Femtosecond optical parametric oscillator frequency combs for coherent pulse synthesis

Richard A. McCracken

Submitted for the degree of Doctor of Philosophy

Heriot-Watt University

School of Engineering and Physical Sciences

September 2013

Abstract

Coherent pulse synthesis takes as its objective the piecewise assembly of a sequence of identical broadband pulses from two or more mutually-coherent sequences of narrowband pulses. The requirements for pulse synthesis are that the parent pulses share the same repetition frequency, are phase coherent and have low mutual timing jitter over the required observation time.

The work carried out in this thesis explored the requirements for broadband coherent pulse synthesis between the multiple visible outputs of a synchronously pumped femtosecond optical parametric oscillator. A femtosecond Ti:sapphire laser was characterised and used to pump a PPKTP-based OPO that produced a number of second-harmonic and sum-frequency mixing outputs across the visible region. Using a novel lock-to-zero CEO stabilisation technique, broadband phase coherence was established between all the pulses on the optical bench, producing the broadest zero-offset frequency comb to date. Employing a common optical path for all the pulses provided common-mode rejection of noise, ensuring less than 150 attoseconds of timing jitter between the pulses over a 1 second observation window. The parent pulses were compressed and their relative delays altered in a quasi-common path prism delay line, allowing pulse synthesis at a desired reference plane.

For Louise and Daniel

ACADEMIC REGISTRY
Research Thesis Submission

Name*:*	Richard Alexander McCracken		
School/PGI:	EPS/IPAQS		
Version: *(i.e. First, Resubmission, Final)*	First	Degree Sought (Award **and** Subject area)	Ph. D. in Physics

Declaration

In accordance with the appropriate regulations I hereby submit my thesis and I declare that:

1) the thesis embodies the results of my own work and has been composed by myself
2) where appropriate, I have made acknowledgement of the work of others and have made reference to work carried out in collaboration with other persons
3) the thesis is the correct version of the thesis for submission and is the same version as any electronic versions submitted*.
4) my thesis for the award referred to, deposited in the Heriot-Watt University Library, should be made available for loan or photocopying and be available via the Institutional Repository, subject to such conditions as the Librarian may require
5) I understand that as a student of the University I am required to abide by the Regulations of the University and to conform to its discipline.

* *Please note that it is the responsibility of the candidate to ensure that the correct version of the thesis is submitted.*

Signature of Candidate*:*		Date:	

Submission

Submitted By *(name in capitals):*	Richard Alexander McCracken
Signature of Individual Submitting:	
Date Submitted:	

For Completion in the Student Service Centre (SSC)

Received in the SSC by *(name in capitals):*			
Method of Submission *(Handed in to SSC; posted through internal/external mail):*			
*E-thesis Submitted (**mandatory for final theses)***			
Signature:		Date:	

Please note this form should bound into the submitted thesis.

Updated February 2008, November 2008, February 2009, January 2011

Acknowledgements

The work carried out over the last 3 and a half years would not have been possible without the generous support and advice of a number of people, so I would like to thank them here.

I owe so much to Professor Derryck Reid, my supervisor and guide through the world of ultrafast photonics. Your knowledge, wisdom and patience have no limit, and I will always be grateful for your tutelage and your vision.

I would like to thank Professor Jinghua Sun, who supported me through the first stages of my project. The work you have put into this group can be seen every day, and we all thank you for your guidance and wish you the best of luck in China.

To Dr Tobias Lamour, who joined me in fatherhood in the final months of his PhD; thank you for your advice, your discussions and your guide to making good coffee.

To Dr Euan Ramsay, who helped Louise and I in no small way when our little surprise came along; thank you for everything.

To the other members of the Ultrafast Optics research group; Dr Christopher Leburn, Dr Carl Farrell, Dr Zhaowei Zhang, Cristtel Ramirez, Thomas Schratwieser and Karolis Balskus. We are an international and diverse bunch, but we are a great team. Thank you all for your help over the years.

To the mechanical workshop technicians, Mark Stewart, Tony Aitken and Peter Heron; your time and effort have been invaluable, especially considering my poor drawing skills.

To Dr Lynn Paterson; thank you for being my ride home and letting me stay late! Let's hope Daniel and Freya become good friends at nursery.

Finally, to my family; when I started my PhD I had no idea what was just around the corner, but you helped Louise and I find our feet in a shaky situation. We couldn't have done it without you, and I couldn't have gotten this far without her. I love you all.

List of publications by the candidate

Peer reviewed journal articles

R. A. McCracken, J. Sun, C. G. Leburn, and D. T. Reid, "Broadband phase coherence between an ultrafast laser and an OPO using lock-to-zero CEO stabilization," *Optics Express* **20** (15), 16269-16274 (2012).

R. A. McCracken and D. T. Reid, "Characterisation of a liquid-crystal pulse shaper over a 0.36-PHz bandwidth," submitted.

R. A. McCracken, I. Gianani, C. Farrell, I. A. Walmsley, and D. T. Reid, "Multi-coloured coherent pulse synthesis," in preparation.

Conference submissions

R. A. McCracken, J. Sun, C. G. Leburn, and D. T. Reid, "Towards optical attosecond pulses: broadband phase coherence between an ultrafast laser and an OPO using lock-to-zero CEO stabilization," XVIIIth International Conference on Ultrafast Phenomena (UP2012), poster (2012).

R. A. McCracken, J. Sun, C. G. Leburn, and D. T. Reid, "Towards optical attosecond pulses: broadband phase coherence between an ultrafast laser and an OPO using lock-to-zero CEO stabilization," Ultrafast Optics (2013).

R. A. McCracken and D. T. Reid, "Characterisation of a liquid-crystal pulse shaper over a 0.36-PHz bandwidth," Ultrafast Optics, poster (2013).

R. A. McCracken, J. Sun, C. G. Leburn, and D. T. Reid, "Broadband phase coherence between an ultrafast laser and an OPO using lock-to-zero CEO stabilization," Conference on Lasers and Electro-Optics (CLEO) Europe (2013).

R. A. McCracken and D. T. Reid, "Characterisation of a liquid-crystal pulse shaper over a 0.36-PHz bandwidth," Conference on Lasers and Electro-Optics (CLEO) Europe (2013), poster.

R. A. McCracken, J. Sun, C. G. Leburn, and D. T. Reid, "Broadband phase coherence between an ultrafast laser and an OPO using lock-to-zero CEO stabilization," Conference on Lasers and Electro-Optics (CLEO) (2013).

Contents

Chapter 1 - Introduction

1.1 Aim

To date, attosecond pulse generation has relied on high-harmonic generation (HHG) in noble gases to produce either trains of attosecond pulses [1] or, more recently, isolated pulses [2] in the XUV / soft-X-ray region. Attosecond science is therefore confined to the ionisation (non-perturbative) regime due to the high photon energy of the probe pulses. Generation of optical attosecond pulses is an appealing prospect, allowing perturbative measurements on atomic timescales and also enabling the use of comparatively cheaper optical components. Optical attosecond pulses can be generated by coherently combining femtosecond pulses across a suitable bandwidth in the UV to near-IR region [3]. In contrast to HHG methods, an all-solid-state approach promises high-efficiency pulse generation at high repetition rates using accessible femtosecond laser technology to create attosecond pulses removed from the XUV / soft-X-ray region and which can therefore propagate freely in air.

The proposed technical approach that will be examined in this thesis is to construct visible attosecond pulses by coherent pulse synthesis from sequences of distinct parent pulses. This method requires four critical prerequisites to be satisfied. First, the parent pulses must have matching repetition frequencies; second, phase coherence must be established across the entire bandwidth of the parent pulses; third, the timing jitter between the pulses must be controlled to attosecond precision; finally, the group delay dispersion must be corrected for each pulse and between different pulses.

The work reported in this thesis will address each of these prerequisites in turn in order to determine the viability of generating optical attosecond pulses through coherent pulse synthesis.

1.2 Coherent pulse synthesis

Coherent pulse synthesis aims to combine two or more sequences of phase coherent narrowband pulses in order to create a new sequence of identical broadband pulses. Coherently synthesised pulses have been generated using a number of complementary methods which are discussed below.

1.2.1 Coherent synthesis through Raman resonances

Few-cycle and single-cycle pulses have been created [4–11] by using an adiabatically-driven Raman resonance [12] to produce monochromatic sidebands that are linearly spaced in frequency and whose phases can be manipulated using a standard Fourier-domain modulator [13]. This approach, first implemented by Sokolov, Yavuz and Harris, required high-energy nanosecond pulses and resulted in a train of around 100,000 synthesised pulses that lasted for the duration of the pump pulses. A typical experimental schematic for producing such a pulse train is shown in Figure 1.1 (a). Two nanosecond fields of frequencies ω_1 and ω_2 drive a Raman transition of frequency ω_{ab} at a detuning of δ. The generated collinear Raman sidebands are compressed in a pulse shaper in order to synthesise a train of transform-limited subfemtosecond pulses with repetition rate ω_{ab}.

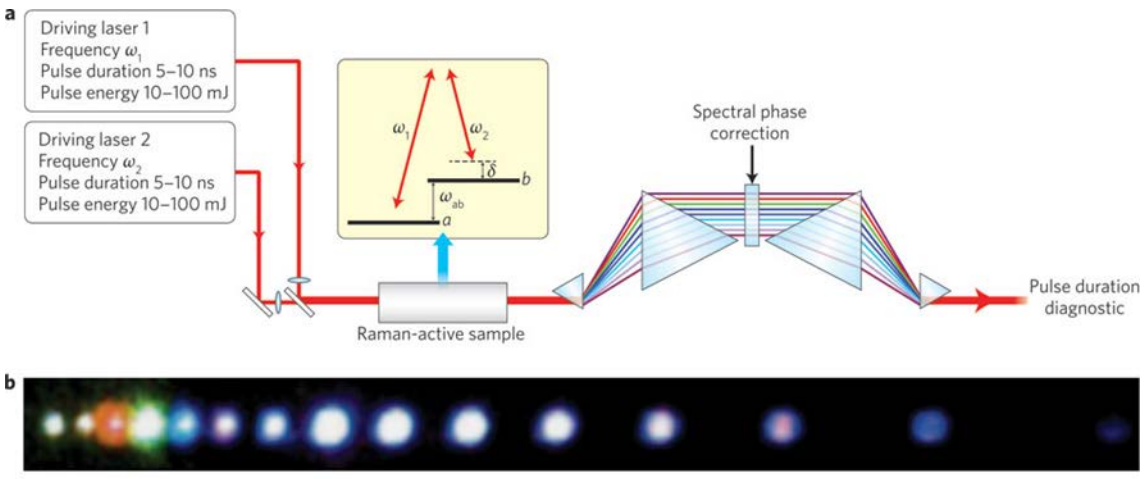

Figure 1.1. (a) A typical experimental schematic for producing few-cycle pulses by molecular modulation; (b) multiple collinear Raman sidebands spanning the UV to near-infrared. Upper image taken from [14]; lower image taken from [4].

1.2.2 Coherent synthesis from mode-locked lasers and amplifiers

Coherent synthesis can also be achieved by combining the pulse sequences from two distinct laser oscillators. This was demonstrated by Shelton et al. in 2001 [15,16], where the authors synchronised the phases and repetition rates of two mode-locked Ti:sapphire lasers operating at different centre wavelengths. The experimental layout from [15] is shown in Figure 1.2. Two phase-locked loops working at different time resolutions were used to synchronise the repetition rates of the lasers. Heterodyne beating between the two lasers at an avalanche photodiode provided the difference in their carrier-envelope offset (CEO) frequencies. A third feedback loop was used to control the CEO frequency of one oscillator, phase locking the lasers. Autocorrelation measurements were performed on the combined beam from the two lasers; when the lasers were synchronised and phase locked the resulting autocorrelation was narrowed compared to the unlocked, unsynchronised measurement, showing that a shorter pulse had been synthesised.

Further work from the same group reduced the timing jitter between the lasers to sub-femtosecond durations, paving the way for attosecond pulse synthesis [17,18]. In complementary work a single Erbium-doped fiber laser was used to seed two EDFAs which generated phase-coherent supercontinua that were combined to synthesise a single-cycle pulse in the near-infrared [19]. Synthesis has also been demonstrated between lasers with different gain media, including Ti:sapphire and Cr:forsterite lasers [20–22], and a Ti:sapphire and an Erbium-doped fiber laser [23]. Few-cycle waveforms have been synthesised in work carried out by [24] and [25]. Coherent synthesis has been demonstrated by combining a number of different optical parametric amplifiers [26–28] for application in strong-field physics experiments.

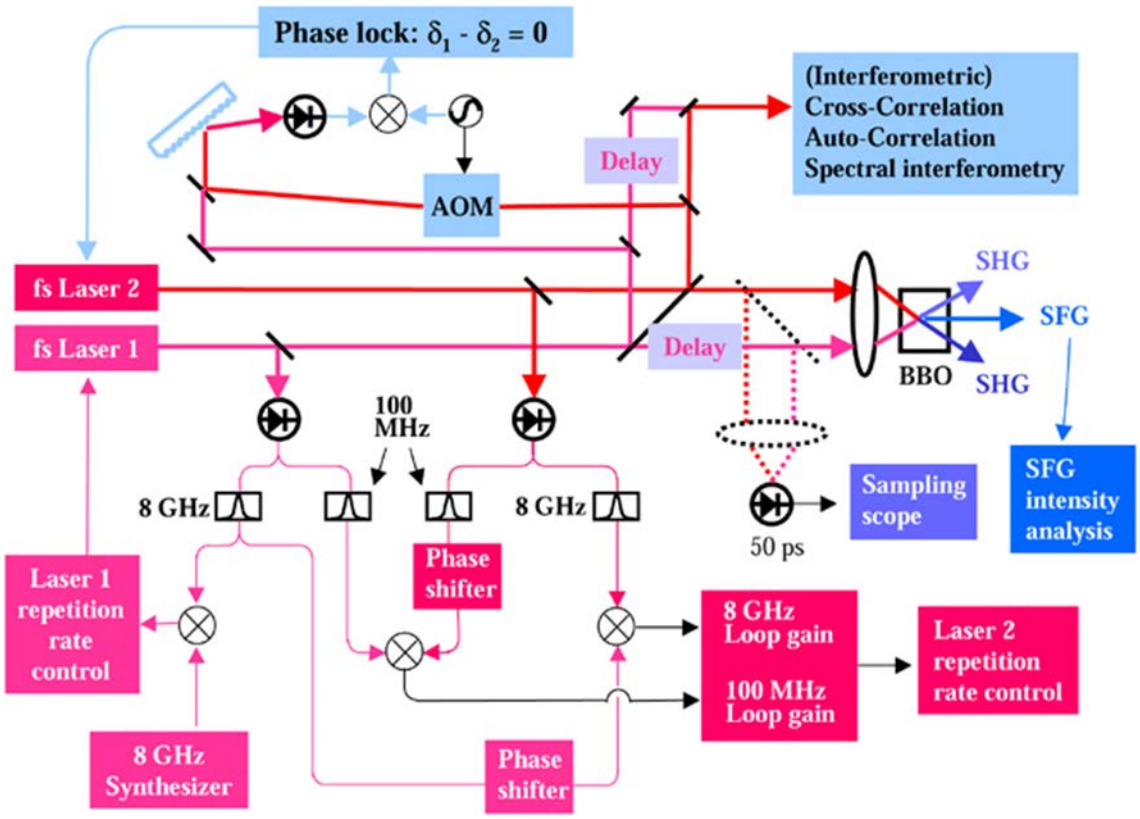

Figure 1.2. Experimental schematic of the first demonstration of coherent pulse synthesis between two independent mode-locked lasers. Taken from supplemental material from [15].

1.2.3 Coherent synthesis between a laser and an optical parametric oscillator

In work previously carried out by the Reid group, a Ti:sapphire-pumped synchronous optical parametric oscillator (OPO) based on a cascaded-grating MgO:PPLN crystal [29] was used to demonstrate pulse synthesis between the depleted pump pulses an intracavity-doubled-signal pulses from the OPO [30–32]. The CEO frequencies of both the pump and doubled signal pulses were locked to the same value. A power spectral density measurement of the CEO frequency stability and an optical cross-correlation measurement between the pump and doubled signal pulses showed that they remained coherent over 1.4 ms with a mutual timing jitter of 30 as in 20 ms. The CEO frequency locking techniques used in this work limited the number of synthesis parent pulses to two, motivating the work carried out in this thesis.

1.2.4 Attosecond pulse synthesis from an optical parametric oscillator

In order to synthesise pulses with attosecond durations an equivalent frequency bandwidth must be spanned. This bandwidth is not achievable directly from any laser oscillator; however it is within the transparency window of several nonlinear crystals.

An OPO is an ideal way to exploit this bandwidth, as multiple sum-frequency-mixing and second-harmonic-generation outputs can be simultaneously generated in the OPO crystal, as listed in Table 1.1. Provided that the prerequisites given above can be met for a set of sufficiently short parent pulses, coherently synthesised attosecond pulses in the visible region could be achievable, as shown in Figure 1.3. Combining the pulses listed in Table 1 would produce a high-contrast 700 as field transient, provided the parent pulses were each of 15 fs duration and were transform limited.

Table 1.1. Fundamental, SFM and SHG mixing outputs available from an OPO crystal with a transparency band from the UV to the mid-IR.

Wavelength (μm)	0.200	0.400	0.456	0.530	0.642	0.800	1.060	3.262
Origin	$4\omega_p$	$2\omega_p$	$\omega_p + \omega_s$	$2\omega_s$	$\omega_p + \omega_i$	ω_p	ω_s	ω_i

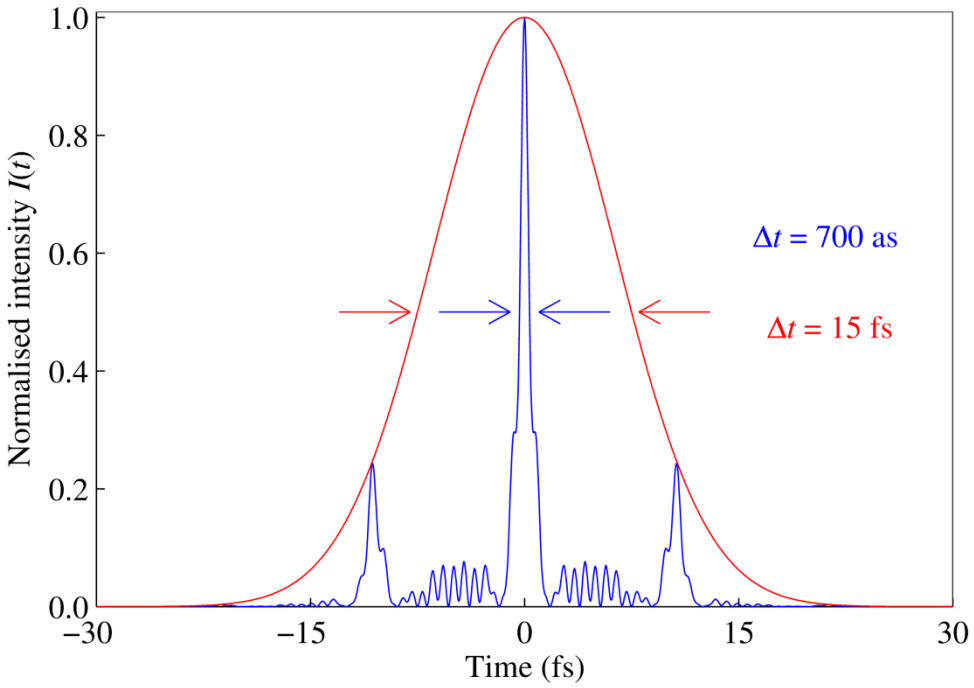

Figure 1.3. Electric field profile of a pulse synthesised from the parent pulses listed in Table 1.

1.3 Thesis outline

In this thesis the results from the design and characterisation of a femtosecond OPO for broadband coherent pulse synthesis are presented. The carrier-envelope-offset frequencies of the pump, signal, idler and various frequency-mixing outputs of the OPO were stabilised and the relative timing jitter between the pulses was calculated. Three pulses were selected as candidate parent pulses for pulse synthesis.

Chapter 2 reviews a number of important concepts in nonlinear optics, including phasematching, parametric generation, ultrashort pulse dynamics and pulse characterisation techniques.

Chapter 3 describes the characterisation of a Ti:sapphire laser used to synchronously pump a femtosecond OPO. The design, construction and characterisation of two OPOs were considered, and relevant results are presented.

Chapter 4 discusses the carrier-envelope-offset (CEO) frequency and various methods for its stabilisation. CEO frequencies from both the pump and OPO were detected using a nonlinear interferometer; zero-offset CEO locking was established through a pair of feedback loops and confirmed through temporal interferometry.

Chapter 5 presents a number of noise analyses from the pump and OPO, including relative intensity noise, CEO-locking phase noise and relative timing jitter between the pulses.

Chapter 6 describes the requirement to compress the parent pulses used for coherent synthesis, as well as controlling their relative group delays. A prism delay line was constructed in order to compress the pulses and subsequent pulse duration results are presented. A liquid-crystal spatial light modulator was characterised and is presented as an alternative to the prism delay line.

Chapter 7 outlines the main conclusions from the work presented in this thesis and establishes whether broadband coherent synthesis from an OPO is a viable path to achieving visible attosecond pulses. Improvements to the work carried out are presented and future experiments are suggested.

1.4 References

1. P. M. Paul, E. S. Toma, P. Breger, G. Mullot, F. Auge, P. Balcou, H. G. Muller, and P. Agostini, "Observation of a train of attosecond pulses from high harmonic generation," Science (80-.). **292**, 1689–92 (2001).

