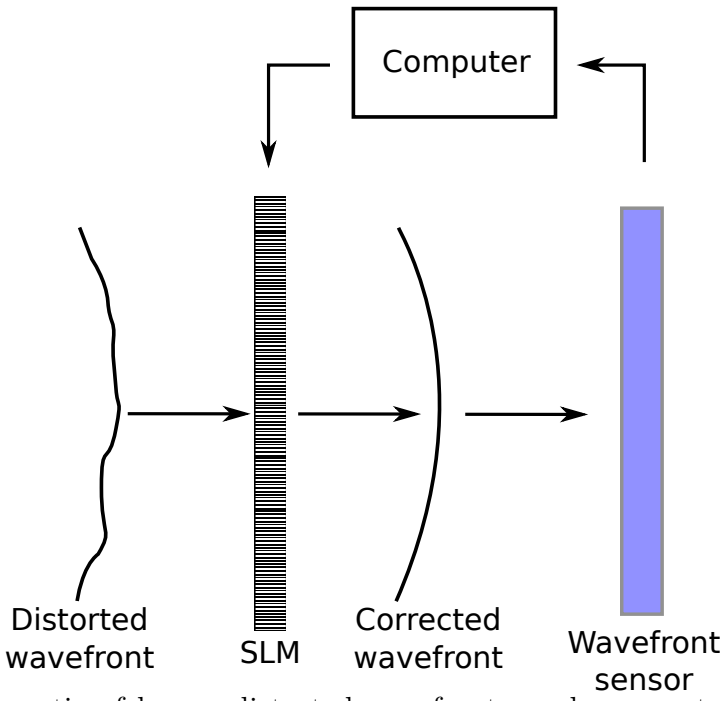


Figure 10.2: A schematic of how a distorted wavefront can be corrected by a spatial light modulator. Feedback from the wavefront sensor allows the computer to calculate the corrections needed to configure the spatial light modulator.

Referring back to the theoretical work on Zernike polynomials, work has been reported where a spatial light modulator was able to reconstruct different Zernike modes and imprint them onto a beam, up to the 15th mode [134]. If these modes could be generated and imprinted onto the beam, then they could also be suppressed. A wavefront sensor was used to measure and deconstruct the beam into a linear combination of Zernike modes, as was demonstrated in Chapter 9. Software could then be used to calculate the phase needed to suppress the identified Zernike modes. Similar to the concept of a variable focal length lens, a spatial light modulator could be used to add a defocus term to give a variable amount of focusing. However, spatial light modulators do have some limitations: Firstly, they are not rated for high power, as the liquid crystal is not 100% transmissive, with transmission efficiencies of $90\% - 98\%$ reported by manufacturers. Without some mechanism for cooling, high power beams will deposit large amounts of heat into the liquid crystal, causing irreparable damage. Secondly, spatial light modulators are limited in the amounts of phase shift they can imprint onto the beam, with typical phase shifts of 2π for double pass systems. Given that the wavefront of the beam cannot be corrected before the fibre

amplifiers, where the beam power is reduced, this wavefront correction technique will be limited to optics such as deformable mirrors.

Looking now at the future for a fibre-laser based laser-wakefield accelerator. The fibre technology that has been assumed for this thesis, was that of photonic crystal fibres (PCFs). This was because of their large core diameters, compared to step index fibres, and hence their larger volume of gain media. When building a facility-scale laser system, the ease of construction, reliability and cost will be important factors in the design. PCF rod-fibres possessing much larger core diameters, are not able to be coiled due to their brittleness. Instead, they are reliant upon v-grooves to support them along their entire length, to avoid bend losses and mode quality degradation. Their inclusion into a coherent combination setup must be done using free-space coherent combination. Smaller PCFs, that can be coiled and bent, can be spliced together to allow many fibres to be coherently combined entirely in-fibre. However, smaller PCFs, have the drawback of requiring a larger quantity of fibres. Firstly, to account for the smaller volume of gain media in each fibre, and secondly, to prevent the intensity in any single fibre from driving large non-linear effects. The number of fibres required to reach similar combined pulse energies and to avoid the end-face damage problem could be two orders of magnitude more over that of a rod-type PCF system. Because of the sheer quantity of fibres that would need to be combined, it would be impractical to use Hänsch-Couillaud-polarisation locking alone due to the number of components required.

Instead, a hybrid system could be used, making use of both tiled aperture and filled aperture setups. One possibility could be a series of 20×20 fibre arrays stabilised using quadriwave lateral shearing interferometry or digital holography. Each array would be coherently combined with other 20×20 arrays in a filled aperture like fashion. A set of 10 such arrays would need 10 locking mechanisms to phase lock the beams in each 20×20 array, and 9 Hänsch-Couillaud polarisation locking mechanisms to phase lock pairs of 20x20 arrays. These two mechanisms working together could allow the coherent combination of 4000 fibres. Under the assumption that each fibre is capable of delivering 30 μJ to the point of combination, 4000 fibres could be combined to give a total pulse energy of 120 mJ. Assuming a hexagonal packing arrangement of the fibres, as discussed in Chapter 3, the pulse energy

retrieved in the central lobe could be $\sim$ 50% which gives a usable pulse energy of 60 mJ. Assuming a 200 fs pulse, this gives an approximate peak power of 300 GW.

Shorter pulses are possible using Yb-doped silica as the gain spectrum is broad enough to support pulses as small as 28 fs [47, 48, 49, 50, 51]. The limitations preventing these pulse durations can be broken down into mechanical limitations and physical limitations. Mechanically, the stretcher and compressor used for the chirped pulse amplification system must be capable of handling the broad bandwidth. Larger bandwidths with larger stretch factors will need larger gratings. Some groups have reported pulse durations longer than the seed pulse from their oscillator simply because the bandwidth of the pulse experienced a hard-cut by the edges of the gratings in the stretcher [33, 93, 84]. The most prominent physical limitation is that due to gain narrowing. Looking at the emission spectrum of Yb-doped silica back in Figure 2.8, the non-flat emission cross section around 1030 nm leads to the central frequency components (for a bandwidth centred at 1032 nm) seeing higher gain than that of the surrounding frequency components [103]. However, there are possible options for maintaining larger bandwidths during amplification, and these are discussed shortly.