2. G. Sansone, E. Benedetti, F. Calegari, C. Vozzi, L. Avaldi, R. Flammini, L. Poletto, P. Villoresi, C. Altucci, R. Velotta, S. Stagira, S. De Silvestri, and M. Nisoli, "Isolated single-cycle attosecond pulses," Science (80-.). **314**, 443–6 (2006).

3. T. W. Hänsch, "A proposed sub-femtosecond pulse synthesizer using separate phase-locked laser oscillators," Opt. Commun. **80**, 71–75 (1990).

4. A. Sokolov, D. Walker, D. Yavuz, G. Yin, and S. Harris, "Femtosecond light source for phase-controlled multiphoton ionization," Phys. Rev. Lett. **87**, 033402 (2001).

5. M. Shverdin, D. Walker, D. Yavuz, G. Yin, and S. Harris, "Generation of a single-cycle optical pulse," Phys. Rev. Lett. **94**, 033904 (2005).

6. D. Yavuz, D. Walker, M. Shverdin, G. Yin, and S. Harris, "Quasiperiodic Raman technique for ultrashort pulse generation," Phys. Rev. Lett. **91**, 233602 (2003).

7. W.-J. Chen, Z.-M. Hsieh, S. Huang, H.-Y. Su, C.-J. Lai, T.-T. Tang, C.-H. Lin, C.-K. Lee, R.-P. Pan, C.-L. Pan, and A. Kung, "Sub-single-cycle optical pulse train with constant carrier envelope phase," Phys. Rev. Lett. **100**, 163906 (2008).

8. A. V Sokolov and S. E. Harris, "Ultrashort pulse generation by molecular modulation," J. Opt. B Quantum Semiclassical Opt. **5**, R1–R26 (2003).

9. M. Wittmann, A. Nazarkin, and G. Korn, "Fs-pulse synthesis using phase modulation by impulsively excited molecular vibrations," Phys. Rev. Lett. **84**, 5508–11 (2000).

10. A. V Sokolov, D. D. Yavuz, and S. E. Harris, "Subfemtosecond pulse generation by rotational molecular modulation," Opt. Lett. **24**, 557–9 (1999).

11. A. V Sokolov and S. E. Harris, "Subfemtosecond pulse generation by rotational molecular modulation," Phys. Rev. B **81**, 2894–2897 (1998).

12. S. Harris, D. Walker, and D. Yavuz, "Raman technique for single-cycle pulses," Phys. Rev. A **65**, 021801 (2002).

13. A. M. Weiner, "Femtosecond pulse shaping using spatial light modulators," Rev. Sci. Instrum. **71**, 1929–1960 (2000).

14. S. Baker, I. A. Walmsley, J. W. G. Tisch, and J. P. Marangos, "Femtosecond to attosecond light pulses from a molecular modulator," Nat. Photonics **5**, 664–671 (2011).

15. L.-S. Ma, R. Shelton, H. Kapteyn, M. Murnane, and J. Ye, "Sub-10-femtosecond active synchronization of two passively mode-locked Ti:sapphire oscillators," Phys. Rev. A **64**, 021802 (2001).

16. R. K. Shelton, L. S. Ma, H. C. Kapteyn, M. M. Murnane, J. L. Hall, and J. Ye, "Phase-coherent optical pulse synthesis from separate femtosecond lasers," Science (80-.). **293**, 1286–9 (2001).

17. R. K. Shelton, S. M. Foreman, L.-S. Ma, J. L. Hall, H. C. Kapteyn, M. M. Murnane, M. Notcutt, and J. Ye, "Subfemtosecond timing jitter between two independent, actively synchronized, mode-locked lasers," Opt. Lett. **27**, 312–4 (2002).

18. J. Ye, S. T. Cundiff, S. Foreman, T. M. Fortier, J. L. Hall, K. W. Holman, D. J. Jones, J. D. Jost, H. C. Kapteyn, K. a. H. v. Leeuwen, L. S. Ma, M. M. Murnane, J. L. Peng, and R. K. Shelton, "Phase-coherent synthesis of optical frequencies and waveforms," Appl. Phys. B **74**, s27–s34 (2002).

19. G. Krauss, S. Lohss, T. Hanke, A. Sell, S. Eggert, R. Huber, and A. Leitenstorfer, "Synthesis of a single cycle of light with compact erbium-doped fibre technology," Nat. Photonics **4**, 33–36 (2009).

20. Z. Wei, Y. Kobayashi, and K. Torizuka, "Relative carrier-envelope phase dynamics between passively synchronized Ti:sapphire and Cr:forsterite lasers," Opt. Lett. **27**, 2121–3 (2002).

21. A. Bartels, N. R. Newbury, I. Thomann, L. Hollberg, and S. A. Diddams, "Broadband phase-coherent optical frequency synthesis with actively linked Ti:sapphire and Cr:forsterite femtosecond lasers," Opt. Lett. **29**, 403–5 (2004).

22. Y. Kobayashi, D. Yoshitomi, M. Kakehata, H. Takada, and K. Torizuka, "Long-term optical phase locking between femtosecond Ti:sapphire and Cr:forsterite lasers," Opt. Lett. **30**, 2496–8 (2005).

23. J. A. Cox, W. P. Putnam, A. Sell, A. Leitenstorfer, and F. X. Kärtner, "Pulse synthesis in the single-cycle regime from independent mode-locked lasers using attosecond-precision feedback," Opt. Lett. **37**, 3579–81 (2012).

24. A. Poppe, R. Holzwarth, A. Apolonski, G. Tempea, C. Spielmann, T. W. Hänsch, and F. Krausz, "Few-cycle optical waveform synthesis," Appl. Phys. B **72**, 373–6 (2001).

25. A. Wirth, M. T. Hassan, I. Grgura, J. Gagnon, A. Moulet, T. T. Luu, S. Pabst, R. Santra, Z. A. Alahmed, A. M. Azzeer, V. S. Yakovlev, V. Pervak, F. Krausz, and E. Goulielmakis, "Synthesized light transients," Science (80-.). **334**, 195–200 (2011).

26. S.-W. Huang, G. Cirmi, J. Moses, K.-H. Hong, S. Bhardwaj, J. R. Birge, L.-J. Chen, E. Li, B. J. Eggleton, G. Cerullo, and F. X. Kärtner, "High-energy pulse synthesis with sub-cycle waveform control for strong-field physics," Nat. Photonics **5**, 475–479 (2011).

27. C. Manzoni, S.-W. Huang, G. Cirmi, P. Farinello, J. Moses, F. X. Kärtner, and G. Cerullo, "Coherent synthesis of ultra-broadband optical parametric amplifiers," Opt. Lett. **37**, 1880–2 (2012).

28. A. Baltuška, T. Fuji, and T. Kobayashi, "Controlling the carrier-envelope phase of ultrashort light pulses with optical parametric amplifiers," Phys. Rev. Lett. **88**, 133901 (2002).

29. F. Ruebel, P. Haag, and J. A. L'huillier, "Synchronously pumped femtosecond optical parametric oscillator with integrated sum frequency generation," Appl. Phys. Lett. **92**, 011122 (2008).

30. Y. Kobayashi, H. Takada, M. Kakehata, and K. Torizuka, "Optical phase locking among femtosecond subharmonic pulses," Opt. Lett. **28**, 1377–9 (2003).

31. B. J. S. Gale, J. H. Sun, and D. T. Reid, "Towards versatile coherent pulse synthesis using femtosecond laser and optical parametric oscillators," Opt. Express **16**, 1616–22 (2008).

32. J. Sun and D. T. Reid, "Coherent ultrafast pulse synthesis between an optical parametric oscillator and a laser," Opt. Lett. **34**, 854–6 (2009).

Chapter 2 - Fundamentals of optical parametric oscillators and ultrashort pulse dynamics

2.1 Introduction

Nonlinear interactions in bulk crystals and dielectric media are fundamental to the generation and propagation of ultrashort pulses. An understanding of the spatial and temporal evolution of such pulses in a resonant cavity, and as they propagate through optical material, is crucial for the fine dispersion control necessary to achieve ultrashort pulse durations. In this chapter the reader is introduced to the basic principles of nonlinear optics, including the fundamental concepts of phasematching and dispersion, as well as a mathematical description of a laser pulse. Consideration is then given to some of the methods of characterizing ultrashort pulses, including autocorrelation, FROG, SPIDER and the powerful MIIPS technique.

2.2 Fundamental nonlinear properties and concepts

This section of the chapter describes the interactions that take place in a nonlinear material when subject to an intense electromagnetic field. Discussion is given to $\chi^{(2)}$ processes through which ultrashort pulses are generated in a nonlinear medium, with emphasis on phasematched parametric processes found in an optical parametric oscillator.

2.2.1 Nonlinear susceptibility

When light propagates through a material the electrons and atoms within it react to the electromagnetic fields of the wave, producing a change in the spatial and temporal distribution of electrical charges. The field causes small displacements of the valence electrons from their normal orbits, perturbations which create electric dipoles that are manifested as the electric polarization of the medium. For small field strengths (such as an unfocused, low energy laser beam) this electric polarization $P(\omega)$ is proportional to the electric field strength $E(\omega)$, expressed as

$$P(\omega) = \varepsilon_0 \chi^{(1)} E(\omega) \tag{2.1}$$

where ε_0 is the vacuum permittivity and $\chi^{(1)}$ is the first-order or linear susceptibility, a measure of the sensitivity of the electric polarization to the applied field [1]. For intense

electric fields a nonlinear polarization of the medium takes place, described as a power series expansion of Equation (2.1) and often written as

$$P(\omega) = \varepsilon_0 \chi^{(1)} E(\omega) + \varepsilon_0 \chi^{(2)} E^2(\omega) + \varepsilon_0 \chi^{(3)} E^3(\omega) + \ldots, \tag{2.2}$$

where $\chi^{(2)}$ and $\chi^{(3)}$ are the second-order and third-order nonlinear susceptibilities respectively.

2.2.2 Properties of the $\chi^{(2)}$ nonlinear susceptibility

Optical field interactions involving the $\chi^{(2)}$ nonlinearity are phase-only, with no photon energy being absorbed into the medium. This makes $\chi^{(2)}$ nonlinear processes appealing as they require no cooling and can be highly efficient. While all transparent materials display first- and third-order nonlinear susceptibility, second-order effects are only observed in non-centrosymmetric crystals. Such nonlinear crystals can be used in power-scalable processes that are limited only by the material properties of the crystal, such as the damage threshold, hygroscopicity and photorefraction.

2.2.3 Nonlinear optical processes

The second-order nonlinear susceptibility can be exploited to produce a number of nonlinear interactions. A formal derivation of these processes will be given in Sub-section 2.2.4; however a more general outline will be given first here. The most common application of the $\chi^{(2)}$ nonlinearity is to facilitate frequency mixing between two electromagnetic waves to produce an interfering field *E*, given by

$$E = E_1 \cos(\omega_1 t) + E_2 \cos(\omega_2 t). \tag{2.3}$$

The field *E* will generate a polarisation response in the material as defined by the first two terms of Equation (2.2), given by

$$\begin{aligned} P(\omega_1, \omega_2) = {} & \varepsilon_0 \chi^{(1)} \left[E_1 \cos(\omega_1 t) + E_2 \cos(\omega_2 t) \right] \\ & + \varepsilon_0 \chi^{(2)} \left[2E_1 E_2 \cos(\omega_1 t) \cos(\omega_2 t) \right. \\ & \left. + E_1^2 \cos^2(\omega_1 t) + E_2^2 \cos^2(\omega_2 t) \right] \end{aligned} \quad . \tag{2.4}$$

Using the trigonometric identities $\cos(A)\cos(B) = \frac{\cos(A-B)+\cos(A+B)}{2}$ and $\cos^2(A) = \frac{1+\cos(2A)}{2}$, Equation (2.4) can be expanded into six independent terms which describe the nonlinear interactions between the two fields, given by

$$P = \varepsilon_0 \chi^{(1)} \left[E_1 \cos(\omega_1 t) + E_1 \cos(\omega_2 t) \right] \text{ (INT)} \tag{2.5}$$

$$+ \frac{\varepsilon_0 \chi^{(2)} E_1 E_2}{2} \left[\cos\left([\omega_1 - \omega_2] t \right) \right] \text{ (DFG)} \tag{2.6}$$

$$+ \frac{\varepsilon_0 \chi^{(2)} E_1 E_2}{2} \left[\cos\left([\omega_1 + \omega_2] t \right) \right] \text{ (SFG)} \tag{2.7}$$

$$+ \frac{\varepsilon_0 \chi^{(2)} E_1^2}{2} \left[\cos\left(2\omega_1 t \right) \right] \text{ (SHG)} \tag{2.8}$$

$$+ \frac{\varepsilon_0 \chi^{(2)} E_2^2}{2} \left[\cos\left(2\omega_2 t \right) \right] \text{ (SHG)} \tag{2.9}$$

$$+ \frac{\varepsilon_0 \chi^{(2)}}{2} \left[E_1^2 + E_2^2 \right] \text{ (EOR).} \tag{2.10}$$

Equation (2.5) describes the interference field (INT) between the incident waves, manifested as the linear polarisation response of the medium. Equations (2.6) and (2.7) respectively describe difference-frequency generation (DFG) and sum-frequency generation (SFG) between the two waves. Equations (2.8) and (2.9) correspond to second harmonic generation (SHG) for each of the incident frequencies ω_1 and ω_2 respectively. The final term, Equation (2.10), describes the DC polarization component known as electro-optic rectification (EOR), a process exploited in the field of terahertz generation.

2.2.4 The coupled wave equations

The previous sub-section introduced the nonlinear polarization response of a medium to an intense electric field and how this interaction creates new frequencies. In order to fully describe how these new frequencies are generated and their relation to the frequencies of the incident frequency fields we must extend the standard wave equation using Maxwell's equations, expressed as

$$\nabla^2 E - \frac{1}{c^2}\frac{\partial^2 E}{\partial t^2} = \frac{1}{\varepsilon_0 c^2}\frac{\partial^2 P}{\partial t^2}. \tag{2.11}$$

This equation can be simplified by considering that the incident waves only travel in one propagation direction, commonly chosen as *z*, and by limiting the interaction to three frequencies ω_1, ω_2 and ω_3, where convention dictates that

$$\omega_1 \le \omega_2 \le \omega_3. \tag{2.12}$$

The simplified equation has a solution that produces three coupled wave equations that describe any second-order frequency mixing process, given by

$$\begin{aligned}
\frac{dE_{1i}}{dz} &= -\frac{i\omega_1}{cn_1} d'_{ijk} E_{3j} E^*_{2k} e^{(-i\Delta kz)}, \\
\frac{dE^*_{2k}}{dz} &= \frac{i\omega_2}{cn_2} d'_{kij} E_{1i} E^*_{3j} e^{(+i\Delta kz)}, \quad \text{and} \\
\frac{dE_{3j}}{dz} &= -\frac{i\omega_3}{cn_3} d'_{jik} E_{1i} E_{2k} e^{(+i\Delta kz)},
\end{aligned} \tag{2.13}$$

where *i*, *j* and *k* refer to Cartesian coordinates and can take on values *x* and *y*, and where d_{ijk} is the tensor element describing the nonlinear coupling between the fields, with a value of half the susceptibility tensor [2]. Δk is the wave-vector mismatch, given by

$$\Delta k = k_3 - k_2 - k_1 \tag{2.14}$$

where

$$k_m = \frac{2\pi n(\lambda_m)}{\lambda_m}. \tag{2.15}$$

Efficient generation of new frequencies is only possible if $\Delta k = 0$, a condition that is achieved through careful choice of propagation directions through the medium and consideration of its material properties. Satisfying this condition is known as phasematching. The traditional method to achieve phasematching is through use of the birefringence of a nonlinear crystal such as lithium niobate ($LiNbO_3$), and is commonly referred to as birefringent phasematching.

2.2.5 Birefringent phasematching

For interactions between three waves there are two possible birefringent phasematching conditions. In Type I phasematching the two longer wavelengths have the same polarization, perpendicular to that of the shorter generated wavelength, such as

$$\begin{matrix} oo \rightarrow e \\ ee \rightarrow o \end{matrix}. \tag{2.16}$$

This phasematching is common in second harmonic generation. Using Equations (2.14) and (2.15) it can be seen that the phasematching condition for Type I SHG is given by

$$\Delta k = 2\pi \left[\frac{n_{2\omega}^{o(e)}}{\lambda_{2\omega}} - \frac{n_{\omega}^{e(o)}}{\lambda_{\omega}} - \frac{n_{\omega}^{e(o)}}{\lambda_{\omega}} \right] = 0. \tag{2.17}$$

This is only possible in materials where $n_{2\omega}^{o(e)} = 2n_{\omega}^{e(o)}$, a condition that can be met in many crystals but is highly limited by the exact dispersion relation of the material, as shown in Figure 2.1. Fortunately some degree of tunability in the phasematching condition can be achieved by adjusting the angle of the crystal relative to the propagation direction. Light polarized in the plane containing the optic axis and the direction of propagation will experience an angle-dependent refractive index that varies between n_e and n_o. The correct angle necessary for satisfying Equation (2.17) can be found by adjusting the crystal angle from normal incidence or by cutting the crystal so that the new normal interface becomes the phasematching angle.

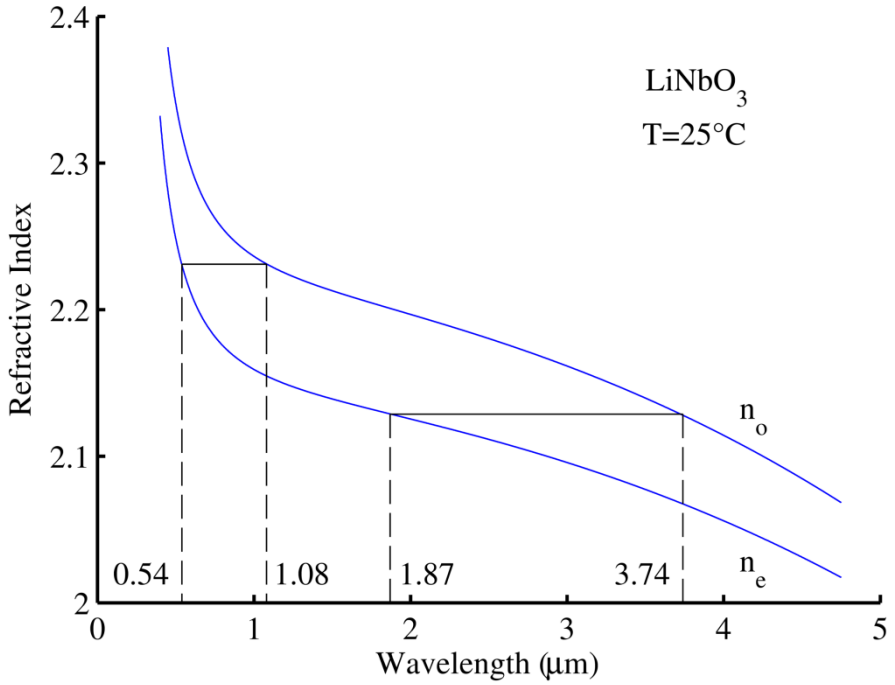

Figure 2.1. Type I birefringent phasematching solutions for SHG in lithium niobate at 25°C. The horizontal lines indicate those interactions which simultaneously satisfy the phasematching and energy conservation conditions of Equation (2.17). Solutions are obtained for SHG of 0.54 μm from 1.08 μm and 1.87 μm from 3.74 μm. Adapted from [3].

Angle tuning introduces walk-off in the Poynting vector of the generated wave which propagates at an angle to the optic axis. Walk-off reduces the spatial overlap between the interacting waves, lowering efficiency. Another downside to angle tuning is that Δk varies rapidly with angle, constraining the acceptance angle for phasematching. An alternative to angle tuning is to utilize the temperature dependence of the refractive indices. In many optical crystals the rate of change of refractive index temperature is different for the ordinary and extraordinary indices, allowing for tuning of the birefringence, however the temperatures required to meet the desired phasematching conditions may be inconvenient.

In Type II phasematching the two longer wavelength waves have perpendicular polarisations, with the shorter generated wavelength having either polarization, the orientation of which depends on the nonlinear crystal properties. Notation is given by

$$\begin{matrix} oe \rightarrow e \\ oe \rightarrow o \end{matrix}. \qquad (2.18)$$

Satisfying Equation (2.14) for Type II phasematching follows the same methodology as for Type I phasematching, but will not be discussed further here.

2.2.6 Quasi-phasematching

An alternative approach to achieving efficient frequency generation in a second-order nonlinear interaction is the case where $\Delta k \neq 0$, known as quasi-phasematching [3]. This process takes advantage of the fact that all $\chi^{(2)}$ materials are dispersive, a concept discussed in Sub-section 2.3.2. Again considering the case of three interacting electromagnetic waves, the relative phase between them after some propagation distance will accumulate to π, given as

$$\Delta\phi = \phi_3 - \phi_2 - \phi_1 = \pi \,. \tag{2.19}$$

The conversion efficiency builds to a maximum over this propagation distance which is known as the coherence length, given by

$$L_c = \frac{\pi}{\Delta k}. \tag{2.20}$$

After the waves propagate beyond the coherence length the converted intensity decreases as the phasematching condition has changed. In fact the condition is such that energy is converted from the generated frequencies back into the fundamental frequency, a process known as back conversion. To prevent back conversion and increase the efficiency of the forward conversion, a phase step of π is added to the interacting fields, returning to the phasematching condition. Physically this phase step is added by periodically flipping the polarity of the nonlinear coefficient tensor d_{ijk} by the quasi-phasematching or grating period Λ_g, given by

$$\Lambda_g = \frac{2\pi}{\Delta k}. \tag{2.21}$$

The quasi-phasematching condition, a modified version of Equation (2.14), is therefore given by

$$\Delta k_{qpm} = k_3 - k_2 - k_1 - \frac{2\pi}{\Lambda_g} = 0 \,. \tag{2.22}$$

Satisfying this condition can be achieved by varying the length of the grating period, under the condition that

$$\Lambda_g = \left[\frac{n(\lambda_3)}{\lambda_3} - \frac{n(\lambda_2)}{\lambda_2} - \frac{n(\lambda_2)}{\lambda_1} \right]^{-1}. \tag{2.23}$$

Periodic inversion of the nonlinear polarity of a material is achieved through the periodic poling technique. Ferroelectric crystals such as $KTiOPO_4$ are exposed to an intense electric field which alters the position of the ions inside the unit cell. Manufacturers can extend this technique across a crystal by using a patterned electrode to create alternating poled domains with few-micron-scale grating periods over a crystal length of several centimetres. The QPM process allows for efficient frequency conversion in nonlinear materials where phasematching is not possible. The relative efficiencies for SHG are shown in Figure 2.2 as an example.

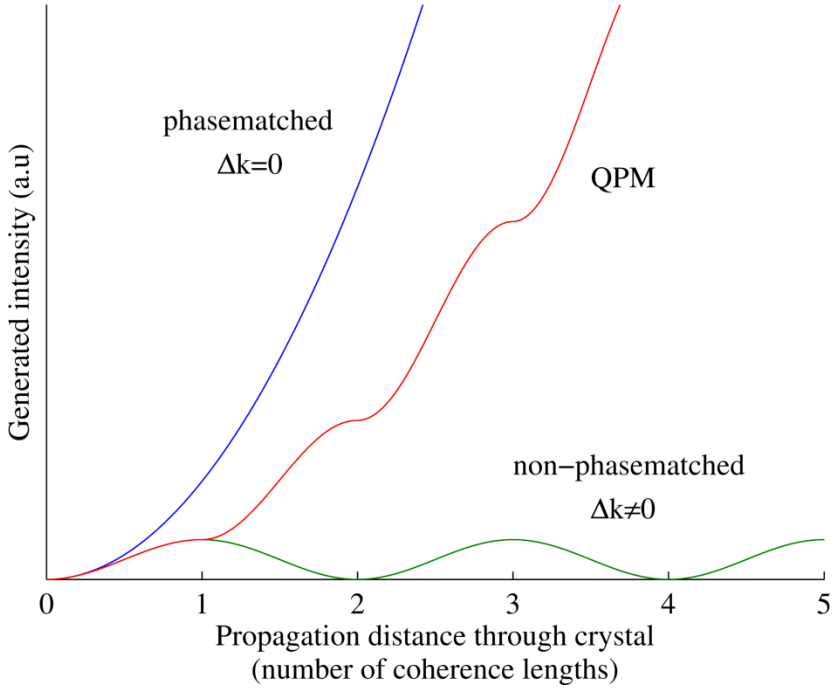

Figure 2.2. Quasi-phasematching allows efficient frequency conversion in non-phasematched crystals by introducing a π phase shift whenever the phase mismatch accrues to π.

2.2.7 Second harmonic generation

Solving the coupled wave equations is a non-trivial process, and often requires simplification before an analytical solution can be found. In the case of second harmonic generation it is assumed that the conversion efficiency is low with no depletion of the interacting waves, giving $\frac{dE_1}{dz} \approx \frac{dE_2}{dz} \approx 0$. This leaves only one equation to be solved,

$$\frac{dE_{3j}}{dz} = -\frac{i\omega_3}{cn_3} d'_{jik} E_1 E_2 e^{(+i\Delta kz)}, \tag{2.24}$$

with

$$E_3 = ae^{(+i\Delta kz)} + b, \tag{2.25}$$

and second harmonic generation constraining the frequencies to $\omega_1 = \omega_2 = \frac{\omega_3}{2}$. Setting the derivative of Equation (2.25) with Equation (2.24) we can find the constant *a*, given by

$$a = -\frac{\omega_3 d'_{jik} E_1 E_2}{cn_3 \Delta k}. \tag{2.26}$$

The analytical expression can be further simplified by noting that no coupling takes place until the waves meet the nonlinear medium, and so $E_3(z=0) = 0$. This allows *b* to be written as

$$b = -a = \frac{\omega_3 d'_{jik} E_1 E_2}{cn_3 \Delta k}, \tag{2.27}$$

giving a final generated electric field of

$$E_3 = \frac{2\omega_{1,2} d'_{jik} E_1 E_2}{cn_3 \Delta k}\left[1 - i\Delta k e^{(+i\Delta kz)}\right]. \tag{2.28}$$

The intensity of the second harmonic output is obtained by multiplying the field by its complex conjugate, that is

$$I_3 = \frac{nc\varepsilon_0 E_3 E_3^*}{2}, \tag{2.29}$$

which provides a final expression

$$I_{2\omega_{1,2}}(z) = \frac{8\omega_{1,2}^2 d'^2_{jik} I_\omega^2 z^2}{c^3 n_1 n_2 n_3 \varepsilon_0} \operatorname{sinc}^2\left[\frac{\Delta kz}{2}\right]. \tag{2.30}$$

It can be seen from Equation (2.30) that the maximum second harmonic output is generated when $\Delta k = 0$, as determined by Equation (2.14), and increases with both

incident intensity and propagation distance. For quasi-phasematched processes the second harmonic intensity follows a sinc^2 function with maximum conversion over one coherence length.

2.2.8 Sum- and difference-frequency generation

As with second harmonic frequency generation, the coupled-wave equations for sum- and difference-frequency generations are complicated to solve without simplification. The assumption is made that pump depletion is negligible and that perfect phasematching conditions are present. Following the convention established in Equation (2.12), sum-frequency mixing between waves ω_1 and ω_2 generates a new wave

$$\omega_3 = \omega_1 + \omega_2 , \tag{2.31}$$

where $\omega_3 > \omega_2 \geq \omega_1$. When considering nonlinear interactions with three different frequencies, the exchange of energy between input waves ω_1 and ω_2 and output wave ω_3 must be taken into account. This energy is transferred backward and forward as the waves propagate through the crystal with a complex gain and coupling relationship, a full derivation of which can be found in a number of nonlinear optics textbooks. A simpler summary is now given for nonlinear interactions with no wavevector mismatch and for the case where one of the applied fields (at frequency ω_2) is strong and the other field (at frequency ω_1) is weak. This example illustrates the sum-frequency mixing process that occurs between the intense pump field and the weak idler field in a signal-resonant OPO, as discussed in Sub-section 2.2.9.

As the amplitude E_2 of frequency ω_2 is much larger than E_1, it is assumed that it remains unaffected by the interaction. The coupled wave equations (2.13) then reduce to

$$\frac{dE_1}{dz} = K_1 E_3 e^{-i\Delta kz} \tag{2.32}$$

$$\frac{dE_3}{dz} = K_3 E_1 e^{+i\Delta kz} \tag{2.33}$$

where $\Delta k = k_1 + k_2 - k_3$ and the quantities $K_{1,3}$ are introduced, where

$$K_1 = \frac{2i\omega_1^2 d_{ijk}}{k_1 c^2} E_2^* \tag{2.34}$$

and

$$K_3 = \frac{2i\omega_3^2 d_{ijk}}{k_3 c^2} E_2 . \tag{2.35}$$

As no wavevector mismatch is assumed, $\Delta k = 0$. By manipulating Equations (2.32) and (2.33), an equation involving only E_1 can be obtained, written as

$$\frac{d^2 E_1}{dz^2} = -\kappa^2 E_1 \tag{2.36}$$

where κ^2 is the positive coupling coefficient, given by

$$\kappa^2 = -K_1 K_3 = \frac{4\omega_1^2 \omega_3^2 d_{ijk} \left|E_2\right|^2}{k_1 k_3 c^4} . \tag{2.37}$$

The general solution to Equation (2.36) is given as

$$E_1(z) = B\cos(\kappa z) + C\sin(\kappa z) \tag{2.38}$$

and can be combined with Equation (2.32) to get a general equation for E_3, given as

$$E_3(z) = \frac{-B\kappa}{K_1}\sin(\kappa z) + \frac{C\kappa}{K_1}\cos(\kappa z) . \tag{2.39}$$

A solution can now be found for the appropriate boundary conditions. It is assumed that the field ω_3 is not present at the input to the nonlinear crystal, and so $E_1(0) = 0$. The solutions for the weak ω_1 field and the generated ω_3 field are therefore given by

$$E_1(z) = E_1(0)\cos(\kappa z) \tag{2.40}$$

and

$$E_3(z) = -E_1(0)\frac{\kappa}{K_1}\sin(\kappa z) \tag{2.41}$$

respectively. With some thought this can be reduced to give

$$E_3(z) = i\sqrt{\frac{n_1\omega_3}{n_3\omega_1}} E_1(0)\sin(\kappa z) e^{i\phi_2} \tag{2.42}$$

where ϕ_2 is the phase of E_2.

For difference-frequency generation the same method of solution applies, with the convention that ω_3 is the pump wave and ω_1 is the mixing wave. The solutions for the weak field ω_1 and generated field ω_2 are given by

$$E_1(z) = E_1(0)\cosh(\kappa z) \tag{2.43}$$

and

$$E_2(z) = i\sqrt{\left(\frac{n_1\omega_2}{n_2\omega_1}\right)}\frac{E_3}{|E_3|}E_1^*(0)\sinh(\kappa z)\,. \tag{2.44}$$

An important result of these equations is that both the mixing wave and generated wave increase in intensity as they propagate through the nonlinear crystal, depleting the pump wave ω_3. This is important for the realisation of parametric frequency generation.

2.2.9 Optical parametric generation

The result described in the previous sub-section indicates that amplification via the $\chi^{(2)}$ nonlinearity is possible. This is achieved through the implementation of a resonant cavity known as an optical parametric oscillator (OPO) that enhances the parametric optical amplification, illustrated in Figure 2.3. The three frequencies of the three interacting waves $\omega_3 > \omega_2 \geq \omega_1$ are denoted as the pump (ω_p), signal (ω_s) and idler (ω_i) respectively. When the cavity mirrors are coated such that they are highly reflective at the signal or idler frequency then the cavity is said to be singly-resonant. If the high reflectivity regions extend to both the signal and idler waves then the cavity is said to be doubly-resonant. For the special case where the resonant signal and idler frequencies overlap at half the pump frequency, the OPO is said to be degenerate.

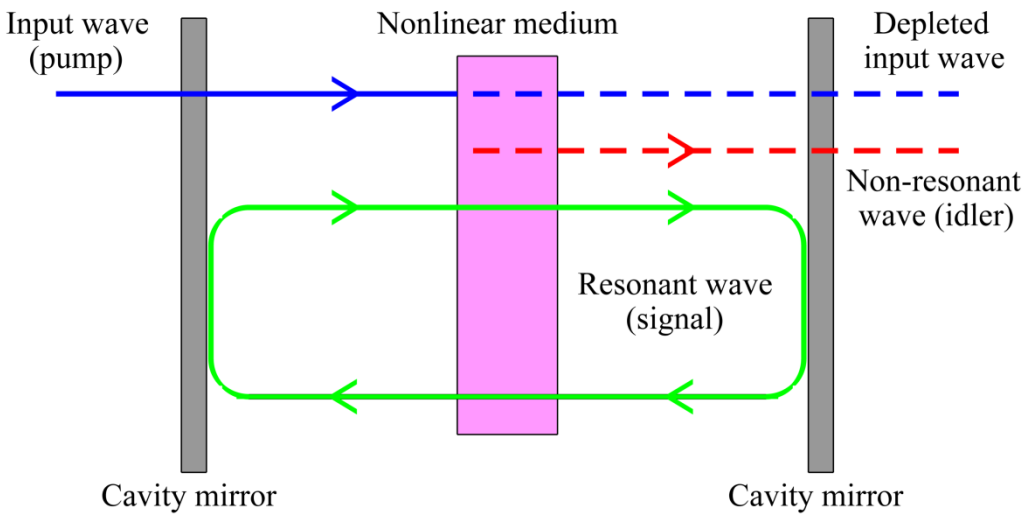

Figure 2.3. A simple illustration of a singly-resonant optical parametric oscillator. The pump wave is coupled into the nonlinear medium through a dichroic cavity mirror. The pump wave is converted into a non-resonant idler wave and a resonant signal wave. The signal wave can be coupled out of the resonator with a standard partially-reflective cavity mirror.

As with the case of difference-frequency generation, both the signal and idler waves increase in intensity as they propagate through the nonlinear crystal.

Optical parametric oscillators are widely used as a highly tunable source, limited by the phasematching conditions of the nonlinear crystal and the reflective coating of the resonator mirrors. OPOs can be operated both intra-and extra-cavity; in CW-mode or in a pulsed regime; with pulse durations ranging from nanoseconds to femtoseconds; all depending on the phasematching conditions and cavity geometry.

An extra requirement must be satisfied for pulsed operation. The OPO must be synchronously pumped, so that the cavity length of the OPO matches the repetition rate of the pump laser. This requirement arises due to the lack of absorption and gain storage in an OPO; the nonlinear gain process is instantaneous and so the generated electric fields only build up if a generated pulse coherently overlaps after one roundtrip with the next incident pump pulse inside the crystal [4]. Synchronously pumped OPOs can also be constructed where the OPO cavity length is a unit fraction [5] or improper fraction [6] of the pump cavity length. Both cases generate high repetition rate OPO pulses, but with a decrease in efficiency.

2.3 Ultrashort pulse dynamics

This section of the chapter describes in formal terms the propagation and mathematical description of ultrashort pulses. The dispersive spectral and temporal evolution of a propagating pulse is considered, as well as methods for its control.

2.3.1 Description of an ultrafast pulse

A pulse with centre frequency ω_0 can be represented by a time-dependent complex envelope function with a real electric field given by

$$E(t) \propto \mathrm{Re}\left[A(t)e^{(i\omega_0 t)} \right] \tag{2.45}$$

where $A(t)$ is the amplitude of the normalized electric field, shown in Figure 2.4 [7]. The spectral amplitude can be obtained through a Fourier transform of the temporal field and is given by

$$A(\omega) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} A(t)e^{(-i\omega t)} dt \,. \tag{2.46}$$

The instantaneous optical power and power spectrum are given by $|A(t)|^2$ and $|A(\omega)|^2$ respectively.

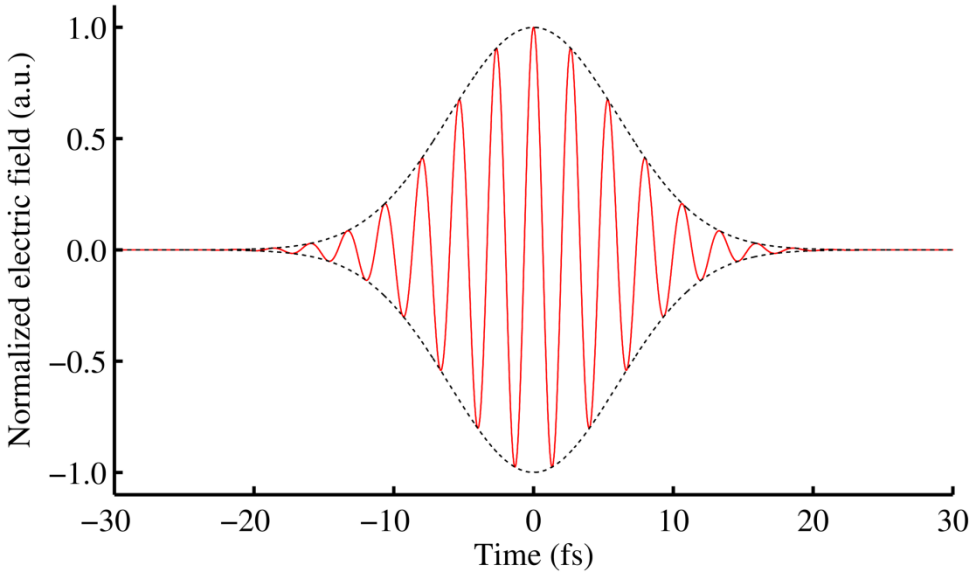

Figure 2.4. The real component of the electric field representing a Gaussian pulse with 10 fs duration at λ_0 = 0.8 µm. The dashed line represents the envelope function that is modulated by the underlying carrier wave, shown in red.

A pulse is characterised by its duration $\Delta\tau_p$ and its spectral bandwidth $\Delta\nu_p$, with a fixed relationship known as the time-bandwidth product, $\Delta\tau_p\Delta\nu_p$. The temporal and spectral widths are measured at the full-width half-maximum (FWHM) point, as shown in Figure 2.5. In ultrafast science there are two important field envelope profiles; the Gaussian shape $\left(e^{\left(\frac{-2\ln 2t^2}{\Delta\tau_p^2}\right)}\right)$ with a minimum time-bandwidth product of 0.441; and the hyperbolic secant shape $\left(\text{sech}\left(\frac{1.763t}{\Delta\tau_p}\right)^2\right)$ with a minimum time-bandwidth product of 0.315 [8]. These two time-bandwidth values correspond to a pulse that has accrued no positive or negative chirp, a process that increases the minimum time-bandwidth product. Such a chirp-free pulse is often described as transform-limited.

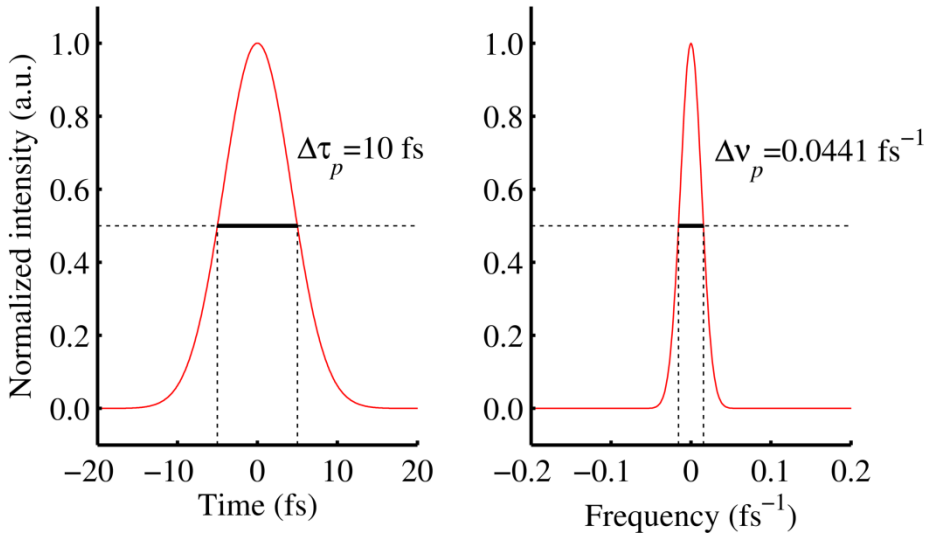

Figure 2.5. The temporal and spectral intensity profiles of a 10 fs Gaussian pulse centred at 0.800 µm. The FWHM bandwidths are shown for each domain. The time-bandwidth product is 0.441.

2.3.2 Dispersion

Dispersion arises from that fact that, excluding a vacuum, the refractive index of a medium is frequency dependent. Different spectral components of the pulse will travel at different velocities, introducing a frequency-dependent phase delay and causing the pulse to change its shape. This can by visualised by imagining a long pulse composed of several smaller wavepackets. The lower frequency side of the spectrum has a frequency $(\omega_0 - y)$ while the higher frequency side has frequency $(\omega_0 + y)$. The

leading edge is therefore moving with a phase velocity which is $2y$ greater than the trailing edge, causing the pulse to broaden.

The spectral phase of a pulse is the phase of the electric field in the frequency domain, and is commonly defined as a Taylor expansion with respect to the mean frequency of the pulse ω_0,

$$\begin{aligned}\varphi(\omega) = \varphi(\omega_0) & \\ &+(\omega-\omega_0)\left(\frac{\partial\varphi}{\partial\omega}\right)\Big|_{\omega_o} \\ &+\frac{1}{2}(\omega-\omega_0)^2\left(\frac{\partial^2\varphi}{\partial\omega^2}\right)\Big|_{\omega_o} \\ &+\frac{1}{6}(\omega-\omega_0)^3\left(\frac{\partial^3\varphi}{\partial\omega^3}\right)\Big|_{\omega_o}+\ldots\end{aligned} \quad . \tag{2.47}$$

The first line in Equation (2.47) is the absolute phase, a measure of the phase accumulated at reference frequency ω_0, and has no effect on the pulse shape. The second line contains a first-order differential term which describes the linear change of the phase with frequency, corresponding to a time-domain delay or group-delay, τ_g. The third line contains a term describing the rate at which the group delay changes with frequency, known as the group-delay dispersion (GDD), with units of fs^2 [9].

The final line in Equation (2.47) contains two important terms in the evolution of the pulse shape. The term $(\omega-\omega_0)^3$ describes cubic spectral phase variations associated with pulse break-up, a result of steepening of one edge of the pulse and stretching of the other. The differential term is known as the third-order dispersion (TOD), with units of fs^3, and is the first of the higher-order spectral phase terms. Figure 2.6 illustrates the effect of dispersion upon an initially transform-limited pulse.