Returning now to the example of a facility-scale system. If rod-fibre PCF technology was used, and if these PCFs could deliver a compressed pulse energy of 1 mJ to the point of combination, then to reach the 60 mJ of the 4000 fibre architecture, only $\sim$ 60 PCFs would be required. Of course, to combine the fibres using Hänsch-Couillaud polarisation locking, N-1 locking systems would be required, which would be a significant cost to be considered. Compared to the cost of the PCF amplifiers, the average price for a 1 kW diode laser is $\sim$\$10k or \$10 per watt. If each PCF amplifier possessed its own diode pump laser, then the cost of 60 diode pump lasers is going to be around \$600k. The cost of the PCFs, assumed to be \$3000 per PCF, will add on \$180k. Therefore, the total cost of the amplifiers and pump lasers will be approximately \$780k. With the v-grooves, mirrors, lenses, optical mounts and other fixings, it is likely to take the total cost over one million dollars. For the locking and stabilisation system, a two-channel PID loop with Hänsch-Couillaud detectors costs approximately \$3500. If a large order was made, a series of 4 channel systems with Hänsch-Couillaud detectors could theoretically be purchased at \$5000 each (assuming discounts for bulk orders), requiring only

15 lockboxes to lock 60 PCFs. This would give an estimated cost for the electronic locking system of $75k, which is not as large as the cost of the amplifiers and pump laser but still significant. It may be possible to make further reductions by combining a larger number of PID locking loops into one lockbox.

Alternatively, Locking of Optical Coherence by Single detector Electronic frequency Tagging technology (LOCSET), could be an alternative for such a facility-scale system, with each fibre given a unique frequency tag to identify its contribution to the final beam power. However, the limiting factor would be the number of available frequency tags. This number will be limited by the ability of the locking equipment to demodulate any given tagged channel and adjust its phase. The more channels there are, the more complicated the locking mechanism will be. Compromising between complexity and the number of locking mechanisms, a group of 10 channels could be locked using LOCSET, with these combined beams then locked together using Hänsch-Couillaud polarisation locking, as shown in Figure 10.3.

This staging of locking domains reduces the number of detectors from 59 to 11; one for each group of 10 channels (6 total) and 5 to combine the combined beams of each group using Hänsch-Couillaud polarisation locking. As mentioned previously, for the beams to optimally combine in the simplest way, the path lengths post-PCF would have to be matched. This is straightforward for small numbers of PCFs in a group, but locking groups of PCFs together may become more problematic, in terms of the space available, the density of the layout and the location of the combination elements. Wavefront correction optics presented in this thesis might, therefore, see application in the combination of groups of channels, as opposed to pairs of channels. This would have two benefits: Firstly, the post-PCF path length matching would need only to be precise within a group of 10 channels; Secondly, costly wavefront correction optics would be needed only between groups of channels, as opposed to every pair of channels, further reducing the cost of building such a laser system.

Comparing the 4000 fibre and 60 PCF architectures, some benefits of the 4000 fibre architecture are: the large redundancy for individual channels becoming damaged; the whole array does not lose lock if one channel drops out; bendy fibres are easier to position and