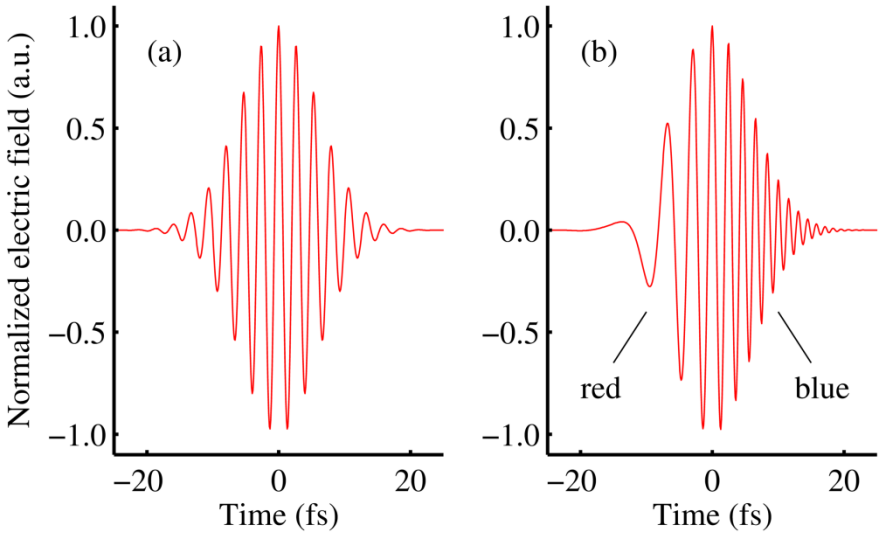

Figure 2.6. A 10 fs pulse (a) before and (b) after positive dispersive broadening. Longer blue frequencies are delayed relative to the shorter red frequencies.

Control of the intracavity dispersion is necessary for the generation of few-cycle ultrashort pulses. The GDD, TOD and higher-order spectral phase terms associated with one cavity roundtrip must be negligible ('flat') across the spectral bandwidth of the pulse. The net dispersion accumulated during one roundtrip can come from a number of sources.

2.3.3 Material dispersion and associated dispersion compensation methods

Material dispersion arises from the wavelength-dependent refractive index, $n(\lambda)$, of all dielectric media, such as optical glasses, gain media and nonlinear crystals [10]. For many materials $n(\lambda)$ is described through the Sellmeier equation, from which the material dispersion can be derived. For practicality, it is useful to describe the GDD in terms of refractive index and wavelength, where

$$GDD = \frac{\lambda^3 L}{2\pi c^2}\frac{d^2 n}{d\lambda^2}. \tag{2.48}$$

A summary of important phase and frequency relationships is given in Table 2.1.

Table 2.1. Higher-order dispersion terms. Adapted from [11].

Dispersion term	Angular frequency notation	Refractive index notation
Phase velocity	$v_p = \frac{\omega}{k_n}$	$\frac{c}{n}$
Group velocity	$v_g = \frac{d\omega}{dk_n}$	$c\left(n - \lambda\frac{dn}{d\lambda}\right)^{-1}$
1st order dispersion (group delay)	$\tau_g = \frac{d\varphi}{d\omega}$	$\frac{L}{c}\left(n - \lambda\frac{dn}{d\lambda}\right)$
2nd order dispersion (GDD)	$\frac{d\tau_g}{d\omega} = \frac{d^2\varphi}{d\omega^2}$	$\frac{\lambda^3 L}{2\pi c^2}\frac{d^2 n}{d\lambda^2}$
3rd order dispersion (TOD)	$\frac{d^3\varphi}{d\omega^3}$	$\frac{-\lambda^4 L}{4\pi^2 c^3}\left(3\frac{d^2 n}{d\lambda^2} + \lambda\frac{d^3 n}{d\lambda^3}\right)$

2.3.3.1 Dispersion compensation through a prism pair

Material dispersion arises from the wavelength-dependent refractive index of dielectric media. In contrast, geometric dispersion arises from the wavelength-dependent path lengths associated with optical systems that exhibit angular dispersion, such as diffraction gratings or prisms. Inserting carefully chosen angularly dispersive elements in a cavity allows for control of the net GDD and TOD, subject to well defined geometric constraints [12].

The use of a prism pair for both intracavity and extracavity dispersion compensation has been well documented, and a full analysis available in many optics textbooks [13–18]. With a prism apex separation l and angular beam deviation β, the negative second- and third-order dispersion due to the optical geometry illustrated in Figure 2.7 is given by

$$\frac{d^2\varphi}{d\omega^2} = \frac{\lambda^3}{2\pi c^2}\frac{d^2 P}{d\lambda^2} \tag{2.49}$$

$$\frac{d^3\varphi}{d\omega^3} = -\frac{\lambda^4}{4\pi^2 c^3}\left(3\frac{d^2 P}{d\lambda^2} + \lambda\frac{d^3 P}{d\lambda^3}\right) \tag{2.50}$$

where

$$\frac{d^2P}{d\lambda^2} = 4l\sin\beta\left(\frac{d^2n}{d\lambda^2} + \left(2n - \frac{1}{n^3}\right)\left(\frac{dn}{d\lambda}\right)^2\right) - 8l\cos\beta\left(\frac{dn}{d\lambda}\right)^2 \tag{2.51}$$

and

$$\frac{d^3P}{d\lambda^3} = l\cos\beta\left(\left(\frac{24}{n^3} - 48n\right)\left(\frac{dn}{d\lambda}\right)^3 - 24\frac{dn}{d\lambda}\frac{d^2n}{d\lambda^2}\right) + l\sin\beta\left(\begin{array}{l}\left(\frac{dn}{d\lambda}\right)^3\left(\frac{12}{n^6} + \frac{12}{n^4} + \frac{8}{n^3} - \frac{16}{n^2} + 32n\right) \\ + \left(24n - \frac{12}{n^3}\right)\frac{dn}{d\lambda}\frac{d^2n}{d\lambda^2} + 4\frac{d^3n}{d\lambda^3}\end{array}\right). \tag{2.52}$$

This negative geometric dispersion, when coupled with the positive material dispersion of the prisms, allows for fine control over the net dispersion experienced by pulses as they propagate through the prism sequence. The net dispersion can be tuned by sliding the prisms along an axis normal to their base, varying the positive material dispersion contribution while keeping the negative geometrical dispersion constant.

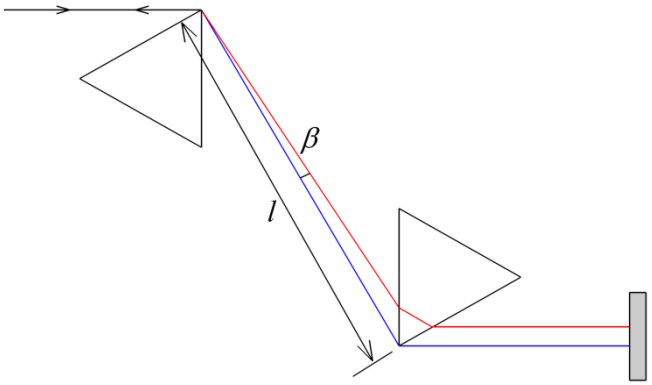

Figure 2.7. Schematic of a double-pass prism compressor for dispersion compensation, with apex separation *l* and beam deviation *β*.

2.3.3.2 Chirped Mirrors

A dielectric mirror coating consists of multiple thin layers of transparent optical materials with alternating low-high refractive indices. Such mirrors routinely offer 99.9% reflectivity over a desired bandwidth and so are commonly used in laser resonators. Mirror dispersion can be understood as arising due to longer wavelengths penetrating deeper into the mirror layers until they are reflected, introducing a wavelength-dependent delay. Alteration of the mirror structure and coating properties allows their use as dispersion control optics, as discussed below.

A progressive increase in the thickness of the alternating mirror layers allows short wavelength components to be reflected before longer wavelength components [19,20], as illustrated in Figure 2.8. The group delay of the reflected beams increases with wavelength, giving negative GDD. Such mirrors are said to be chirped, and for an appropriate design the GDD can be approximately constant over a selected bandwidth, giving zero TOD. In practice chirped mirrors are often implemented in pairs with complementary coatings to remove modulations around the desired GDD value [21–23].

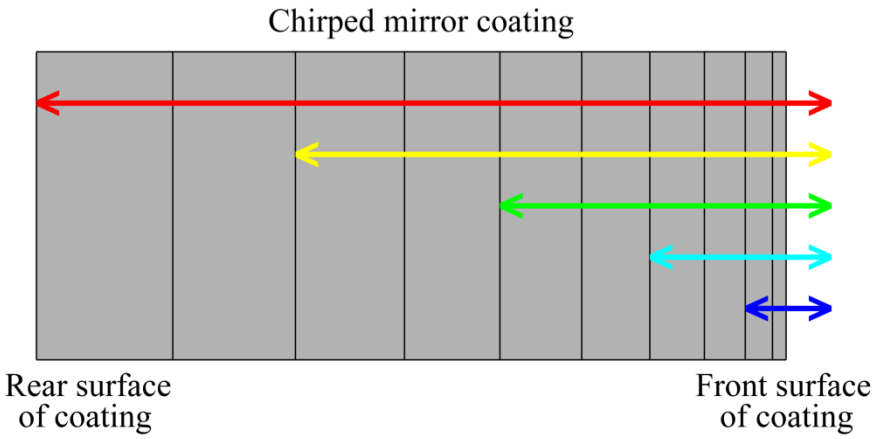

Figure 2.8. Exaggerated diagram of the structure of a chirped mirror.

2.3.3.3 GTI mirrors

A Gires-Tournois interferometer (GTI) mirror contains an etalon-like structure, the monolithic equivalent of an optical standing-wave resonator [24,25]. While they offer substantially greater negative GDD than a chirped mirror pair, the reflectivity bandwidth is narrower.

Both mirror types provide fixed GDD values, and so are often considered for coarse compensation of pulse dispersion. For finer control, a prism pair can be implemented as well.

2.3.4 Self-phase modulation

The optical Kerr effect is caused by a time- and intensity-dependent refractive index variation, described by

$$n = n_0 + n_2 I(t), \tag{2.53}$$

where n_2 is the nonlinear refractive index of the medium [7]. The time-dependent intensity profile of a pulse introduces a higher refractive index change at the peak of the profile, causing the more intense part of the pulse to experience a greater phase shift than weaker pulse components. This broadening of the spectrum is known as self-phase modulation (SPM), and the resulting instantaneous frequency shift is given by

$$\Delta\omega = -\frac{d\phi_{SPM}(t)}{dt} = -\frac{\omega n_2 L}{c}\frac{dI(t)}{dt}. \tag{2.54}$$

SPM is a $\chi^{(3)}$ process similar to TOD, where the leading edge of an intense pulse is red-shifted while the trailing edge is blue-shifted. In the pulse centre there is an approximately linear positive chirp. SPM is commonly used to spectrally broaden pulses, for example in a length of highly nonlinear optical fibre. Subsequent compression of the broadened pulse through dispersion compensation allows for very short pulses to be realised, often shorter than the minimum duration achievable through intracavity dispersion management.

2.4 Pulse characterisation techniques

Section 2.3 described the propagation and evolution of an ultrashort pulse. This section will discuss the complex task of characterising the intensity and phase profiles of such pulses, but first a question must be asked: how you measure an event without comparing it to something shorter? For an ultrashort pulse this is usually not possible, as there is no faster event to compare it against. This section will detail some of the most common techniques that have been developed to measure the temporal properties of ultrashort pulses, including some that also retrieve their phase profile, a critical element in the work carried out in this thesis.

2.4.1 Autocorrelation

Autocorrelation (or more precisely SHG-autocorrelation) was one of the earliest techniques used to measure the temporal features of ultrashort pulses [26]. It does not directly recover any phase information, but is still commonly used to estimate pulse duration due to its simple optical arrangement and data recovery.

2.4.1.1 Interferometric autocorrelation

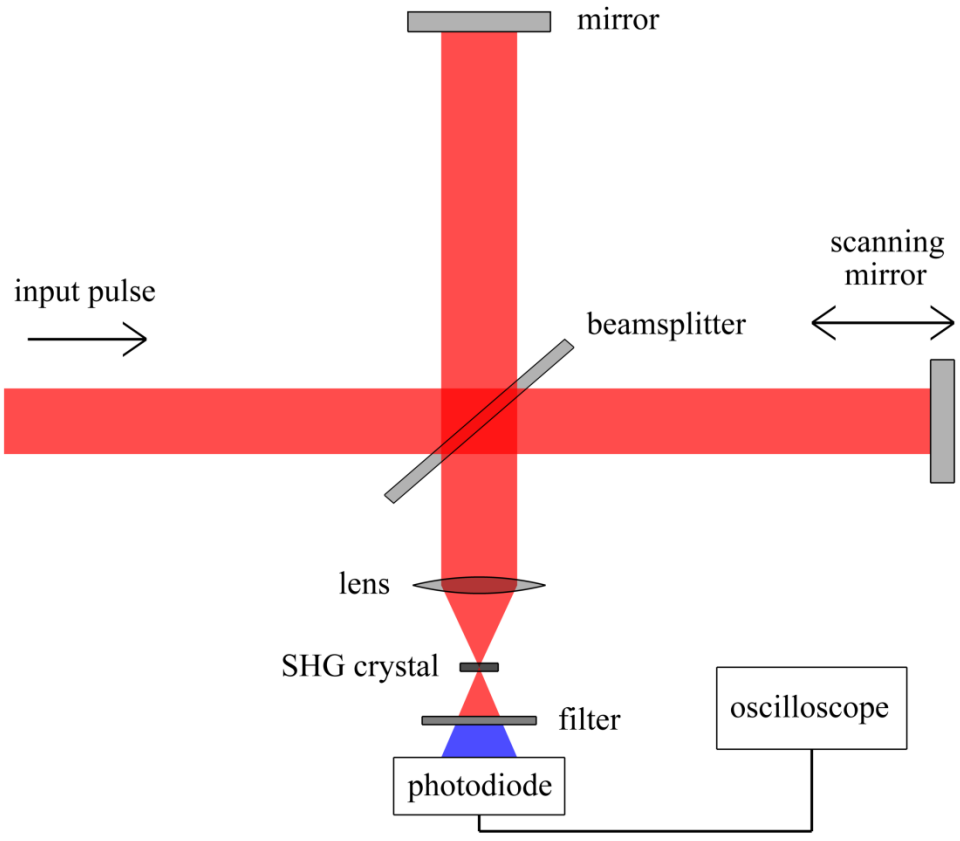

Figure 2.9. A simple optical arrangement for an SHG autocorrelator.

A Michelson interferometer (Figure 2.9) is constructed, where the test pulse is split into two copies which travel down separate arms. A relative delay τ between the two replicas is introduced by moving one arm parallel to the incoming beam, before the pulses are recombined upon exiting the interferometer. The pulses are spatially overlapped within a nonlinear crystal chosen for efficient second harmonic generation. The intensity within the crystal is given by

$$G(\tau) = \int_{-\infty}^{\infty} \left| \left[E(t) + E(t+\tau) \right]^2 \right|^2 dt \,, \tag{2.55}$$

A filter is used to block the fundamental frequencies, and a photomultiplier tube or photodiode is used to detect the second harmonic signal, given by

$$g_2(\tau) = \frac{\int_{-\infty}^{\infty} \left| \left[E(t) + E(t+\tau) \right]^2 \right|^2 dt}{2\int_{-\infty}^{\infty} \left| E(t) \right|^4 dt} \,. \tag{2.56}$$

This signal is a fringe pattern known as an interferometric autocorrelation. For a transform-limited pulse the maximum and minimum value of each fringe half-period is given by a delay change of

$$\Delta\tau = \pm \frac{\pi}{\omega_0} \,. \tag{2.57}$$

This produces upper $\left(g_u\right)$ and lower $\left(g_l\right)$ envelope functions for the fringe pattern, given by

$$g_u(\tau) = \frac{\int_{-\infty}^{\infty} \left| E(t) + E(t-\tau) \right|^4 dt}{2\int_{-\infty}^{\infty} \left| E(t) \right|^4 dt} \tag{2.58}$$

and

$$g_l(\tau) = \frac{\int_{-\infty}^{\infty} \left| E(t) - E(t-\tau) \right|^4 dt}{2\int_{-\infty}^{\infty} \left| E(t) \right|^4 dt} \tag{2.59}$$

respectively. From Equations (2.58) and (2.59) it can be shown that the maximum value corresponds to a delay of zero $(\tau = 0)$ and where the delay becomes increasingly large $(\tau \to \pm\infty)$. Equation (2.56) then becomes

$$g_2(\tau) = \frac{\int |2E(t)|^4 dt}{2\int |E(t)|^4 dt} = \frac{16}{2} = 8 \tag{2.60}$$

in the first instance and

$$g_2(\tau) = \frac{\int |E(t)|^4 dt + \int |E(t)|^4 dt}{2\int |E(t)|^4 dt} = 1 \tag{2.61}$$

in the second instance. An interferometric autocorrelation will therefore have a contrast ratio of 8:1, as shown in Figure 2.10.

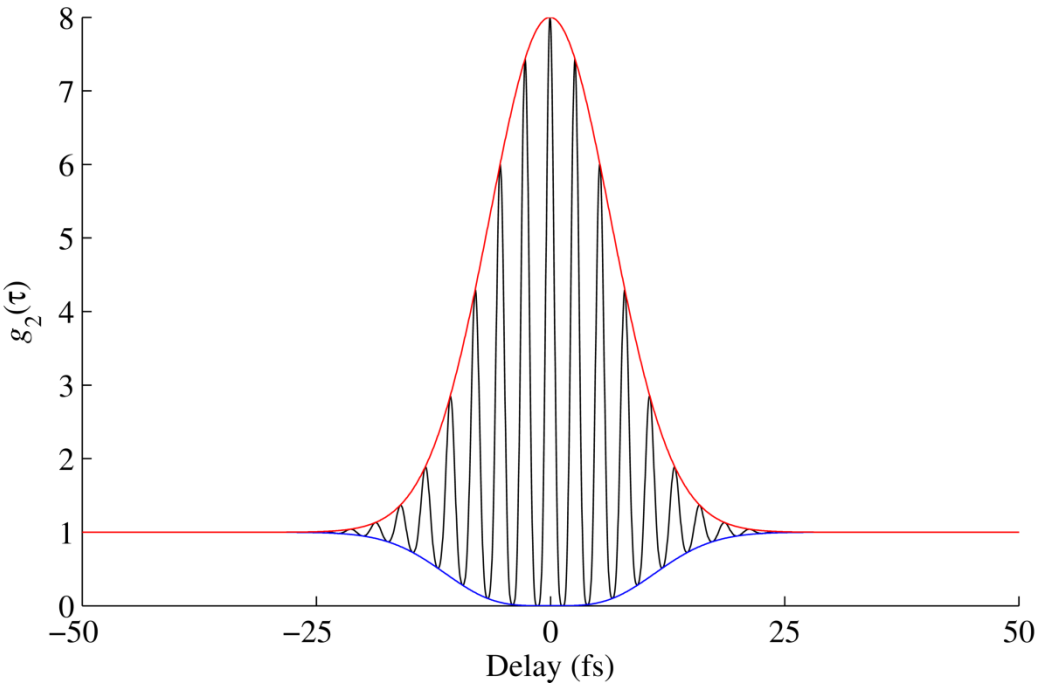

Figure 2.10. An interferometric trace of a transform-limited 10 fs Gaussian pulse centred at 0.800 µm. g_2 is the black line, g_u is the red line and g_l is the blue line.

A chirped pulse will produce a distorted interferogram as the two electric fields will most strongly interfere when the delay is zero. This changes the shape of the upper and lower envelopes, allowing a quick visual inspection to show if the pulse is chirped An interferometric autocorrelation therefore contains some amount of phase information, however as Figure 2.11 demonstrates the trace only records the magnitude of the chirp added to the pulse. The envelope distortions created by the presence of chirp can lead

to experimental errors when determining the pulse duration from the FWHM point [27]. This technique does not recover the absolute phase across the pulse.

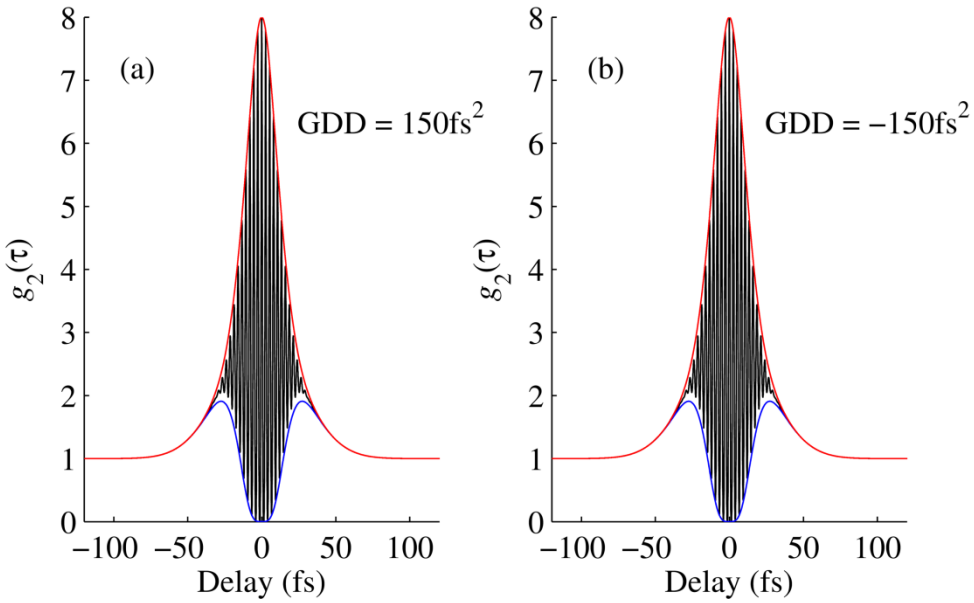

Figure 2.11. Interferometric autocorrelation traces of a chirped 10 fs Gaussian pulse centred at 0.800 µm. The addition of (a) 150fs^2 and (b) -150fs^2 of GDD produces identical traces, and the envelope becomes distorted.