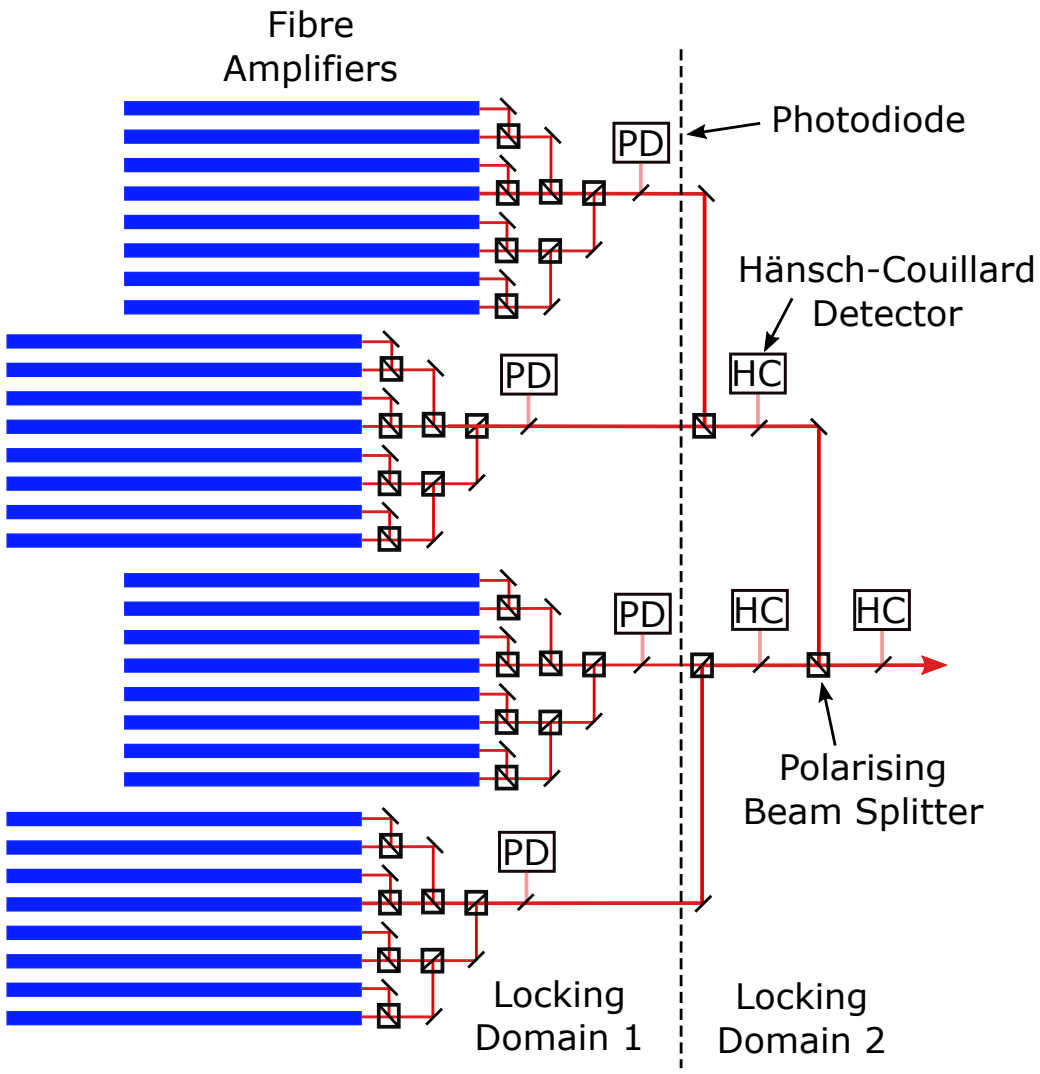


Figure 10.3: The coherent combination of many fibres achieved using a combination of locking techniques performed in stages. The first stage combines groups of channels using a single detector (per group) using a locking technique, such as digital holography or LOCSET (pictured), as seen in Locking Domain 1. The combined beams are combined in a pair like fashion, using Hänsch-Couillaud polarisation locking as seen in Locking Domain 2. By doing it this way, as opposed to Hänsch-Couillaud polarisation locking every pair of channels, the total detector count required is reduced significantly. However, this reduction will be useful only if the replacement electronics do not increase the cost or complexity of the locking system beyond what is technically feasible.

align than rod-type fibres; and the possibility of splitting pulses in fibre, as opposed to free space. Some of the drawbacks of the 4000 fibre architecture include: the difficulty in replacing individual channels in the centre of an array when damage occurs and; more fibres are required to make up for the loss of energy, seen at the point of combination, due to tiled aperture coherent combination.

The benefits of the rod-fibre architectures are: the channels have large core diameters; fewer channels are required to produce the target energy; the damaged channels are more

spatially separated, therefore, they are easier to replace and; rod-fibres, kept straight in v-grooves do not need stress rods to enable polarisation maintenance, making a simpler design. However, the drawbacks of the rod-fibre architecture are; the fibres cannot be bent and must be kept straight, impacting the size of the system and; a single damaged channel has a larger effect on the combined output power, than one from the 4000 fibre system.

To reach higher peak powers, more energy, and shorter pulses are needed. Figure 6.7 shows that the spectrum of the seed laser was narrowed as it propagated through its Yb-doped pre-amplifiers, limiting the compressed pulse duration to ~ 370 fs. Other groups have demonstrated solutions to this problem [47, 48, 49, 50, 51], producing pulses with durations as short as 28 fs using Yb-doped amplifiers. The authors of these papers set out to increase the width of the spectrum recovered into the final combined pulse. The general idea across all of the literature was to split a wide spectrum into smaller overlapping sections and to treat each spectral slice as a separate pulse. Then, these pulses could be separated using a diffraction grating, to couple pulses with different spectra into different amplifiers, or to treat the pulse as a pulse train of different colour pulses, using the principle of divided pulse amplification. All of these papers reported larger recovered bandwidths in the amplified combined pulse, than was originally recorded in the amplified pulse with an unsliced spectrum. Alternative methods to reduce the pulse duration would be to incorporate some non-linear pulse compression, as mentioned earlier in Section 4.2.2 and demonstrated by other groups [43, 44, 45]. One group reported non-linear pulse compression from 280 fs down to 26 fs with pulse energies of 1.1 mJ and 0.5 mJ before and after compression, respectively. This gave give peak powers of 2.2 GW and > 11 GW, respectively. Even though some energy has been lost, the peak power had been improved significantly.

The experiment reported in this thesis achieved a maximum combined pulse energy of $\sim 3.1 \pm 0.01$ μJ , from the coherent combination of two mismatched PCFs. In addition to using larger pump lasers, it would be possible to improve the pulse energy by altering the repetition rate of the seed laser. In this thesis, the repetition rate was set at 6.49 MHz, due to legacy reasons, whereas the rep-rate found in similar experiments [38, 42] was 40 kHz and 55.9 kHz. Lower repetition rates allow the pump laser more time to build up the population

inversion in the amplifier, before the next pulse arrives. Therefore the next step in advancing the experiments in this thesis, would be to lower the repetition rate to ∼10 kHz.

Another idea that may help to improve the pulse energy is an enhancement cavity. The pulses are sent into a very long ring based cavity, with subsequent pulses entering and stacking on top of the first [135, 136], as they arrive back at the entry point. The stacking of pulses is coherent combination in the sense that, the first pulse must carry sufficient phase delay to arrive at the first cavity mirror, at the same time as the second pulse arrives to enter the cavity. The conceptual work [136] discussed the idea of stacking > 600 pulses in a cavity. Once sufficient energy had been accumulated, the combined pulse would then be switched out. This stacking of pulses could technically enable ~ 600 pulses with energies of ~ 1.6 mJ to be stacked to produce pulses of ~ 1 J. The main benefit of this pulse stacking architecture would be that high pulse energies could be reached without having to amplify pulses harder and risk acquiring more non-linear phase. The energy gain of the combined pulse can be spread out over time, with the trade-off being a reduction in the system's repetition rate. With fibre based systems running at MHz repetition rates from the outset, enhancement cavities could be one way to enable fibre lasers to produce terawatt peak powers.

10.3 Closing remarks

Looking forward to terawatt peak power applications, the direction of development for a facility-scale, fibre-laser system is to use rod-type, large pitch fibres with divided pulse amplification, a pulse stacking enhancement cavity and non-linear compression. The current state of the art for PCF-based laser systems is to produce 12 mJ pulses, at a peak power of 35 GW, and a repetition rate of 55.9 kHz. These pulse energies were achieved using a combination of divided pulse amplification and the coherent combination of 8 PCFs [42]. This 12 mJ pulse energy is similar to that used by Milchberg [31], who used a Ti:Sapphire laser producing < 10 mJ, 30 fs pulses at 1 kHz to demonstrate electron acceleration at the $\sim$ MeV level with 1 pC of charge. Based on the similarities between these two experiments, the advance in technology, it is likely that a fibre based laser-plasma wakefield demonstration accelerator producing MeV electrons, will be constructed within the next five years. For now, higher pulse energies and smaller pulse durations will be required to directly compete with

Ti:Sapphire systems producing GeV electron beams. Further research into large pitch fibres will allow PCFs to hold more energy, and reduce the number of PCFs required. This reduces the burden on locking mechanisms, as the same final pulse energy can be attained using a smaller number of channels. Smaller pulse durations below the 200 fs, already demonstrated, will be required to allow a system to achieve terawatt peak powers without having to use larger amounts of pulse energy to compensate for the longer pulse durations. With research groups like those at the Friedrich-Schiller-University, Jena, now focused on high-intensity pulses possessing pulse durations of 100 fs and below, it is likely that a fibre based terawatt laser facility could be built within the next 15 years, if funding is available.