2.4.1.2 Intensity autocorrelation

A complementary technique to interferometric autocorrelation is intensity autocorrelation. When the signal rise time of the detector is slower than the time taken to scan through one fringe then the fringes cannot be resolved and an intensity profile is recorded instead. The intensity autocorrelation trace is time averaged, with all phase information being lost, and is given by

$$g_2(\tau) = 1 + \frac{2\int_{-\infty}^{\infty} I(t)I(t-\tau)dt}{\int_{-\infty}^{\infty} I(t)^2 dt}. \tag{2.62}$$

When $\tau = 0$ and $\tau = \pm\infty$, Equation (2.62) simplifies to 3 and 1 respectively, giving a contrast ratio of 3:1, as shown in Figure 2.12.

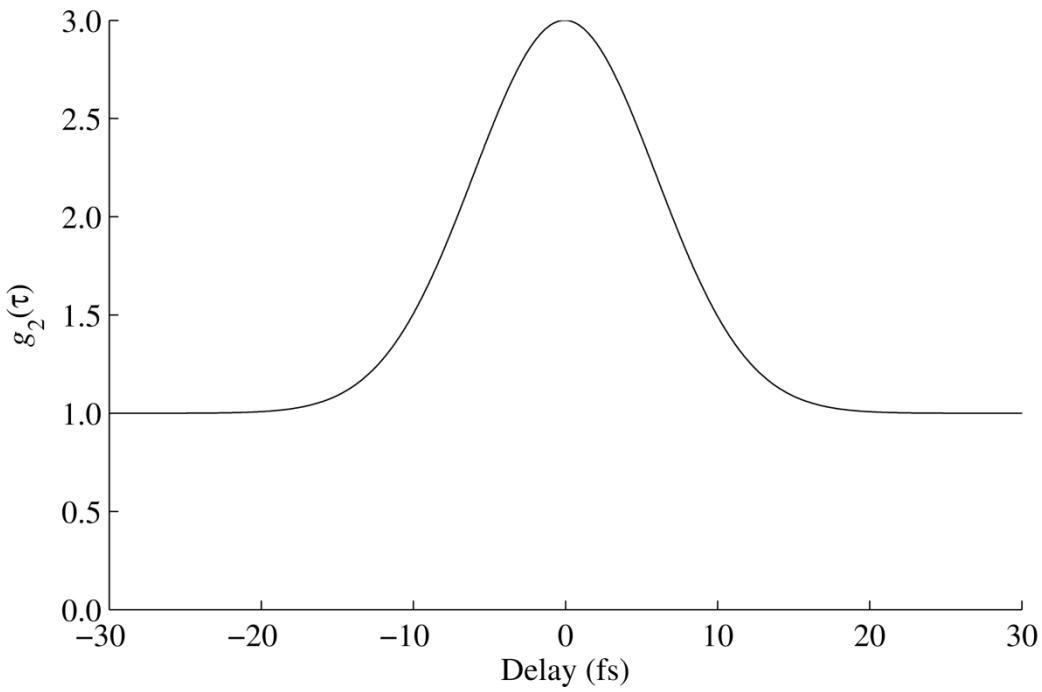

Figure 2.12. Intensity autocorrelation of a 10 fs Gaussian pulse centred at 0.800 µm.

An intensity autocorrelation will broaden in the presence of additional chirp, allowing for accurate pulse duration measurements to be performed when the pulse shape is known with certainty. As with an interferometric autocorrelation, the trace cannot be used to differentiate between positive and negative dispersion, as shown in Figure 2.13.

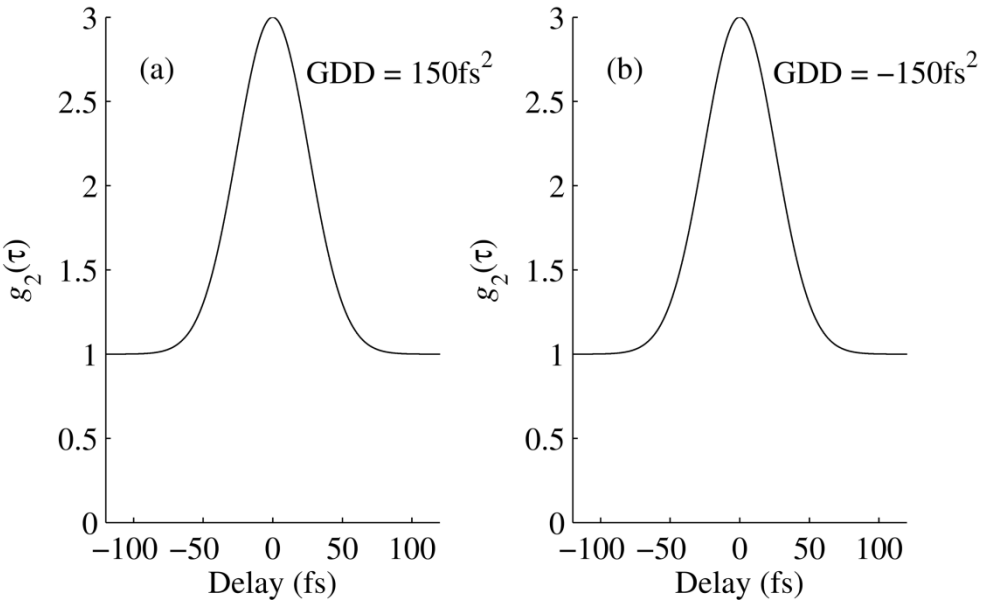

Figure 2.13. Intensity autocorrelation traces of a chirped 10 fs Gaussian pulse centred at 0.800 µm. The addition of (a) 150fs^2 and (b) -150fs^2 of GDD produces identical traces, and the envelope remains undistorted.

2.4.1.3 Pulse duration retrieval

The pulse duration is calculated directly from the resulting autocorrelation trace. The FWHM of the autocorrelation trace Δt is related to the actual pulse duration $\Delta\tau_p$ by a conversion factor k in the expression

$$\Delta\tau_p = \frac{\Delta t}{k}. \tag{2.63}$$

The value of k depends on the pulse shape and the type of autocorrelation used. A summary of the values of k for two ultrashort pulse profiles is given in Table 2.2.

Table 2.2. Interferometric and intensity autocorrelation conversion factors for Gaussian and sech^2 pulses.

	Conversion factor *k*	
Pulse profile	***Intensity***	***Interferometric***
Gaussian	1.414	1.697
sech^2	1.543	1.897

While the autocorrelation technique can retrieve the duration of a pulse, it is ill-suited for recovering information about its phase. As pulse durations become shorter the phase across the spectral profile of the pulse becomes more important. There are several techniques that have been developed over the past few decades that allow the intensity and phase profiles of a pulse to be characterised in a single experiment. Frequency-resolved optical gating (FROG) was the first of these to be introduced, and works in the time-frequency domain. Spectral-phase interferometry for direct electric-field reconstruction (SPIDER) was pioneered several years later, and is an interferometric technique. More recently, and of some importance to the work carried out in this thesis, the multiphoton intrapulse interference phase scan (MIIPS) method was created to both characterise and shape a pulse simultaneously. These techniques and some of their variants will now be discussed.

2.4.2 Frequency-resolved optical gating (FROG) and its variants

Experiments to measure the pulse spectrum or autocorrelation are intrinsically constrained to the frequency or time domain respectively. A FROG measurement takes place in both regimes simultaneously in what is known as the time-frequency domain.

While this may be non-intuitive to visualise, it is most familiar to people in the form of a music score (Figure 2.14). The frequency of the note depends on the vertical position in the scale (y-axis), and the time at which the note is played is given by its horizontal position (x-axis). The volume or intensity of the note is also indicated in the score.

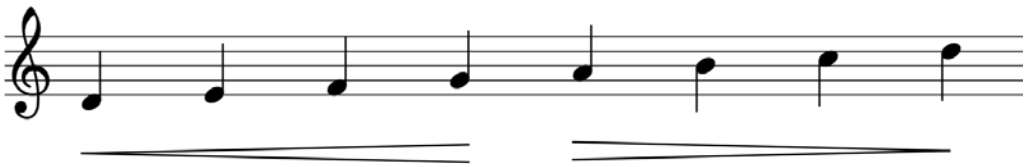

Figure 2.14. A musical score, a common example of the time-frequency domain.

A musical score is an example of a spectrogram, a visual representation of the time-frequency domain where the frequency is recorded as a function of time, thus indicating which frequencies arrive first. This is known as time-gating, and an optical spectrogram is given by

$$SP(\omega,\tau)=\left|\int_{-\infty}^{\infty}E(t)G(t-\tau)e^{-i\omega t}dt\right|^2. \tag{2.64}$$

$G(t-\tau)$ is known as the time gating function, producing a set of spectra of gated sections of $E(t)$ as the delay τ is varied. FROG was the first technique to experimentally record a pulse spectrogram, where the signal beam produced by a frequency-mixing autocorrelation is spectrally resolved [28–31].

2.4.2.1 Second harmonic generation FROG (SHG-FROG)

SHG-FROG is the most commonly used FROG variant, and requires the same optical arrangement as an SHG autocorrelator [30,32–35]. Instead of recording the second harmonic output on a photodiode however, a spectrometer is used to spectrally resolve the output as a function of delay. A prism or grating is often used to spatially separate the second harmonic from the fundamental frequency. Alternatively a non-collinear geometry can be employed so that the generated signal exits the nonlinear crystal at a different angle from the input, however an optical filter would also suffice as material dispersion introduced after the doubling crystal does not affect the measurement. The SHG-FROG trace is mathematically described by

$$I_{FROG}^{SHG}(\omega,\tau)=\left|\int_{-\infty}^{\infty}E(t)E(t-\tau)e^{-i\omega t}dt\right|^2. \tag{2.65}$$

Comparing Equation (2.65) with Equation (2.64), it can be seen that gating pulse is identical to the test pulse, as expected from an SHG process. SHG-FROG traces for a negatively-chirped, chirp-free and positively-chirped pulse are shown in Figure 2.15.

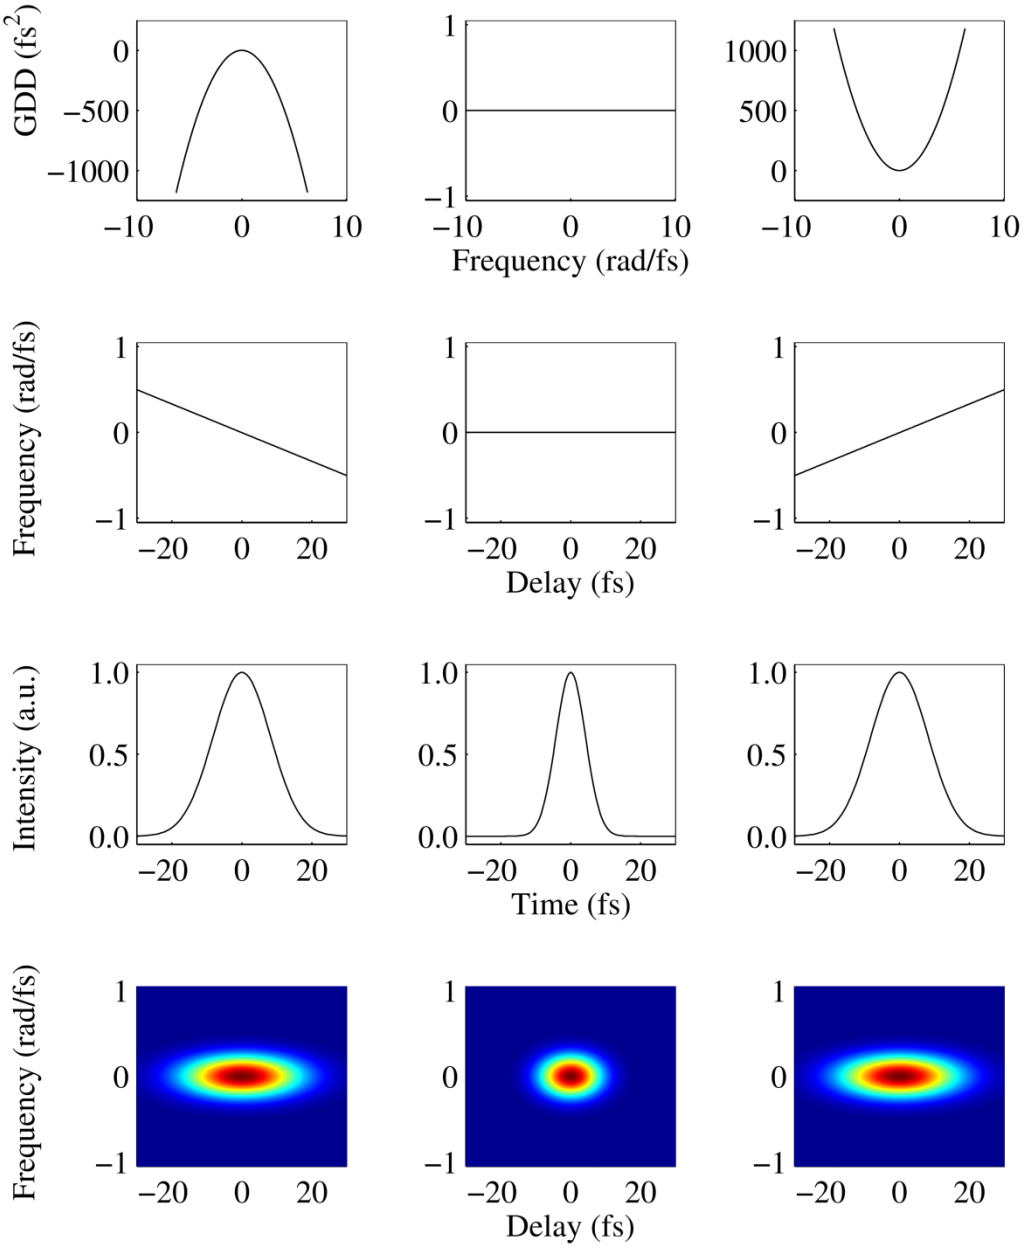

Figure 2.15. SHG-FROG traces for an unchirped pulse (centre column), a negatively chirped pulse (left column) and a positively chirped pulse (right column). In the chirped cases the pulse has visibly broadened in time while the spectrum remains unaltered. A comparison between traces for positive and negative chirp show the time direction ambiguity associated with SHG-FROG.

The SHG-FROG technique has many practical advantages that make it preferable over other FROG variants. The optical arrangement introduces little material dispersion into the beam path and so is suitable for characterising very short pulses; material dispersion introduced after the nonlinear crystal does not alter the measurement as it does not

affect the spectral domain. As SHG is a $\chi^{(2)}$ process the measurement is also very sensitive.

One issue with this technique is the introduction of ambiguities in the time domain. The pulse $E(t)$ and its time reversed conjugate $E^*(-t)$ both produce the same SHG-FROG trace. The trace is therefore symmetrical with respect to delay, and as a result the sign of the higher order phase terms cannot be determined, as illustrated in Figure 2.15.

2.4.2.2 Cross FROG (XFROG)

XFROG differs from the standard FROG technique in that instead of the same pulse acting as the gate and the test, two distinctly different pulses are used [36,37], such as two spectrally separate pulses or a two spectrally identical pulses with different amounts of chirp. These two pulses interact in the nonlinear crystal as either SFG or DFG [38,39], with the resulting spectrum resolved to recover the XFROG trace, as shown in Figure 2.16. This technique allows the characterisation of pulses that cannot be frequency-doubled, such as those in the UV spectral region. With use of the right retrieval algorithm, both the test and gate pulse can be characterised with this technique.

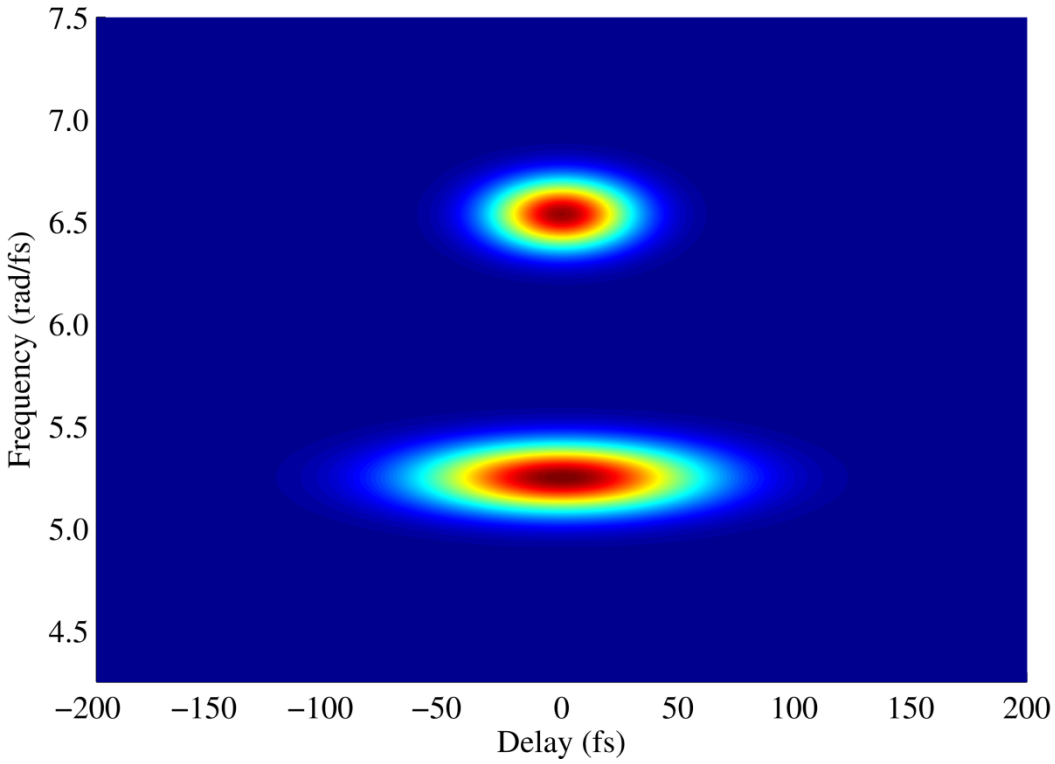

Figure 2.16. A trace from a simulated XFROG experiment operating in SFG mode. A 10 fs pulse centred at 0.800 µm is used to gate a 50 fs pulse centred at 0.480 µm (upper trace) and a 100 fs pulse centred at 0.650 µm. The resulting output wavelengths are 0.288 µm and 0.359 µm respectively.

2.4.2.3 Data retrieval from a FROG trace

While the general shape of a FROG trace will offer some indication of the level of chirp of the pulse(s), a more thorough analysis is necessary for an accurate characterisation. The retrieval of information from a FROG trace is non-trivial, however one algorithm has become the standard technique due to its robustness and fast computational speed. This technique is referred to as Principal Component Generalised Projections (PCGP), and is the method used to analyse the FROG traces acquired throughout this thesis [40,41].

The retrieval algorithm is used to find the form of the complex amplitude of the signal field $E_{sig}(t,\tau)$, which must satisfy two constraints. The intensity constraint specifies that the Fourier transform of $E_{sig}(t,\tau)$ must have an intensity equal to the experimentally measured trace, given by

$$I_{FROG}(\omega,\tau) = \left| \int_{-\infty}^{\infty} E_{sig}(t,\tau) e^{(-i\omega t)} dt \right|^2 . \tag{2.66}$$

The physical constraint is the mathematical form of the signal field produced when two pulses are combined in a nonlinear medium, given by

$$E_{sig}(t,\tau) = E(t)G(t-\tau) \tag{2.67}$$

where $E(t)$ is the test pulse and $G(t-\tau)$ is the gate function. The exact form of the gate function depends on the FROG geometry used. Combining Equations (2.66) and (2.67) provides the following expression that needs to be solved

$$\sqrt{I_{FROG}(\omega,\tau)} = \int_{-\infty}^{\infty} E(t)G(t-\tau)e^{(-i\omega t)} dt . \tag{2.68}$$

The measured trace $I_{FROG}(\omega,\tau)$ is a real quantity and so contains no phase information. The algorithm aims to determine this phase by minimising the difference between the measured FROG trace and the FROG trace calculated from the modelled pulse. Most algorithms follow the outline shown in Figure 2.17.

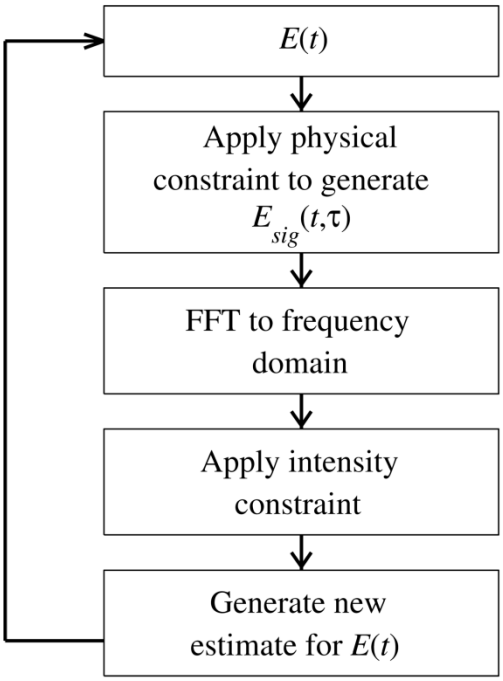

Figure 2.17. The generalised technique for FROG retrieval algorithms

The iterative process works as follows. An initial guess of $E(t)$ is made with reasonable values for $\Delta\omega$ and $\Delta\tau_p$, and the resulting guess value of $E_{sig}(t,\tau)$ is calculated. This function is Fourier transformed to obtain the calculated FROG amplitude $\sqrt{I_{calc}(\omega,\tau)}\phi_{calc}(\omega,\tau)$. The square root of the measured FROG trace $\left(\sqrt{I_{FROG}(\omega,\tau)}\right)$ is then substituted for the calculated intensity. Reversing this process by using the PCGP technique a new estimate for $E(t)$ is obtained from $\sqrt{I_{FROG}(\omega,\tau)}\phi_{calc}(\omega,\tau)$. The algorithm minimises the difference between $I_{calc}(\omega,\tau)$ and $I_{FROG}(\omega,\tau)$ by using identical test and gate pulses for SHG-FROG, and unconstrained test and gate pulses for XFROG.

2.4.2.4 Principal component generalised projections (PCGP)

The PCGP technique makes no special assumption of any connection between the test and gate pulses, generally not the case for many FROG techniques. Retrieval using this method is known as blind-deconvolution, and the blind-FROG trace $I_{FROG}(\omega,\tau)$ is given by Equation (2.65).

At any given time delay τ, the blind-FROG trace is the intensity spectrum resulting from the product of the test pulse and gate pulse. By sampling the intensity spectrum at

different values of t with constant spacing Δt, the gate pulse and test pulse can be written in vector form as

$$E_{test} = \left[E\left(\left(\frac{-N}{2}\right)\Delta t\right), E\left(\left(\frac{-N}{2}-1\right)\Delta t\right), \ldots E\left(\left(\frac{N}{2}-2\right)\Delta t\right), E\left(\left(\frac{N}{2}-1\right)\Delta t\right)\right] \quad (2.69)$$

$$E_{gate} = \left[G\left(\left(\frac{-N}{2}\right)\Delta t\right), G\left(\left(\frac{-N}{2}-1\right)\Delta t\right), \ldots G\left(\left(\frac{N}{2}-2\right)\Delta t\right), G\left(\left(\frac{N}{2}-1\right)\Delta t\right)\right] \quad (2.70)$$

where N is the length of the vector. These equations can be simplified to

$$E_{test} = \left[E_1, E_2 \ldots\ldots E_{N-1}, E_N\right] \quad (2.71)$$

$$E_{gate} = \left[G_1, G_2 \ldots\ldots G_{N-1}, G_N\right]. \quad (2.72)$$

An outer product (O) matrix can be created from these two vectors which consists of all possible products of the gate and test pulses, written as

$$O = \begin{bmatrix} E_1G_1 & E_1G_2 & E_1G_3 & \cdots & E_1G_{N-1} & E_1G_N \\ E_2G_1 & E_2G_2 & E_2G_3 & \cdots & E_2G_{N-1} & E_2G_N \\ E_3G_1 & E_3G_2 & E_3G_3 & \cdots & E_3G_{N-1} & E_3G_N \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ E_{N-1}G_1 & E_{N-1}G_2 & E_{N-1}G_3 & \cdots & E_{N-1}G_{N-1} & E_{N-1}G_N \\ E_NG_1 & E_NG_2 & E_NG_3 & \cdots & E_NG_{N-1} & E_NG_N \end{bmatrix}. \quad (2.73)$$

The FROG trace in the time domain can be formed by simple matrix manipulation. Shifting each row by one place in incremental steps forms a new matrix, where the columns are made up from the gate pulse and test pulse products that have equal temporal delay, written as

$$\begin{bmatrix} E_1G_1 & E_1G_2 & E_1G_3 & \cdots & E_1G_{N-1} & E_1G_N \\ E_2G_2 & E_2G_3 & E_2G_4 & \cdots & E_2G_N & E_2G_1 \\ E_3G_3 & E_3G_4 & E_3G_5 & \cdots & E_3G_1 & E_3G_2 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ E_{N-1}G_{N-1} & E_{N-1}G_N & E_{N-1}G_1 & \cdots & E_{N-1}G_{N-3} & E_{N-1}G_{N-2} \\ E_NG_N & E_NG_1 & E_NG_2 & \cdots & E_NG_{N-2} & E_NG_{N-1} \\ \tau = 0 & \tau = -\Delta t & \tau = -2\Delta t & \cdots & \tau = 2\Delta t & \tau = \Delta t \end{bmatrix}. \quad (2.74)$$

The rows are shifted until the resulting temporal delay varies incrementally from negative to positive, providing a matrix form of $E_{sig}(t,\tau)$. Taking the modulus squared

of the Fourier transform of this matrix produces the generated FROG trace. The magnitude of this trace is then replaced with the square root of the magnitude of the experimental FROG trace. The entire procedure is then reversed to recover a pseudo outer product O_p from which values of the gate and test pulse can be extracted. This can easily be achieved by using matrix-vector products, where the next guesses for the pulses are obtained by multiplying the previous guesses by the pseudo outer product matrix and its transpose. This is the formal PGCP step, and is written as

$$E_{new} = E_{old} O_p O_p^{'} \tag{2.75}$$

and

$$G_{new} = G_{old} O_p O_p^{'} . \tag{2.76}$$

The new values E_{new} and G_{new} are used to create a new outer product matrix and the whole process is repeated until the computed FROG trace matches well with the experimentally retrieved trace. For the case of SHG-FROG, in which the test and gate pulses are identical, an additional constraint can be applied to the guess pulse that reflects this. The entire PCGP algorithm is outlined in Figure 2.18.

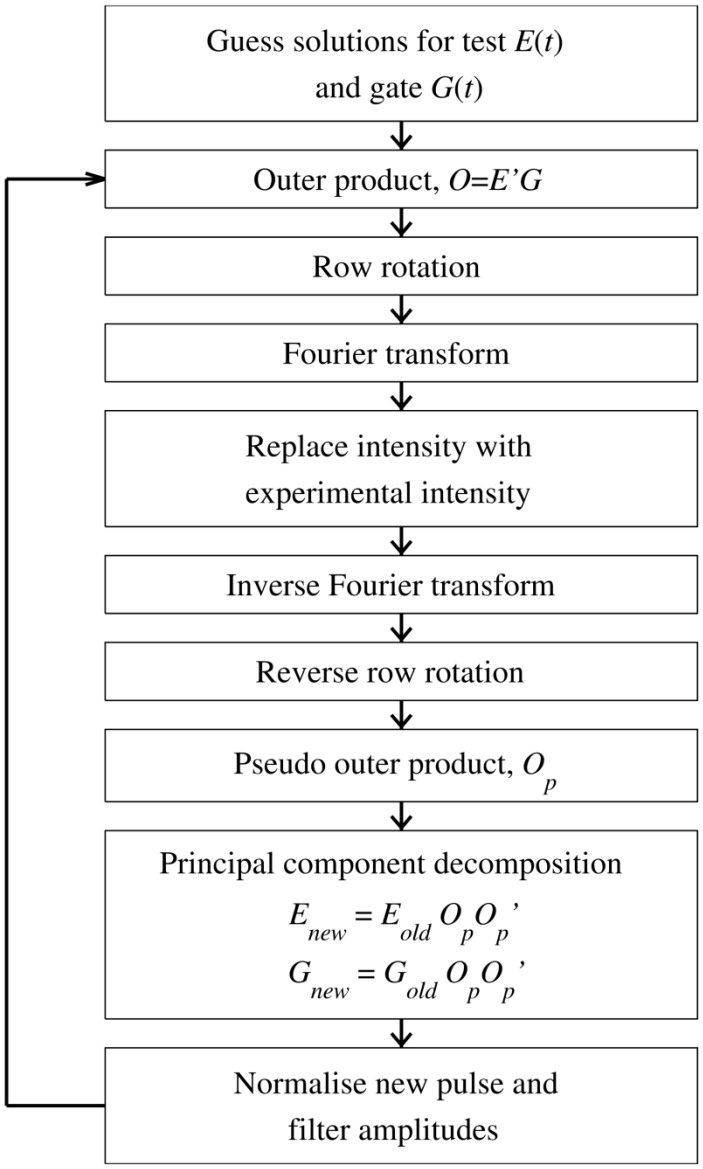

Figure 2.18. Schematic of the PCGP retrieval algorithm

While the FROG technique is extremely useful for determining the duration of a pulse it is less suited for recovery of its spectral phase. This drawback, combined with the lack of the ability to perform a single-shot measurement, led to the development of SPIDER, a complementary interferometric technique that allows for the full characterisation of an ultrashort pulse.

2.4.3 Spectral-phase interferometry for direct electric-field reconstruction (SPIDER)

SPIDER recovers the spectral phase of a pulse through spectral shearing interferometry, where the interference pattern between two pulses separated in time is recorded [42–44]. The pulses are identical in duration and intensity, however they have a different

centre frequency and so are said to be spectrally sheared. The interferogram generated by these two pulses is given by

$$S(\omega) = I(\omega+\Omega) + I(\omega) + 2\sqrt{I(\omega+\Omega)I(\omega)}\cos\left[\phi(\omega+\Omega) - \phi(\omega) + \omega\tau\right], \quad (2.77)$$

where $I(\omega)$ is the pulse spectrum, Ω is the difference between the centre frequencies of the two pulses (spectral shear), $\phi(\omega)$ is the spectral phase of the pulse and τ is the delay between the two pulses.

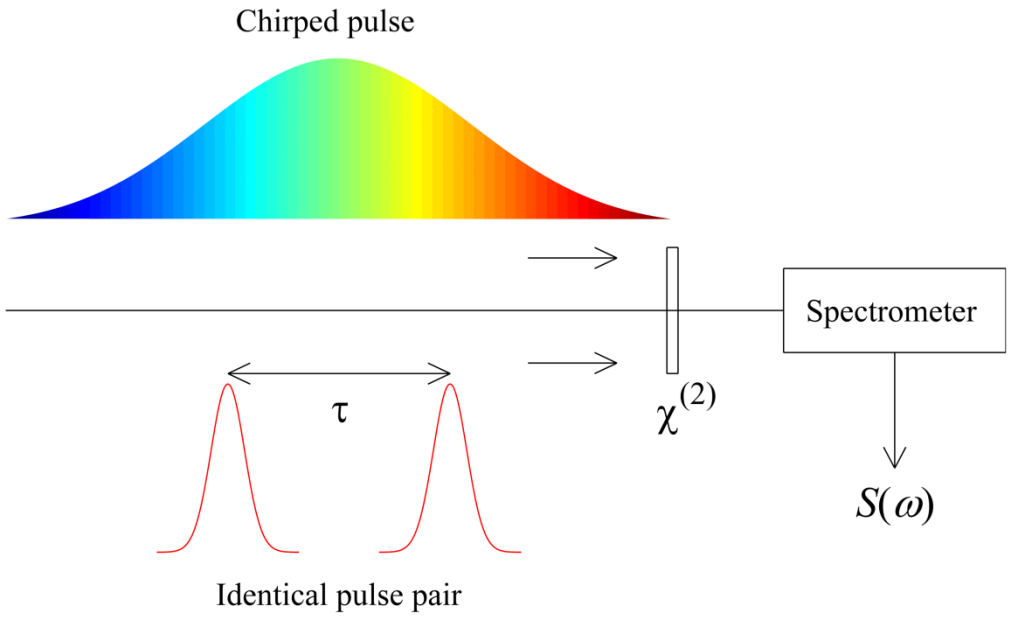

Figure 2.19. Schematic of a SPIDER experimental set-up. A heavily chirped pulse is mixed in a nonlinear crystal with an identical pulse pair separated by delay τ. The delay is sufficiently small (or the chirp sufficiently large) that the identical pulses interact with different frequency components of the chirped pulse.

Experimentally, the spectral shearing of the pulses is obtained by mixing two well-conditioned beams in a nonlinear crystal, as illustrated in Figure 2.19. The first beam consists of a pair of identical pulses separated by a known delay τ. The second beam consists of a highly chirped pulse that must satisfy two conditions. Firstly, the duration of the chirped pulse must be longer than the known delay τ. Secondly, the chirp must be sufficiently large to ensure that when the chirped pulse mixes with the delayed pair of pulses, each pulse mixes with a different frequency within the chirped pulse. When both these conditions are met and the experiment is correctly configured, each mixing pulse produced in the crystal will have a different centre frequency and so will be spectrally sheared. The SPIDER interferogram is obtained by measuring the output from the nonlinear crystal on a spectrometer.

2.4.4 Multi-photon intrapulse interference phase scan (MIIPS)

The newest common technique for pulse measurement combines spectral phase characterisation with pulse shaping in a single set-up to produce transform-limited pulses. While numerous attempts have been made to combine pulse characterisation with shaping [45], MIIPS is the first method that analytically retrieves the phase of the pulse across the whole spectrum and compensates for phase distortions [23]. The process is iterative, with increasing accuracy with subsequent iterations, and can be tailored to suit almost any test pulse. The technique can also distinguish between positive and negative chirp (unlike FROG) and is a single-line measurement with no spatial and temporal overlap requirements and no moving parts (unlike both FROG and SPIDER).

A simple experimental arrangement for a MIIPS measurement is shown in Figure 2.21. The output from an ultrafast laser is directed into a pulse shaper containing a spatial light modulator (SLM) before being frequency doubled in a nonlinear crystal. The SHG spectrum is recorded on a spectrometer with the intensity of the spectrum proportional to the chirp of the laser, as discussed in Sub-section 2.4.4.1. The chirp is retrieved (as detailed in Sub-section 2.4.4.2) and the inverse phase applied to the pulses using the SLM. This process is iterated to produce a near-transform-limited pulse at the plane of the SHG crystal. This procedure is outlined in Figure 2.20.

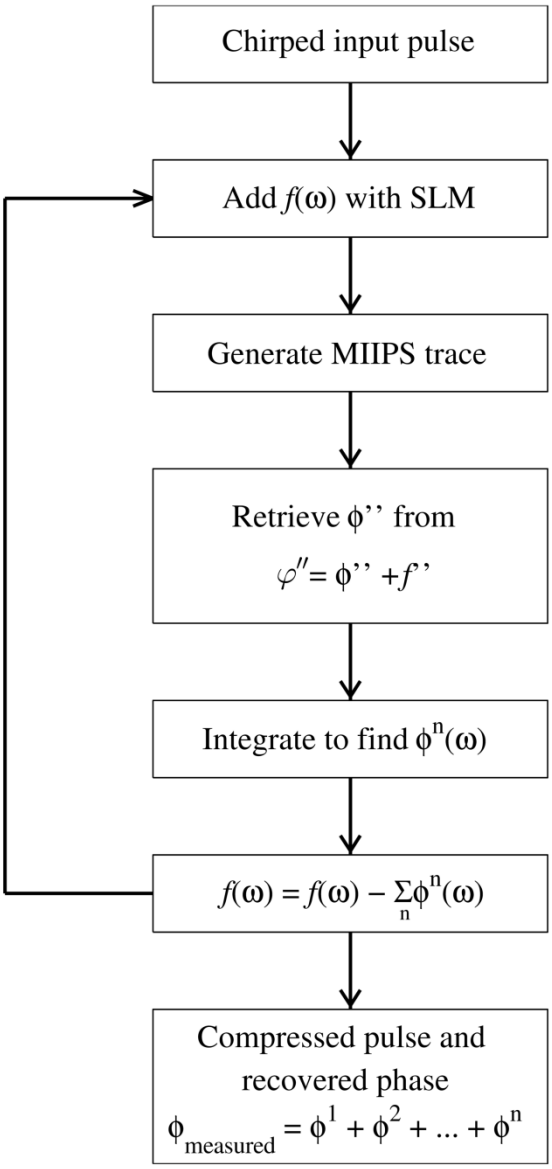

Figure 2.20. Outline of the MIIPS phase retrieval and compensation procedure.

There are several important considerations that must be made when selecting components for a MIIPS measurement. The shaper geometry may be transmissive or reflective; however the SLM must be able to introduce sufficient phase modulation at the desired wavelength, a minimum of 2π. The phase-frequency response of the SLM must be properly characterised in order to produce accurate spectral phase modulation [46,47]. The additional dispersion introduced by the prisms or gratings in a standard pulse shaper configuration may require additional compensation as the upper phase response of an SLM is finite. The phasematching bandwidth of the SHG crystal must be sufficient to allow for conversion of the entire pulse spectrum, and the spectrometer must have sufficient resolution to detect variations in the intensity of the

SHG spectrum. As with a FROG measurement, any optical elements located after the nonlinear crystal do not affect the measurement.

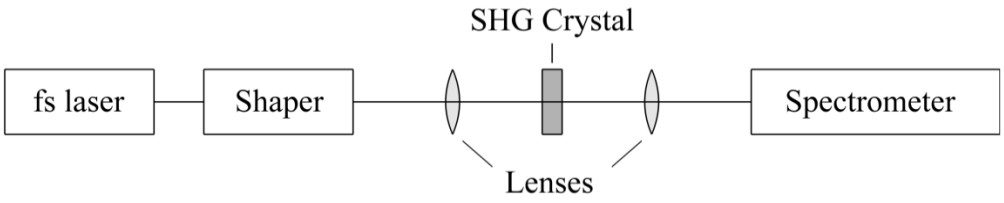

Figure 2.21. Experimental MIIPS setup.

2.4.4.1 The effect of phase on the SHG process

MIIPS exploits the fact that the efficiency of a two-photon process (such as SHG or two-photon absorption) is influenced by the phase relationship between different frequency components within the input pulse. Linear chirp has been shown to reduce the signal intensity of two-photon microscopy [48,49], while quadratic, cubic and higher order chirp has long been understood as a large contributing factor to the efficiency of nonlinear frequency mixing processes [50].

The SHG intensity at frequency 2ω is proportional to an integral over the spectral amplitude $|E(\omega)|$ and spectral phase $\varphi(\omega)$ of the pulse, written as

$$S^{(2)}(2\omega) \propto \left| \int |E(\omega+\Omega)||E(\omega-\Omega)| \, e^{i[\varphi(\omega+\Omega)+\varphi(\omega-\Omega)]} \, d\Omega \right|^2. \tag{2.78}$$

From Equation (2.78) it can be seen that transform-limited pulses (where $\varphi = 0$) generate the maximum SHG intensity. A Taylor expansion of the sum of the spectral phases around ω produces the expression

$$\begin{aligned} \varphi(\omega+\Omega)+\varphi(\omega-\Omega) = 2\varphi(\omega) & \\ +\left(\frac{d^2\varphi}{d\omega^2}\right)\Omega^2 & \\ +\frac{1}{12}\left(\frac{d^4\varphi}{d\omega^4}\right)\Omega^4 + \ldots & \\ +\frac{2}{(2n)!}\left(\frac{d^{2n}\varphi}{d\omega^{2n}}\right)\Omega^{2n}. & \end{aligned} \tag{2.79}$$

To a first approximation and excluding higher-order even terms, the SHG intensity has a local maximum when $\varphi(\omega+\Omega)+\varphi(\omega-\Omega)=0$, which is the case when $\left(\frac{d^2\varphi}{d\omega^2}\right)=\varphi''(\omega)=0$.

2.4.4.2 The formal MIIPS method

Phase retrieval through MIIPS involves the introduction of a reference phase function $f(\omega)$ which locally reduces or cancels phase distortions in the spectral phase of the input pulse [51,52], resulting in a maximum SHG signal with the minimisation of

$$\varphi(\omega)=\phi(\omega)+f(\omega). \tag{2.80}$$

From the Taylor expansion described in Equation (2.79), this maximum occurs when

$$\varphi''(\omega)=\phi''(\omega)+f''(\omega)=0. \tag{2.81}$$

For a well-known reference phase function $f(\omega)$ the second derivative $f''(\omega)$ is also known, allowing both the unknown parameters $\phi''(\omega)$ and $\phi(\omega)$ to be determined. This reference function typically takes the form

$$f(\omega)=\alpha\sin\left(\gamma\omega-\delta\right), \tag{2.82}$$

where $\alpha=1.5\pi$ and $\gamma=\tau_p$, the pulse duration. The phase shift δ is scanned over from 0 to 4π. It should be noted that the MIIPS technique is not limited to periodic functions, however they are useful in constraining the maximum and minimum values of $f(\omega)$. The SHG spectrum is acquired for each step in the phase shift, producing a 2D plot of phase against frequency (or more commonly wavelength). The value of $\delta(\omega)$ for each frequency is found and used to calculate the second derivative of the unknown phase, given by

$$\phi''(\omega)=-f''(\omega)=\alpha\gamma^2\sin\left[\gamma\omega-\delta^n_{\max}(\omega)\right]. \tag{2.83}$$

Equation (2.83) has *n* solutions, and so the search area over which the maximum of the SHG spectrum is found must be predefined. As the goal of a MIIPS measurement is to create transform-limited pulses, the case of $\phi''(\omega)=0$ is considered, providing the solution

$$\delta^{n}_{\max}(\omega) = \gamma(\omega - \omega_0) + n\pi\,, \tag{2.84}$$

where $n = 0, \pm 1, \pm 2\ldots$, $\omega_0 = \dfrac{2\pi c}{\lambda_0}$ (the carrier frequency) and λ_0 is the centre wavelength of the spectrum. Transform-limited pulses therefore produce a MIIPS trace where the SHG signal is greatest on a periodic basis, forming parallel lines separated by π. The search area for $\delta_{\max}(\omega)$ when measuring non-transform-limited pulses is therefore limited to parallel areas separated by π, shown as dashed red lines in Figure 2.22. Simulated MIIPS traces for a transform-limited pulse and a chirped pulse are also shown in this figure.