Appendix A

Datasheets

A.1 70 µm Photonic Crystal Fibre Datasheet

NKT Photonics
the power of light

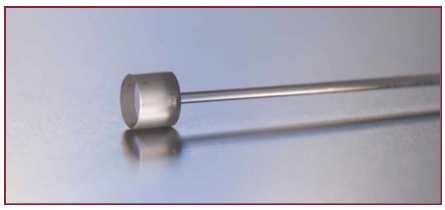
Crystal Fibre • *aero*LASE • Koheras • SuperK

DC-200/70-PM-Yb-ROD

Large Mode-area · Rare Earth Doped · Rod Waveguide · Polarization Maintaining

Ytterbium doped PM ROD fiber with ultra large mode area

- **Extremely large mode area**
- **Polarization maintaining**
- **High pump absorption**



The DC-200/70-PM-Yb-ROD represents the pinnacle of active gain fibers. The rod is polarization-maintaining and has an effective area above 2200 µm². The very large effective area combined with extremely high pump-absorption of up to 30dB/m pushes the nonlinear thresholds to a new regime in fibers. The rod is aimed at high peak-power pulsed amplifiers and has been shown to handle mega-Watt peak-power levels. For even higher power handling, we recommend our DC-285/100-PM-Yb-ROD fiber.

The multimode pump light is guided by our proven airclad technology, ensuring low loss, high damage threshold and a large numerical aperture (NA). The large NA relaxes tolerances on coupling optics and facilitates the use of lower brightness diodes.

The DC-200/70-PM-Yb-ROD is a rod fiber with an outer diameter of 1.7 mm – bridging the DPSS and fiber laser world. The large outer diameter ensures that the extremely large core will not suffer from bending loss. Moreover, the rod is coating free to facilitate thermal management in high-power setups. The rods are available with large AR coated end-caps in a range of standard lengths. Other configurations are available upon request.

More info on using ROD fibers can be found on our website at www.nktphotonics.com/application_notes

Applications
- Ultra high-power amplifiers
- Ultra high-power polarized lasers
- Ultra high-power Q-Switched lasers

Optical properties	
Signal core	
Mode field diameter	55 ± 5 µm
Mode field area	2200 ± 200 µm²
NA @ 1060 nm	~ 0.02
Multimode pump core	
Numerical aperture @ 950 nm	0.6 ± 0.05
Pump absorption @ 920 nm	~ 10 dB/m
Pump absorption @ 976 nm	~ 30 dB/m
Slope efficiency	~60%
Polarization Parameters	
Birefringence Δn	› $1 \cdot 10^{-4}$
Polarization Extinction Ratio	15-25dB*

* Depending on system configuration. The fiber is moderately polarizing with a DOP of 80-90% depending on the configuration. The fiber should be operated in the slow-axis.

Physical properties	
Core material	Yb-doped silica
Outer cladding diameter	1.7 ± 0.1 mm
Coating	None
Signal core diameter	70 ± 5 µm
Pump-cladding diameter	200 ± 5 µm
Pump-cladding shape	Circular

End cap option	
Material	Pure silica
Length/diameter	8 mm / 8.2 mm
AR coating, R @ 1040-1080 nm	‹ 0.1 %
AR coating, R @ 976 nm	‹ 0.5 %
Spot size at facet ($1/e^2$)	› 140 µm
Signal return loss	‹ -35 dB
NA supported	› 0.6

Standard configuration	Rod length	End 1	End 2
Type I	36 cm	End cap	End cap
Type II	55 cm	End cap	End cap
Type III	55 cm	End cap	Sealing
Type IV	80 cm	End cap	End cap
Type V	80 cm	End cap	Sealing
Type VI	120 cm	Sealing	Sealing

Sealed ends are angle-cleaved to 3° ± 1° to minimize reflection from the facet. Sealing length is 50 µm ± 25 µm (measured on the airclad).
The standard configurations are available from stock. Other configurations can be manufactured upon request.

DC-200-70-PM-Yb-ROD-110616

NKT Photonics A/S (Headquarters)
Blokken 84 • 3460 Birkerød • Denmark
Phone: +45 4348 3900
Fax: +45 4348 3901
www.nktphotonics.com

NKT Photonics GmbH
Schanzenstrasse 39 • Bldg D9-D13
51063 Cologne • Germany
Phone: +49 221 99511-0
Fax: +49 221 99511-650

NKT Photonics Inc.
1400 Campus Drive West • Morganville
NJ 07751 • USA
Phone: +1 732 972 9937
Fax: +1 732 414 4094

A.2 85 µm Photonic Crystal Fibre Datasheet

NKT Photonics
the power of light

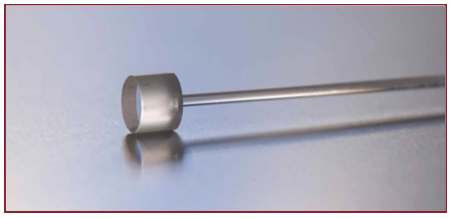
Crystal Fibre • *aero*LASE • Koheras • SuperK

DC-200/85-Yb-ROD

Large Mode-area · Rod Waveguide · Rare Earth Doped

Ytterbium doped ROD fiber with ultra large mode area

- **Extremely large mode area**
- **High pump absorption**
- **Can be operated in PM mode**



With an effective area above 3000 µm², the DC-200/85-Yb-ROD represents the pinnacle of active gain fibers. The very large effective area combined with extremely high pump-absorption of up to 30dB/m pushes the nonlinear thresholds to a new regime in fibers. The rod is aimed at high peak-power pulsed amplifiers and has been shown to handle mega-Watt peak-power levels.

The multimode pump light is guided by our proven airclad technology, ensuring low loss, high damage threshold and a large numerical aperture (NA). The large NA relaxes tolerances on coupling optics and facilitates the use of lower brightness diodes.

The DC-200/85-Yb-ROD is a rod fiber with an outer diameter of 1.7 mm – bridging the DPSS and fiber laser world. The large outer diameter ensures that the extremely large core will not suffer from bending loss. Moreover, the rod is coating free to facilitate thermal management in high-power setups. The rods are available with large AR coated end-caps in a range of standard lengths. Other configurations are available upon request.