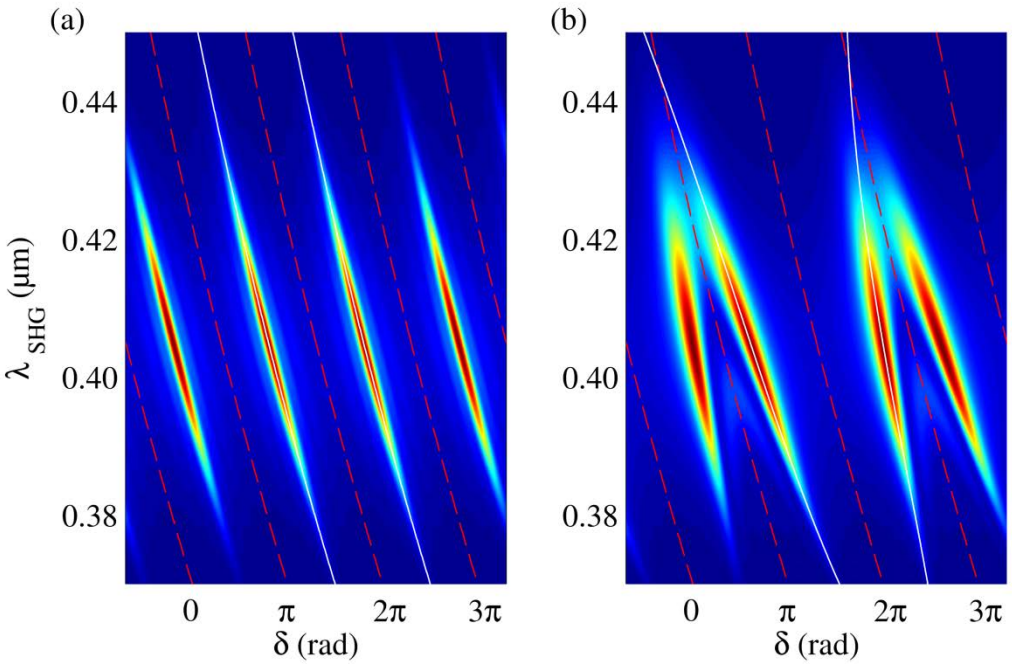

Figure 2.22. Simulated MIIPS traces for a 10 fs pulse centred at 810 nm. (a) A transform-limited pulse. (b) A pulse with GDD = -200 fs^2 and TOD = 1000 fs^3. The dashed red lines are used to define the regions for searching for $\delta_{\max}(\omega)$. Fitted curves of $\delta_{\max}(\omega)$ are shown in white.

Analysis of a MIIPS trace takes place in two neighbouring regions. A mask is created to remove the data from all but one region, leaving a single SHG intensity trace. The phase corresponding to the maximum intensity for each frequency $\left(\delta^{n}_{\max}(\omega)\right)$ is recorded and a polynomial fitted to the data. This polynomial is used to calculate the second derivative of the unknown phase over two regions, written as

$$\phi''(\omega) = \frac{1}{2}\alpha\gamma^2\left(\sin\left[\gamma\omega - \delta^{1}_{\max}(\omega)\right] + \sin\left[\gamma\omega - \delta^{2}_{\max}(\omega)\right]\right). \tag{2.85}$$

The unknown phase $\phi^{I}(\omega)$ is recovered by double integration of $\phi''(\omega)$. This phase distortion is compensated for by the addition of $-\phi^{I}(\omega)$ using the pulse shaper. The process is then repeated and the second iteration phase $\phi^{II}(\omega)$ measured. The correction function written by the shaper now becomes $-\left[\phi^{I}(\omega)+\phi^{II}(\omega)\right]$. Convergence to an accurate result is fast, with increasing absolute accuracy as the magnitude of the phase distortions is reduced [53]. A properly calibrated MIIPS setup can converge to a near-transform-limited pulse within five iterations, with the first iteration giving a good approximation over most of the spectrum. Finally, the unknown phase function is given by the sum of the determined compensation functions, $\phi_{measured}=\phi^{I}+\phi^{II}+...+\phi^{N}$. Note that in almost every case the compensated absolute phase $\phi(\omega)$ will be a ramp, and so it is more instructive to observe the variation of the group delay $\phi'(\omega)$ over the bandwidth of the pulse. An example of the MIIPS phase retrieval and compensation process for a chirped pulse is shown in Figure 2.23.

2.5 Conclusions

This chapter introduced the reader to the fundamentals of nonlinear interactions and pulse dynamics, as well providing a review of the most common techniques for ultrashort pulse characterisation. Phasematching and optical parametric generation will be used in Chapter 3 to construct a synchronously-pumped femtosecond OPO, the outputs of which will be characterised using XFROG. The visible outputs from the OPO will be simultaneously compressed and characterised in Chapter 6.

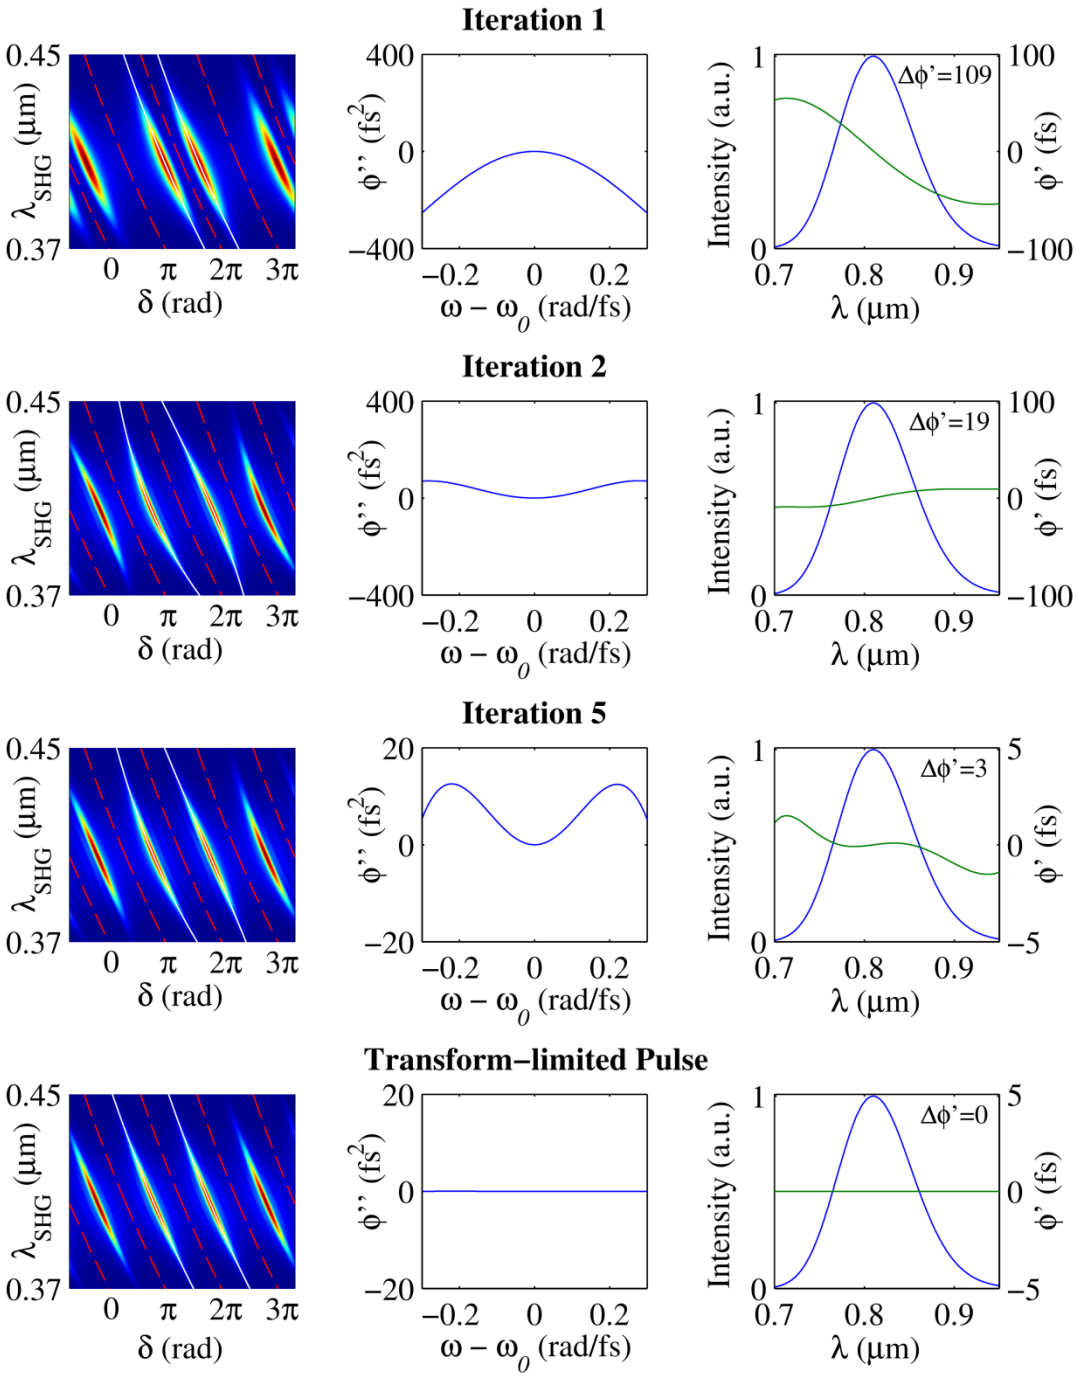

Figure 2.23. Simulation of the MIIPS iterative phase retrieval process. Left panels show the recovered MIIPS traces at each iteration. Centre panels show the retrieved second-order phase. Right panels show the calculated group delay over the wavelength region of the input spectrum. After only one phase retrieval and compensation loop the group delay has dropped by a factor of five, and after five iterations the retrieved group delay varies by less than 3 fs over the spectral bandwidth. The lower row of panels shows the retrieved results for a transform-limited pulse.

2.6 References

1. A. Yariv, *Quantum Electronics*, Third edition (Wiley, 1988).

2. J. A. Armstrong, N. Bloembergen, J. Ducuing, and P. S. Pershan, "Interactions between light waves in a nonlinear dielectric," Phys. Rev. **127**, 1918–1939 (1962).

3. L. E. Myers, "Periodically Poled Materials for Nonlinear Optics," in *Proceedings of SUSSP 52 - Advances in Lasers and Applications*, First edition (Institute of Physics Publishing, 1999), pp. 141–180.

4. G. J. Hall, M. Ebrahimzadeh, A. Robertson, G. P. A. Malcolm, and A. I. Ferguson, "Synchronously pumped optical parametric oscillators using all-solid-state pump lasers," J. Opt. Soc. Am. B **10**, 2168–79 (1993).

5. J. Jiang and T. Hasama, "Harmonic repetition-rate femtosecond optical parametric oscillator," Appl. Phys. B Lasers Opt. **74**, 313–317 (2002).

6. O. Kokabee, A. Esteban-Martin, and M. Ebrahim-Zadeh, "Extended-cavity, tunable, GHz-repetition-rate femtosecond optical parametric oscillator pumped at 76 MHz," Opt. Express **17**, 15635–15640 (2009).

7. A. E. Siegman, *Lasers*, First edition (University Science Books, 1986).

8. H. A. Haus, J. G. Fujimoto, and E. P. Ippen, "Structures for additive pulse mode locking," J. Opt. Soc. Am. B **8**, 2068–76 (1991).

9. I. Walmsley, L. Waxer, and C. Dorrer, "The role of dispersion in ultrafast optics," Rev. Sci. Instrum. **72**, 1–29 (2001).

10. E. Hecht, *Optics (International Edition)*, Fourth edition (Pearson, 2002).

11. U. Keller, "Ultrashort Pulse Generation," in *Proceedings of SUSSP 52 - Advances in Lasers and Applications*, First edition (Institute of Physics Publishing, 1999), pp. 83–115.

12. B. E. Lemoff and C. P. Barty, "Cubic-phase-free dispersion compensation in solid-state ultrashort-pulse lasers," Opt. Lett. **18**, 57–9 (1993).

13. R. L. Fork, O. E. Martinez, and J. P. Gordon, "Negative dispersion using pairs of prisms," Opt. Lett. **9**, 150–2 (1984).

14. O. E. Martinez, J. P. Gordon, and R. L. Fork, "Negative group-velocity dispersion using refraction," J. Opt. A Pure Appl. Opt. **1**, 1003–06 (1984).

15. R. Zhang, D. Pang, J. Sun, Q. Wang, S. Zhang, and G. Wen, "Analytical expressions of group-delay dispersion and cubic phase for four-prism sequence used at other than Brewster's angle," Opt. Laser Technol. **31**, 373–9 (1999).

16. S. Yang, K. Lee, Z. Xu, X. Zhang, and X. Xu, "An accurate method to calculate the negative dispersion generated by prism pairs," Opt. Lasers Eng. **36**, 381–7 (2001).

17. F. J. Duarte, "Generalized multiple-prism dispersion theory for laser pulse compression: higher order phase derivatives," Appl. Phys. B **96**, 809–814 (2009).

18. R. E. Sherriff, "Analytic expressions for group-delay dispersion and cubic dispersion in arbitrary prism sequences," J. Opt. Soc. Am. B **15**, 1224–1230 (1998).

19. R. Szipocs, K. Ferencz, C. Spielmann, and F. Krausz, "Chirped multilayer coatings for broadband dispersion control in femtosecond lasers," Opt. Lett. **19**, 201–3 (1994).

20. M. Yamashita, M. Ishikawa, K. Torizuka, and T. Sato, "Femtosecond-pulse laser chirp compensated by cavity-mirror dispersion," Opt. Lett. **11**, 504–6 (1986).

21. F. X. Kärtner, N. Matuschek, T. Schibli, U. Keller, H. A. Haus, C. Heine, R. Morf, V. Scheuer, M. Tilsch, and T. Tschudi, "Design and fabrication of double-chirped mirrors," Opt. Lett. **22**, 831–3 (1997).

22. N. Matuschek, F. X. Kartner, and U. Keller, "Analytical design of double-chirped mirrors with custom-tailored dispersion characteristics," IEEE J. Quantum Electron. **35**, 129–137 (1999).

23. G. Steinmeyer, "Femtosecond dispersion compensation with multilayer coatings: toward the optical octave," Appl. Opt. **45**, 1484–1490 (2006).

24. F. Gires and P. Tournois, "Interferometre utilisable pour la compression d'impulsions lumineuses modulees en frequence (An interferometer useful for the compression of a frequency modulated light pulse)," C. R. Acad. Sci. Paris **258**, 6112–15 (1964).

25. R. Szipöcs, A. Köhazi-Kis, S. Lako, P. Apai, A. P. Kovacs, G. DeBell, I. Mott, A. W. Louderback, A. V Tikhonravov, and M. K. Trubetskov, "Negative dispersion mirrors for dispersion control in femtosecond lasers : chirped dielectric mirrors and multi-cavity Gires – Tournois interferometers," Appl. Phys. B **70**, S51–S57 (2000).

26. J. A. Armstrong, "Measurement of picosecond laser pulse widths," Appl. Phys. Lett. **10**, 16–18 (1967).

27. A. M. Weiner, "Effect of group velocity mismatch on the measurement of ultrashort optical pulses via second harmonic generation," IEEE J. Quantum Electron. **19**, 1276–1283 (1983).

28. R. Trebino and D. J. Kane, "Using phase retrieval to measure the intensity and phase of ultrashort pulses: frequency-resolved optical gating," J. Opt. Soc. Am. A **10**, 1101–11 (1993).

29. D. J. Kane and R. Trebino, "Single-shot measurement of the intensity and phase of an arbitrary ultrashort pulse by using frequency-resolved optical gating," Opt. Lett. **18**, 823–5 (1993).

30. D. J. Kane and R. Trebino, "Characterization of arbitrary femtosecond pulses using frequency-resolved optical gating," IEEE J. Quantum Electron. **29**, 571–9 (1993).

31. R. Trebino, K. W. DeLong, D. N. Fittinghoff, J. N. Sweetser, M. A. Krumbügel, B. A. Richman, and D. J. Kane, "Measuring ultrashort laser pulses in the time-frequency domain using frequency-resolved optical gating," Rev. Sci. Instrum. **68**, 3277–95 (1997).

32. G. Taft, A. Rundquist, M. M. Murnane, H. C. Kapteyn, K. W. Delong, R. Trebino, and I. P. Christov, "Ultrashort optical waveform measurements using frequency-resolved optical gating," Opt. Lett. **20**, 743–5 (1995).

33. K. W. DeLong, R. Trebino, J. Hunter, and W. E. White, "Frequency-resolved optical gating with the use of second-harmonic generation," J. Opt. Soc. Am. B **11**, 2206–2215 (1994).

34. A. Baltuska, M. S. Pshenichnikov, and D. A. Wiersma, "Second-harmonic generation frequency-resolved optical gating in the single-cycle regime," IEEE J. Quantum Electron. **35**, 459–478 (1999).

35. Z. E. Penman, T. Schittkowski, W. Sleat, D. T. Reid, and W. Sibbett, "Experimental comparison of conventional pulse characterisation techniques and second-harmonic-generation frequency-resolved optical gating," Opt. Commun. **155**, 297–300 (1998).

36. S. Linden, H. Giessen, and J. Kuhl, "XFROG - A new method for amplitude and phase characterization of weak ultrashort pulses," Phys. Status Solidi B - Basic Res. **206**, 119–124 (1998).

37. D. T. Reid, P. Loza-Alvarez, C. T. Brown, T. Beddard, and W. Sibbett, "Amplitude and phase measurement of mid-infrared femtosecond pulses by using cross-correlation frequency-resolved optical gating," Opt. Lett. **25**, 1478–80 (2000).

38. K. W. DeLong, R. Trebino, and W. E. White, "Simultaneous recovery of two ultrashort laser pulses from a single spectrogram," J. Opt. Soc. Am. B **12**, 2463–66 (1995).

39. S. Linden, J. Kuhl, and H. Giessen, "Amplitude and phase characterization of weak blue ultrashort pulses by downconversion," Opt. Lett. **24**, 569–71 (1999).

40. D. J. Kane, "Real-time measurement of ultrashort laser pulses using principal component generalized projections," IEEE J. Sel. Top. Quantum Electron. **4**, 278–284 (1998).

41. D. J. Kane, G. Rodriguez, A. J. Taylor, and T. S. Clement, "Simultaneous measurement of two ultrashort laser pulses from a single spectrogram in a single shot," J. Opt. Soc. Am. B **14**, 935–43 (1997).

42. C. Iaconis and I. A. Walmsley, "Spectral phase interferometry for direct electric-field reconstruction of ultrashort optical pulses," Opt. Lett. **23**, 792–4 (1998).

43. C. Iaconis and I. A. Walmsley, "Self-referencing spectral interferometry for measuring ultrashort optical pulses," IEEE J. Quantum Electron. **35**, 501–509 (1999).

44. L. Gallmann, G. Steinmeyer, D. H. Sutter, T. Rupp, C. Iaconis, I. A. Walmsley, and U. Keller, "Spatially resolved amplitude and phase characterization of femtosecond optical pulses," Opt. Lett. **26**, 96–8 (2001).

45. T. Binhammer, E. Rittweger, R. Ell, F. X. Kärtner, S. Member, and U. Morgner, "Prism-based pulse shaper for octave spanning spectra," IEEE J. Quantum Electron. **41**, 1552–57 (2005).

46. A. M. Weiner, "Ultrafast optical pulse shaping: A tutorial review," Opt. Commun. **284**, 3669–3692 (2011).

47. J. W. Wilson, P. Schlup, and R. A. Bartels, "Ultrafast phase and amplitude pulse shaping with a single, one-dimensional, high-resolution phase mask," Opt. Express **15**, 8979–87 (2007).

48. B. Broers, H. B. v. Linden v. d. Heuvell, and L. D. Noordam, "Large interference effects of small chirp observed in two-photon absorption," Opt. Commun. **91**, 57–61 (1992).

49. I. Pastirk, J. M. Dela Cruz, K. A. Walowicz, and V. V Lozovoy, "Selective two-photon microscopy with shaped femtosecond pulses," Opt. Express **11**, 1695–1701 (2003).

50. M. Hacker, R. Netz, M. Roth, G. Stobrawa, T. Feurer, and R. Sauerbrey, "Frequency doubling of phase-modulated, ultrashort laser pulses," Appl. Phys. B **277**, 273–7 (2001).

51. V. V Lozovoy, I. Pastirk, and M. Dantus, "Multiphoton intrapulse interference. IV. Ultrashort laser pulse spectral phase characterization and compensation," Opt. Lett. **29**, 775–7 (2004).

52. V. V Lozovoy, I. Pastirk, K. A. Walowicz, and M. Dantus, "Multiphoton intrapulse interference. II. Control of two- and three-photon laser induced fluorescence with shaped pulses," J. Chem. Phys. **118**, 3187–3196 (2003).

53. B. Xu, J. M. Gunn, J. M. Dela Cruz, V. V. Lozovoy, and M. Dantus, "Quantitative investigation of the multiphoton intrapulse interference phase scan method for simultaneous phase measurement and compensation of femtosecond laser pulses," J. Opt. Soc. Am. B **23**, 750–9 (2006).

Chapter 3 - Design and characterisation of a Ti:sapphire-pumped optical parametric oscillator

3.1 Introduction

The cornerstone of the few-cycle pulse synthesis project was the synchronously-pumped optical parametric oscillator (OPO), a resonant cavity capable of producing femtosecond pulses at multiple wavelengths across the visible and infrared spectral regions. This chapter briefly discusses the properties of the Ti:sapphire pump laser and its suitability for ultrashort pulse generation. The design, construction and characterisation of two OPO cavities for the generation of visible light pulses are detailed, and a comparison of the two cavities is considered.

3.2 The Ti:sapphire laser

The mode-locked Ti:sapphire laser became synonymous with ultrashort pulse generation after the discovery in 1991 of self-mode-locking through the optical Kerr effect [1–3], yielding a 60 fs pulse. A decade of intense research followed, reducing the pulse duration directly from the laser oscillator to less than 6 fs through careful control of the intracavity dispersion, either through the use of prisms and double-chirped mirrors [4] or through the addition of semiconductor saturable-absorber mirrors (SESAMs) [5]. Commercially available Ti:sapphire lasers are now routinely capable of producing mode-locked bandwidths that support 10 fs pulses, however external management of the pulse duration through dispersion compensation is non-trivial. This section will outline the theory of mode-locked lasers, and discuss Kerr-lens mode-locking and its application towards femtosecond pulse generation. A detailed analysis of these concepts is available in many standard laser textbooks [6,7].