For polarization maintaining RODs, we recommend our DC-285/100-PM-Yb-ROD.

More info on using ROD fibers can be found on our website at www.nktphotonics.com/support

Applications
- Ultra high-power amplifiers
- Ultra high-power polarized lasers
- Ultra high-power Q-switched lasers

Optical properties	
Signal core	
Mode field diameter	65 ± 5 µm
NA @ 1060 nm	~ 0.02
Multimode pump core	
Numerical aperture @ 950 nm	0.55 ± 0.05
Pump absorption @ 920 nm	~ 10 dB/m
Pump absorption @ 976 nm	~ 30 dB/m
System performance	
NL FOM [1]	~ 8 [rad·ns/mJ]
Slope efficiency [2]	~ 60 %

[1] The nonlinear figure-of-merit is calculated as:

$$\text{NL FOM} = \frac{\gamma}{g_0^{\max}} \cdot 10^6 \left[rad \cdot \frac{ns}{mJ} \right]$$

[2] With respect to pump power in the fiber

Physical properties	
Core material	Yb-doped silica
Outer cladding diameter	1.7 ± 0.1 mm
Coating	None
Signal core diameter	85 ± 5 µm
Pump-cladding diameter	200 ± 5 µm
Pump-cladding shape	Circular
Outer cladding material	Pure silica
Inner cladding material	Pure silica

Standard configuration	Rod length	End 1	End 2
Type I	36 cm	End cap	End cap
Type II	55 cm	End cap	End cap
Type III	55 cm	End cap	Sealing
Type IV	80 cm	End cap	End cap
Type V	80 cm	End cap	Sealing
Type VI	120 cm	Sealing	Sealing

Sealed ends are angle-cleaved to 3° ± 1° to minimize reflection from the facet. Sealing length is 75 µm ± 25 µm (measured on the airclad).
The standard configurations are available from stock. Other configurations can be manufactured upon request.

DC-200/85-Yb-ROD-120427

NKT Photonics A/S (Headquarters)
Blokken 84 • 3460 Birkerød • Denmark
Phone: +45 4348 3900
Fax: +45 4348 3901
www.nktphotonics.com

NKT Photonics GmbH
Schanzenstrasse 39 • Bldg D9-D13
51063 Cologne • Germany
Phone: +49 221 99511-0
Fax: +49 221 99511-650

NKT Photonics Inc.
1400 Campus Drive West • Morganville
NJ 07751 • USA
Phone: +1 732 972 9937
Fax: +1 732 414 4094

References

[1] L Evans & P Bryant. "LHC machine". *Journal of instrumentation*, **3**, 08 (2008) S08001.

[2] PJ Bryant & K Johnsen. "The principles of circular accelerators and storage rings". Cambridge University Press (2005).

[3] M Benedikt et al. "Challenges for highest energy circular colliders". *IPAC2014: Proceedings of the 5th International Particle Accelerator Conference*.

[4] J Wenninger et al. "Future Circular Collider Study Lepton Collider Parameters". *CERN EDMS no*, **1346082**.

[5] TP Wangler. "RF Linear accelerators". John Wiley & Sons (2008).

[6] M Koratzinos et al. "TLEP: A High-Performance Circular e+ e-Collider to Study the Higgs Boson". *arXiv preprint arXiv:1305.6498*.

[7] E Esarey et al. "Physics of laser-driven plasma-based electron accelerators". *Reviews of Modern Physics*, **81**, 3 (2009) 1229.

[8] S Kneip et al. "Near-GeV Acceleration of Electrons by a Nonlinear Plasma Wave Driven by a Self-Guided Laser Pulse". *Physical Review Letters*, **103** (2009) 035002.

[9] WP Leemans et al. "Multi-GeV electron beams from capillary-discharge-guided subpetawatt laser pulses in the self-trapping regime". *Physical review letters*, **113**, 24 (2014) 245002.

[10] NAM Hafz et al. "Laser acceleration of electron beams to the GeV-class energies in gas jets". *Journal of the Optical Society of Korea*, **13**, 1 (2009) 8–14.

[11] P Muggli et al. "Physics of the AWAKE Project" (2013).

[12] YF Li et al. "Generation of 20 kA electron beam from a laser wakefield accelerator". *Physics of Plasmas*, **24**, 2 (2017) 023108.

[13] M Mirzaie et al. "Demonstration of self-truncated ionization injection for GeV electron beams". *Scientific reports*, **5** (2015) 14659.

[14] G Golovin et al. "Intrinsic beam emittance of laser-accelerated electrons measured by x-ray spectroscopic imaging". *Scientific reports*, **6** (2016) 24622.

[15] R Weingartner et al. "Ultralow emittance electron beams from a laser-wakefield accelerator". *Physical Review Special Topics-Accelerators and Beams*, **15**, 11 (2012) 111302.

[16] CB Schroeder et al. "Physics considerations for laser-plasma linear colliders". *Physical Review Special Topics-Accelerators and Beams*, **13**, 10 (2010) 101301.

[17] J Isaacs & P Sprangle. "Proton acceleration in a slow wakefield". *Applied Physics Letters*, **110**, 2 (2017) 024101.

[18] G Malka & JL Miquel. "Experimental confirmation of ponderomotive-force electrons produced by an ultrarelativistic laser pulse on a solid target". *Physical Review Letters*, **77**, 1 (1996) 75.

[19] E Esarey et al. "Electron injection into plasma wakefields by colliding laser pulses". *Physical Review Letters*, **79**, 14 (1997) 2682.

[20] M Chen et al. "Electron injection and emittance control by transverse colliding pulses in a laser-plasma accelerator". *Physical Review Special Topics-Accelerators and Beams*, **17**, 5 (2014) 051303.

[21] A Pak et al. "Injection and trapping of tunnel-ionized electrons into laser-produced wakes". *Physical Review Letters*, **104**, 2 (2010) 025003.

[22] C McGuffey et al. "Ionization induced trapping in a laser wakefield accelerator". *Physical review letters*, **104**, 2 (2010) 025004.

[23] CE Clayton et al. "Self-guided laser wakefield acceleration beyond 1 GeV using ionization-induced injection". *Physical Review Letters*, **105**, 10 (2010) 105003.

[24] BB Pollock et al. "Demonstration of a narrow energy spread, 0.5 GeV electron beam from a two-stage laser wakefield accelerator". *Physical review letters*, **107**, 4 (2011) 045001.

[25] CGR Geddes et al. "Plasma-density-gradient injection of low absolute-momentum-spread electron bunches". *Physical review letters*, **100**, 21 (2008) 215004.

[26] K Zhao et al. "Tailoring a 67 attosecond pulse through advantageous phase-mismatch". *Optics letters*, **37**, 18 (2012) 3891–3893.

[27] S Hädrich et al. "Energetic sub-2-cycle laser with 216 W average power". *Optics Letters*, **41**, 18 (2016) 4332–4335.

[28] JH Sung et al. "High-contrast 0.1-Hz 4-PW laser at CoReLS". In "Photonics Conference (IPC), 2016 IEEE", pages 456–457. IEEE (2016).

[29] TJ Yu et al. "Generation of high-contrast, 30 fs, 1.5 PW laser pulses from chirped-pulse amplification Ti: sapphire laser". *Optics express*, **20**, 10 (2012) 10807–10815.

[30] SM Hooker et al. "Multi-Pulse Laser Wakefield Acceleration: A New Route to Efficient, High-Repetition-Rate Plasma Accelerators and High Flux Radiation Sources". *J. Phys.*, **B47** (2014) 234003.

[31] F Salehi et al. "MeV electron acceleration at 1 kHz with < 10 mJ laser pulses". *Optics Letters*, **42**, 2 (2017) 215–218.

[32] J Limpert et al. "Yb-doped large-pitch fibres: effective single-mode operation based on higher-order mode delocalisation". *Light: Science & Applications*, **1**, 4 (2012) e8.

[33] T Eidam et al. "Fiber chirped-pulse amplification system emitting 3.8 GW peak power". *Optics Express*, **19**, 1 (2011) 255–260.

[34] J Limpert et al. "High-power air-clad large-mode-area photonic crystal fiber laser". *Optics Express*, **11**, 7 (2003) 818–823.

[35] C Wei et al. "All-fibre ytterbium-doped photonic crystal fibre laser with high efficiency". *Chinese Physics Letters*, **25**, 3 (2008) 960.

[36] J Li et al. "High-power coherent beam combining of two photonic crystal fiber lasers". *IEEE Photonics Technology Letters*, **20**, 11 (2008) 888–890.

[37] P Crump et al. "> 360 W and > 70% efficient GaAs-based diode lasers". In "Lasers and Applications in Science and Engineering", pages 21–29. International Society for Optics and Photonics (2005).

[38] A Klenke et al. "22GW peak power femtosecond fiber CPA system". *Proceedings of SPIE*, **9344** (2015) 93441A–93441A–6.

[39] DN Schimpf et al. "Optimization of high performance ultrafast fiber laser systems to > 10 GW peak power". *Journal of the Optical Society of America B*, **27**, 10 (2010) 2051–2060.

[40] BP Abbott et al. "Observation of gravitational waves from a binary black hole merger". *Physical review letters*, **116**, 6 (2016) 061102.

[41] WS Brocklesby et al. "ICAN as a new laser paradigm for high energy, high average power femtosecond pulses". *The European Physical Journal Special Topics*, **223**, 6 (2014) 1189–1195.

[42] M Kienel et al. "12 mJ kW-class ultrafast fiber laser system using multidimensional coherent pulse addition". *Optics Letters*, **41**, 14 (2016) 3343–3346.

[43] S Hädrich et al. "Nonlinear compression to sub-30-fs, 0.5 mJ pulses at 135 W of average power". *Optics letters*, **38**, 19 (2013) 3866–3869.

[44] C Xie et al. "Generation of 25-fs high energy pulses by SPM-induced spectral broadening in a photonic crystal fiber laser system". *IEEE Photonics Technology Letters*, **24**, 5 (2012) 551.

[45] J Limpert et al. "SPM-induced spectral compression of picosecond pulses in a single-mode Yb-doped fiber amplifier". *Applied Physics B: Lasers and Optics*, **74**, 2 (2002) 191–195.

[46] J Bourderionnet et al. "Coherent phase combining of 64 fibers". In "The European Conference on Lasers and Electro-Optics", page PDA_4. Optical Society of America (2011).

[47] X Zhou et al. "Generation of 28-fs pulses from a mode-locked ytterbium fiber oscillator". *Optics express*, **16**, 10 (2008) 7055–7059.

[48] Z Zhang et al. "Sub-50 fs Yb-doped laser with anomalous-dispersion photonic crystal fiber". *Optics letters*, **38**, 6 (2013) 956–958.

[49] M Baumgartl et al. "High average and peak power femtosecond large-pitch photonic-crystal-fiber laser". *Optics letters*, **36**, 2 (2011) 244–246.

[50] F Guichard et al. "Two-channel pulse synthesis to overcome gain narrowing in femtosecond fiber amplifiers". *Optics letters*, **38**, 24 (2013) 5430–5433.

[51] P Rigaud et al. "Spectral division amplification of a 40 nm bandwidth in a multicore Yb doped fiber and femtosecond pulse synthesis with in-fiber delay line". *Optics express*, **23**, 21 (2015) 27448–27456.

[52] J Limpert et al. "Performance scaling of ultrafast laser systems by coherent addition of femtosecond pulses". *IEEE Journal of Selected Topics in Quantum Electronics*, **20**, 5 (2014) 268–277.

[53] G Mourou et al. "The future is fibre accelerators". *Nature Photonics*, **7**, 4 (2013) 258–261.

[54] K Yamakawa et al. "100-TW sub-20-fs Ti: sapphire laser system operating at a 10-Hz repetition rate". *Optics letters*, **23**, 18 (1998) 1468–1470.

[55] M Aoyama et al. "0.85-PW, 33-fs Ti: sapphire laser". *Optics letters*, **28**, 17 (2003) 1594–1596.

[56] JH Sung et al. "0.1 Hz 1.0 PW Ti: sapphire laser". *Optics letters*, **35**, 18 (2010) 3021–3023.

[57] NV Zamfir. "Extreme Light Infrastructure–Nuclear Physics (ELI-NP) European Research Centre". In "EPJ Web of Conferences", volume 66, page 11043. EDP Sciences (2014).

[58] Y Nabekawa et al. "All-solid-state high-peak-power Ti: sapphire laser system above 5-kHz repetition rate". *Applied Physics B*, **70**, 1 (2000) S171–S179.

[59] Z Ling et al. "All-solid-state dual end pumped YVO4: Nd/LBO blue laser with 21.8 W output power at 457 nm". *Optics and Spectroscopy*, **116**, 3 (2014) 470–472.

[60] GP Agrawal. "Nonlinear fiber optics". Academic press (2013).

[61] JA Buck. "Fundamentals of optical fibers". John Wiley & Sons (2004).

[62] "Hole Control in Photonic Crystal Fibres. Thesis". Y Chen. Ph.D. thesis, University of Bath (2013).

[63] "Performance scaling of laser amplifiers via coherent combination of ultrashort pulses". A Klenke. Ph.D. thesis, Friedrich-Schiller-Universitt Jena (2016).

[64] M-J Li et al. "Limit of effective area for single-mode operation in step-index large mode area laser fibers". *Journal of Lightwave Technology*, **27**, 15 (2009) 3010–3016.

[65] NA Mortensen et al. "Modal cutoff and the V parameter in photonic crystal fibers". *Optics letters*, **28**, 20 (2003) 1879–1881.

[66] JC Knight et al. "Large mode area photonic crystal fibre". *Electronics Letters*, **34**, 13 (1998) 1347–1348.

[67] TA Birks et al. "Endlessly single-mode photonic crystal fiber". *Optics letters*, **22**, 13 (1997) 961–963.

[68] G Gu et al. "Ytterbium-doped large-mode-area all-solid photonic bandgap fiber lasers". *Optics express*, **22**, 11 (2014) 13962–13968.

[69] J Limpert et al. "Extended single-mode photonic crystal fiber lasers". *Optics Express*, **14**, 7 (2006) 2715–2720.

[70] TT Alkeskjold et al. "Single-mode ytterbium-doped large-mode-area photonic bandgap rod fiber amplifier". *Optics express*, **19**, 8 (2011) 7398–7409.

[71] F Stutzki et al. "High-speed modal decomposition of mode instabilities in high-power fiber lasers". *Optics letters*, **36**, 23 (2011) 4572–4574.

[72] C Jauregui et al. "Impact of modal interference on the beam quality of high-power fiber amplifiers". *Optics express*, **19**, 4 (2011) 3258–3271.

[73] B Ward et al. "Origin of thermal modal instabilities in large mode area fiber amplifiers". *Optics express*, **20**, 10 (2012) 11407–11422.

[74] C Jauregui et al. "Physical origin of mode instabilities in high-power fiber laser systems". *Optics express*, **20**, 12 (2012) 12912–12925.

[75] R Paschotta et al. "Ytterbium-doped fibre amplifiers". *IEEE Journal of quantum electronics*, **33**, 7 (1997) 1049–1056.

[76] NKT Photonics, "NKT Photonics products" (Accessed 20/05/2018). `http://www.nktphotonics.com/`

[77] O Schmidt et al. "Single-polarization ultra-large-mode-area Yb-doped photonic crystal fiber". *Optics express*, **16**, 6 (2008) 3918–3923.

[78] G Brooker. "Modern classical optics". volume 8. Oxford University Press (2003).

[79] Y Takushima et al. "Polarization-stable and single-frequency fiber lasers". *Journal of lightwave technology*, **16**, 4 (1998) 661–669.

[80] J Lhermite et al. "Coherent combining of 49 laser beams from a multiple core optical fiber by a spatial light modulator". *Optics express*, **18**, 5 (2010) 4783–4789.

[81] J Bourderionnet et al. "Collective coherent phase combining of 64 fibers". *Optics express*, **19**, 18 (2011) 17053–17058.

[82] C Bellanger et al. "Coherent fiber combining by digital holography". *Optics letters*, **33**, 24 (2008) 2937–2939.

[83] E Seise et al. "Coherent addition of fiber-amplified ultrashort laser pulses". *Optics Express*, **18**, 26 (2010) 27827–27835.

[84] A Klenke et al. "530 W, 1.3 mJ, four-channel coherently combined femtosecond fiber chirped-pulse amplification system". *Optics letters*, **38**, 13 (2013) 2283–2285.

[85] L 0Daniault et al. "Coherent beam combining of two femtosecond fiber chirped-pulse amplifiers". *Optics Letters*, **36**, 5 (2011) 621–623.

[86] SM Redmond et al. "Diffractive coherent combining of a 2.5 kW fiber laser array into a 1.9 kW Gaussian beam". *Optics letters*, **37**, 14 (2012) 2832–2834.

[87] PA Thielen et al. "Two-dimensional diffractive coherent combining of 15 fiber amplifiers into a 600 W beam". *Optics letters*, **37**, 18 (2012) 3741–3743.

[88] EC Cheung et al. "Diffractive-optics-based beam combination of a phase-locked fiber laser array". *Optics letters*, **33**, 4 (2008) 354–356.

[89] L Daniault et al. "Passive coherent beam combining of two femtosecond fiber chirped-pulse amplifiers". *Optics letters*, **36**, 20 (2011) 4023–4025.

[90] Y Zaouter et al. "Passive coherent combination of two ultrafast rod type fiber chirped pulse amplifiers". *Optics letters*, **37**, 9 (2012) 1460–1462.

[91] E Hecht. "Optics". Addison-Wesley (2002). ISBN 9780805385663.

[92] M Kienel et al. "Analysis of passively combined divided-pulse amplification as an energy-scaling concept". *Optics express*, **21**, 23 (2013) 29031–29042.

[93] M Kienel et al. "Energy scaling of femtosecond amplifiers using actively controlled divided-pulse amplification". *Optics letters*, **39**, 4 (2014) 1049–1052.

[94] S Zhou et al. "Divided-pulse amplification of ultrashort pulses". *Optics letters*, **32**, 7 (2007) 871–873.

[95] A Klenke et al. "Basic considerations on coherent combining of ultrashort laser pulses". *Optics Express*, **19**, 25 (2011) 25379–25387.

[96] F Guichard et al. "High-energy chirped-and divided-pulse Sagnac femtosecond fiber amplifier". *Optics letters*, **40**, 1 (2015) 89–92.

[97] TM Shay. "Theory of electronically phased coherent beam combination without a reference beam". *Optics Express*, **14**, 25 (2006) 12188–12195.

[98] A Flores. "Coherent beam combining of fiber amplifiers in a kW regime". In "CLEO: Science and Innovations", page CFE3. Optical Society of America (2011).

[99] "LOCSET phase locking: operation, diagnostics, and applications (Thesis)". B Pulford. Ph.D. thesis (2012).

[100] TW Hänsch & B Couillaud. "Laser frequency stabilization by polarization spectroscopy of a reflecting reference cavity". *Optics Communications*, **35**, 3 (1980) 441 – 444.

[101] O Martinez. "3000 times grating compressor with positive group velocity dispersion: Application to fiber compensation in 1.3-1.6 μm region". *IEEE Journal of Quantum Electronics*, **23**, 1 (1987) 59–64.

[102] S Backus et al. "High power ultrafast lasers". *Review of Scientific Instruments*, **69**, 3 (1998) 1207–1223.

[103] S Hooker & C Webb. "Laser Physics". Oxford Master Series in Physics. OUP Oxford (2010). ISBN 9780198506928.

[104] I Walmsley et al. "The role of dispersion in ultrafast optics". *Review of Scientific Instruments*, **72**, 1 (2001) 1–29.

[105] B Tatian. "Fitting refractive-index data with the Sellmeier dispersion formula". *Appl. Opt.*, **23**, 24 (1984) 4477–4485.

[106] G Ghosh. "Sellmeier coefficients and dispersion of thermo-optic coefficients for some optical glasses". *Applied optics*, **36**, 7 (1997) 1540–1546.

[107] M Ditmire Perry et al. "Self-phase modulation in chirped-pulse amplification". *Optics letters*, **19**, 24 (1994) 2149–2151.

[108] SP Singh & N Singh. "Nonlinear effects in optical fibers: origin, management and applications". *Progress In Electromagnetics Research*, **73** (2007) 249–275.

[109] S Zhou et al. "Compensation of nonlinear phase shifts with third-order dispersion in short-pulse fiber amplifiers". *Optics express*, **13**, 13 (2005) 4869–4877.

[110] SJ Russell et al. "Artificial intelligence: a modern approach". volume 2. Prentice hall Upper Saddle River (2003).

[111] C-R Hwang. "Simulated annealing: theory and applications". *Acta Applicandae Mathematicae*, **12**, 1 (1988) 108–111.

[112] M Raghuramaiah et al. "A second-order autocorrelator for single-shot measurement of femtosecond laser pulse durations". *Sadhana*, **26**, 6 (2001) 603–611.

[113] IT Sorokina & KL Vodopyanov. "Solid-state mid-infrared laser sources". volume 89. Springer Science & Business Media (2003).

[114] CASTECH, "CASTECH Product Catalogue" (Accessed 20/05/2018). `http://www.castech.com/`

[115] LJ Nevay et al. "High Power Fibre Laser System for a High Repetition Rate Laserwire". *Physical Review. Special Topics - Accelerators and Beams*, **17** (2014) 072801.

[116] CERN, "International Linear Collider (ILC) baseline configuration document" (Accessed 10/06/2018). `https://www.linearcollider.org/wiki/doku.php`

[117] C Barnard et al. "Analytical model for rare-earth-doped fiber amplifiers and lasers". *IEEE Journal of Quantum Electronics*, **30**, 8 (1994) 1817–1830.

[118] D Milam. "Review and assessment of measured values of the nonlinear refractive-index coefficient of fused silica". *Appl. Opt.*, **37**, 3 (1998) 546–550.

[119] D Felice. "A Study of a Nonlinear Schrödinger Equation for Optical Fibers". *arXiv preprint arXiv:1612.00358*.

[120] R Deiterding et al. "A reliable split-step Fourier method for the propagation equation of ultra-fast pulses in single-mode optical fibers". *Journal of Lightwave Technology*, **31**, 12 (2013) 2008–2017.

[121] S Butterworth. "On the theory of filter amplifiers". *Wireless Engineer*, **7**, 6 (1930) 536–541.

[122] F Jansen et al. "Thermally induced waveguide changes in active fibers". *Optics express*, **20**, 4 (2012) 3997–4008.

[123] DataRay. "Dataray User Manual - WinCam D series". **Rev. 1207a**.

[124] MK Davis et al. "Thermal effects in doped fibers". *Journal of lightwave technology*, **16**, 6 (1998) 1013.

[125] DN Schimpf et al. "Self-phase modulation compensated by positive dispersion in chirped-pulse systems". *Optics express*, **17**, 7 (2009) 4997–5007.

[126] F Jansen et al. "Avoided crossings in photonic crystal fibers". *Optics express*, **19**, 14 (2011) 13578–13589.

[127] D Kabeya et al. "Efficient phase-locking of 37 fiber amplifiers by phase-intensity mapping in an optimization loop". *Optics Express*, **25**, 12 (2017) 13816–13821.

[128] P Fricker, "Zernike polynomial MATLAB code" (Accessed 10/06/2018). `https://uk.mathworks.com/matlabcentral/fileexchange/7687-zernike-polynomials`

[129] JC Wyant. "Optical Imaging and Aberrations, Part I, Ray Geometrical Optics, by Virendra N. Mahajan". *Optics & Photonics News*, **10** (1999) 54–58.

[130] G-M Dai. "Wavefront expansion basis functions and their relationships". *Journal of the Optical Society of America A*, **23**, 7 (2006) 1657–1668.

[131] AK Kirby et al. "Adaptive lenses based on polarization modulation". In "Proceedings of SPIE", volume 6018, page 14 (2005).

[132] RS Eng & KG Leib. "Multiple imagery with birefringent lenses". *Applied optics*, **8**, 10 (1969) 2117–2120.

[133] D Schmid et al. "Adjustable and robust methods for polarization-dependent focusing". *Optics express*, **21**, 13 (2013) 15538–15552.

[134] GD Love. "Wave-front correction and production of Zernike modes with a liquid-crystal spatial light modulator". *Applied optics*, **36**, 7 (1997) 1517–1524.

[135] I Pupeza et al. "Power scaling of a high-repetition-rate enhancement cavity". *Optics letters*, **35**, 12 (2010) 2052–2054.

[136] S Breitkopf et al. "Stack and dump: Peak-power scaling by coherent pulse addition in passive cavities". *The European Physical Journal Special Topics*, **224**, 13 (2015) 2573–2577.