3.2.1 Ultrashort pulse generation

In an inhomogeneously broadened laser, the longitudinal modes that will oscillate are those that experience a greater net gain than loss during a cavity round trip. These modes have no fixed phase relationship, producing a continuous-wave (CW) or noise-like output. When the modes are fixed in their relative phase the laser is said to be mode-locked, and interference between these resonant cavity modes produces a pulsed output, as illustrated in Figure 3.1. The repetition frequency of the pulses is determined by the cavity length, given by

$$f_{REP} = \frac{c}{2L_{eff}} \tag{3.1}$$

where L_{eff} is the effective cavity length, taking into account different refractive indices within the cavity.

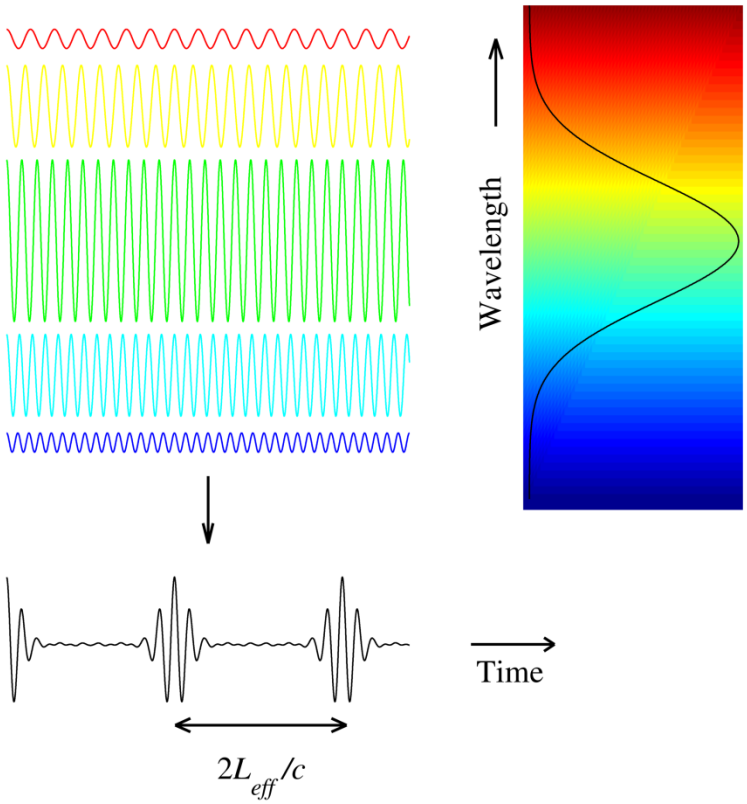

Figure 3.1. Illustration of the concept of mode-locking. Several longitudinal modes that are in phase can interfere to form a pulsed output [8].

Under normal operating conditions a laser will produce a CW output; a phase or amplitude modulation must be applied within the cavity to achieve mode-locking. This modulation can either be introduced through an external source (active mode-locking) or through use of a nonlinear intracavity element (passive mode-locking). Alternatively the laser can be mode-locked the use of additional components (self-mode-locking).

3.2.1.1 Active mode-locking

In active mode-locking a phase or amplitude modulation is applied to the laser, with a modulation frequency equal to a multiple of the cavity repetition rate, given by Equation (3.1). A pulse that makes one round trip of the cavity at this frequency will encounter the modulating device when its gain or transmission is at a maximum,

ensuring that only multiples of this fundamental mode can propagate within the cavity without significant loss. Examples include amplitude or frequency modulation with an intracavity acousto-optic or electro-optic modulator [9], and synchronous pumping with an external pump laser [10]. Due to the frequency bandwidth limitations of the active device, the mode-locked output is typically narrow in spectrum and long in pulse duration.

3.2.1.2 Passive mode-locking

Passive mode-locking relies on nonlinear effects to create modulations in the laser output, most commonly through use of a saturable absorber. A saturable absorber is a material that introduces intracavity loss at low light intensities but becomes transparent at high intensities. Types of saturable absorber include organic dyes and SESAMs [11,12], and more recently carbon nanotubes and graphene [13,14].

3.2.1.3 Self-mode-locking

In Chapter 2 the optical Kerr effect was introduced, a phenomenon that arises when an intense electric field passes through an isotropic dielectric medium. Equation 2.53 showed that a laser beam of sufficient intensity causes a change in the refractive index, and is given here again for clarity

$$n = n_0 + n_2 I(t). \tag{3.2}$$

The phase velocity $\left(v = \frac{c}{n} \right)$ of the pulse is directly proportional to the pulse intensity, so the Kerr effect causes the outer wavefronts of the pulse to have different speeds than those nearer the central, more intense part of the pulse. This creates a virtual lens within the medium leading to self-focusing, which is illustrated in Figure 3.2.

The optical Kerr effect can be exploited to produce a self-focusing within the gain medium of a laser, and was first demonstrated in Ti:sapphire [1]. Introducing a large noise spike to a laser cavity causes self-focusing, which can be exploited to favour pulsed operation by modifying the cavity so that the noise spike experiences less loss the resonant CW mode. Cavity modification takes the form of an aperture; a hard aperture physically constrains the mode size to favour pulsed operation, while a soft aperture alters the cavity alignment so that pulsed operation experiences greater overlap with the pump source than CW operation. Both cases lead to larger gain for pulsed

operation, causing the laser to remain mode-locked. This technique is often called Kerr-lens mode-locking (KLM).

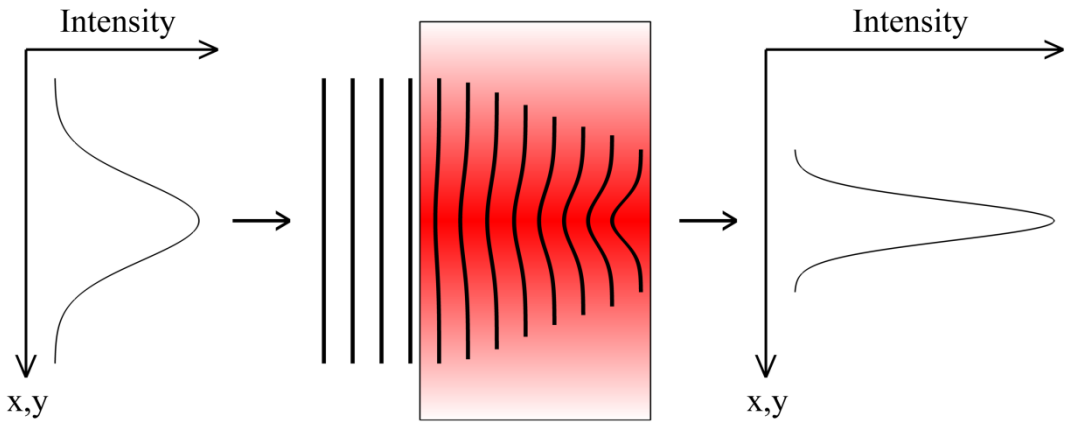

Figure 3.2. Graphical representation of self-focusing.

KLM has the added benefit of generating self-phase modulation in the gain medium (see Subsection 2.3.4), spectrally broadening the pulse and allowing for shorter pulse durations. Ti:sapphire has the largest gain bandwidth of any laser gain material and can therefore support the shortest pulses, provided the intracavity dispersion is carefully managed.

3.2.2 The KM Labs Griffin I pump laser

The Ti:sapphire ($Ti:Al_2O_3$) laser used throughout this project was a custom-built Griffin I model supplied by KM Labs. The pump source for the Ti:sapphire laser was a Verdi V10 manufactured by Coherent, a diode-pumped, frequency-doubled $Nd^{3+}:YVO_4$ laser producing up to 10 W of single-mode CW power at 0.532 µm.

A schematic of the Ti:sapphire laser is shown in Figure 3.3. The cavity was designed around a 4.75 mm Titanium-doped sapphire crystal and a pair of dispersion-compensating fused silica prisms. Each prism can be controlled by a stepper-motor allowing for fine tuning of the intra cavity dispersion. The end mirror was affixed to a pair of piezoelectric actuators which were used for locking of the carrier-envelope offset frequency, as detailed in Section 4.2. The output coupler in mirror M6 was 20% transmitting at 0.820 µm.

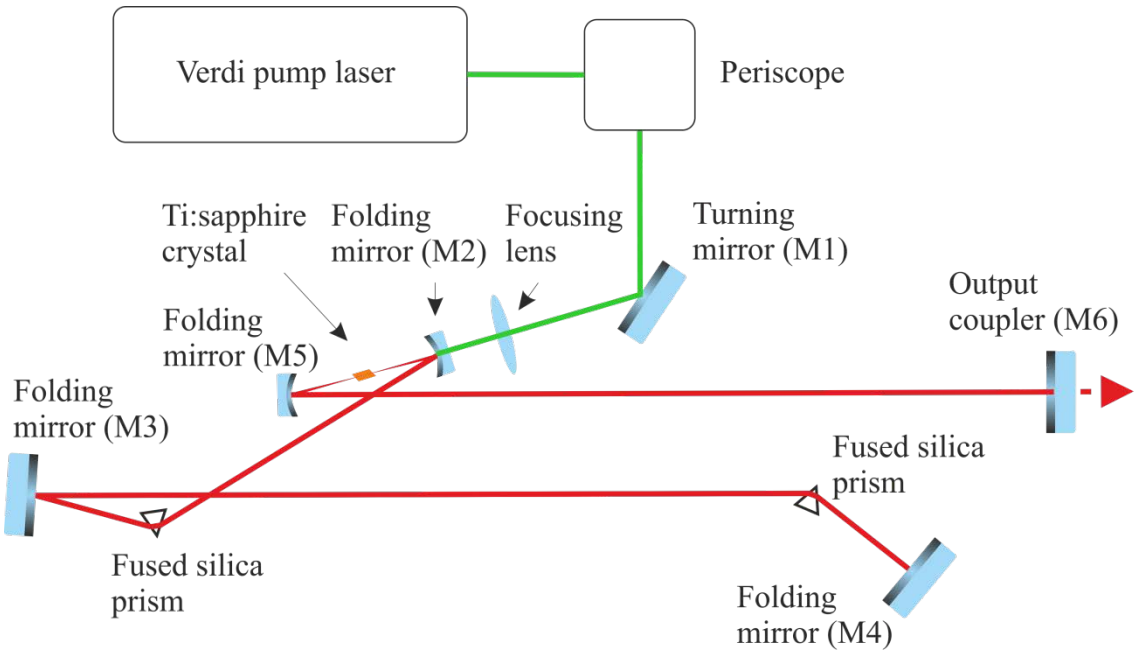

Figure 3.3. Schematic of the Ti:sapphire pump laser.

The design of the Ti:sapphire laser allowed for the beam exiting the cavity to be parallel to the holes on the optical bench. For reasons of stability the two steerable end mirror mounts (M4 and M6) were replaced with sturdier Newport mounts. For added thermal control the breadboard to which the Ti:sapphire optomechanics were fixed was temperature controlled. It was very important that the steering provided by the periscope assembly is fixed and stable, as the beam pointing of the Verdi pump laser affected the performance of the Ti:sapphire oscillator. For this reason, as well as one of environmental control, the periscope assembly was boxed. The entire assembly of Verdi pump, periscope and Ti:sapphire oscillator were encased in a second box to further remove environmental fluctuations.

3.2.3 Characterisation

The Ti:sapphire laser was pumped with 8.8 W from the Verdi laser. The average CW power was 1.5 W, with 1.4 W mode-locked output. A typical mode-locked spectrum is shown in Figure 3.4. The full-width half-maximum (FWHM) bandwidth was 0.042 µm with a centre wavelength of 0.802 µm. From Section 2.3.1, this bandwidth supports pulse durations of 16.1 (22.5) fs for sech^2 (Gaussian) pulse shapes. The pulse repetition frequency was measured as 98.7 MHz, equating to an effective cavity length of 1.52 m.

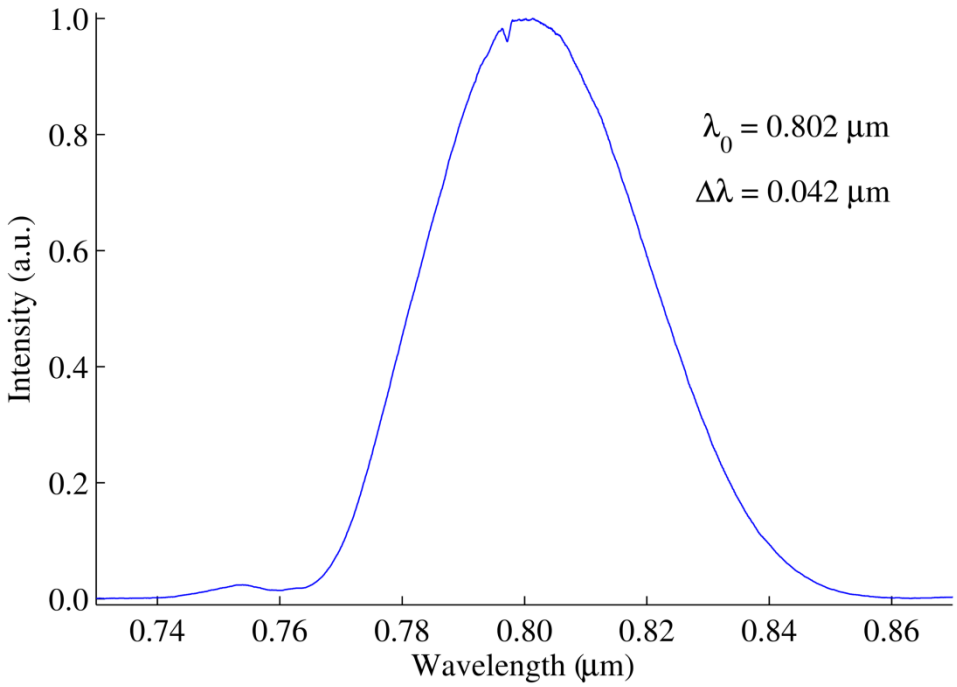

Figure 3.4. Typical output spectrum from the Griffin I Ti:sapphire pump laser. The small notch in the spectrum was due to a defect in the spectrometer.

The centre wavelength, FWHM bandwidth and output power of the laser could be tuned by altering the insertion length of the dispersion-compensating prism pair. The position of each prism could be altered independently, allowing for fine control of the pump characteristics. A 'map' of laser performance as a function of prism position is shown in Figure 3.5. The broadest bandwidths (i.e. those capable of supporting the shortest pulse durations) were achieved by inserting more glass into the beam path with prism 1 and removing glass with prism 2. Figure 3.5 also shows however that an increase in stable mode-locked bandwidth (where no CW breakthrough occurs) resulted in a decrease in average power. This was a balance that must be maintained in order to generate sub-25 fs pulses with a reasonably high average power, and is probably the main reason that commercial lasers are not available with these specifications. Over a period of many months the laser performance would alter and the prism map would drift in the direction indicated by the white arrows in Figure 3.5. This was primarily caused by a change in the pointing of the Verdi pump laser, a consequence of using a periscope to steer the beam into the Ti:sapphire laser box. A complete realignment of the Ti:sapphire laser cavity was therefore necessary on occasion to bring the laser back up to specification.

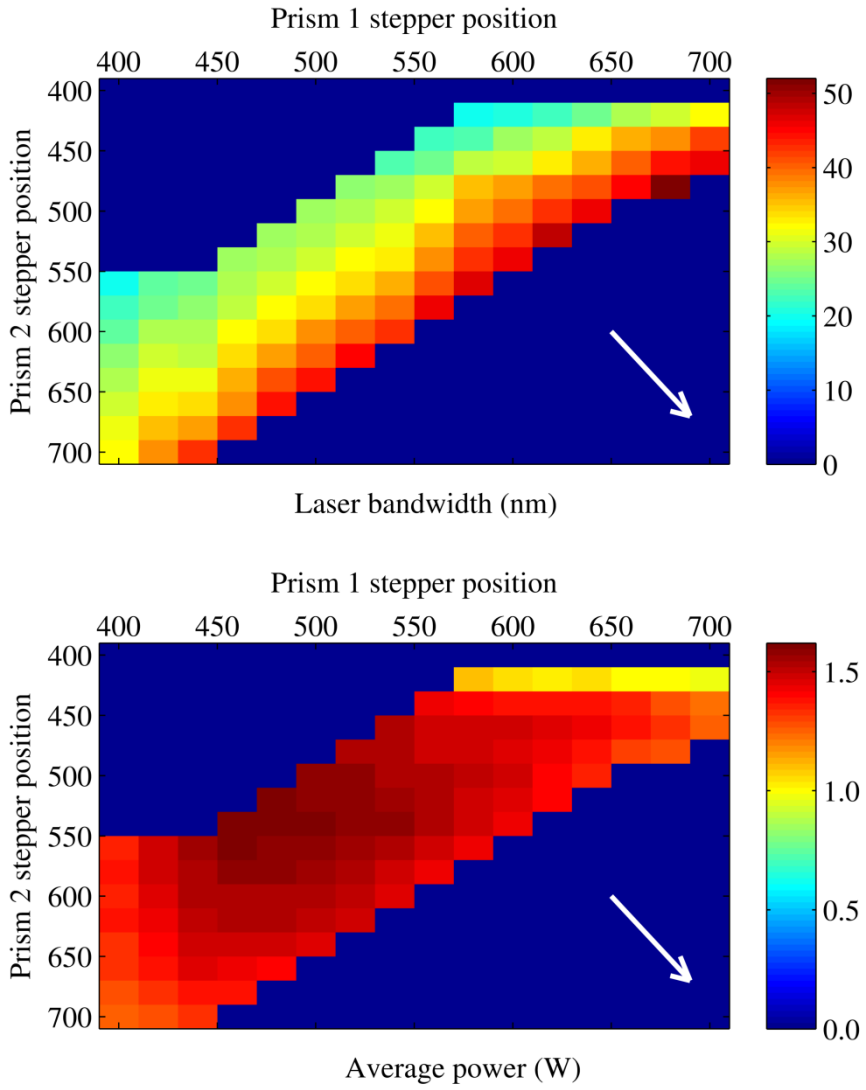

Figure 3.5. 'Maps' of the Ti:sapphire laser performance as a function of prism position. The upper figure shows an increase in laser bandwidth as more of prism 1 is inserted into the beam path and prism 2 is withdrawn. The lower figure shows a decrease in average power for the same prism movements. Blue areas show where no data were recorded, which was selected as the point where CW breakthrough occurred. No resolution was available for the stepper motors; the stepper position is provided by the computer interface of the prisms.

As the pump pulses were output coupled from the cavity they acquired a small amount of chirp from the 6.3 mm of substrate they passed through. In order to measure the pulse duration a small portion of the pump power was reflected from the front surface of a microscope slide and steered into a Timewarp autocorrelator [15]. The resultant autocorrelation is shown in Figure 3.7.

As discussed in Chapter 2, Sub-section 2.4.1, the pulse duration can be recovered from an autocorrelation by dividing the FWHM of the trace by a factor that is dependent on the assumed pulse shape. The FWHM of the trace in Figure 3.7 is 65 fs; assuming a sech^2 pulse shape, this corresponds to a calculated pulse duration of 34 fs. The

time-bandwidth-limited pulse duration can be calculated from the FWHM of the spectral intensity using the equation

$$\Delta\tau_p = \frac{k\lambda_0^2}{c\Delta\lambda} \tag{3.3}$$

where k is the time-bandwidth product for a given pulse shape, λ_0 is the centre wavelength of the spectrum and $\Delta\lambda$ is the FWHM bandwidth. Using 0.315 as the time-bandwidth product for a sech^2 pulse, the minimum pulse duration over this bandwidth is 16.5 fs (see inset of Figure 3.7).

The broadening of a Gaussian pulse of known duration due to propagation over distance z through a material with second-order dispersion $\frac{d^2\phi}{d\omega^2}$ can be calculated by [16]

$$\frac{\tau_p(z)}{\tau_p(0)} = \sqrt{1+\left(\frac{4\ln 2\left(\frac{d^2\phi}{d\omega^2}\right)}{\tau_p^2(0)}\right)}. \tag{3.4}$$

This expression can be used to calculate the additional second-order dispersion that was added to the pulse as it passed through the output coupler of the oscillator. The standard thickness of a 1-inch laser optic is 6.3 mm; the two most common substrates for visible-near-IR transmission are fused silica and BK7. Using these parameters the additional chirp can be calculated as 220 fs^2, with BK7 providing the largest broadening factor of 1.54. The second-order dispersion calculated above is only valid for certain assumptions, primarily that the pulse shape is Gaussian. While an autocorrelation trace cannot be used to directly recover the chirp of a pulse, with some work it can be used to find the chirp to a good approximation. The method is outlined in Figure 3.6. A synthetic pulse is generated in the time domain, again with an assumed pulse shape, with a duration that matches that calculated using the spectral intensity. The synthetic pulse is then Fourier transformed into the frequency domain where its spectrum is replaced with the measured spectrum. The pulse is inverse Fourier transformed to the time domain and a simulated autocorrelation carried out. The synthetic autocorrelation is then compared with the measured autocorrelation. By adding additional second-order phase (and third-order phase if necessary) to the synthetic pulse in the frequency domain, the shape of the resulting autocorrelation can be varied. This process is iterated until the synthetic trace is in good agreement with the measured trace. Viewing

overlapping interferometric traces can be difficult, and so the upper and lower bounds of the synthetic autocorrelation can be used for clarity. Some ambiguity exists however in where the synthetic envelope should sit relative to the interferometric trace.

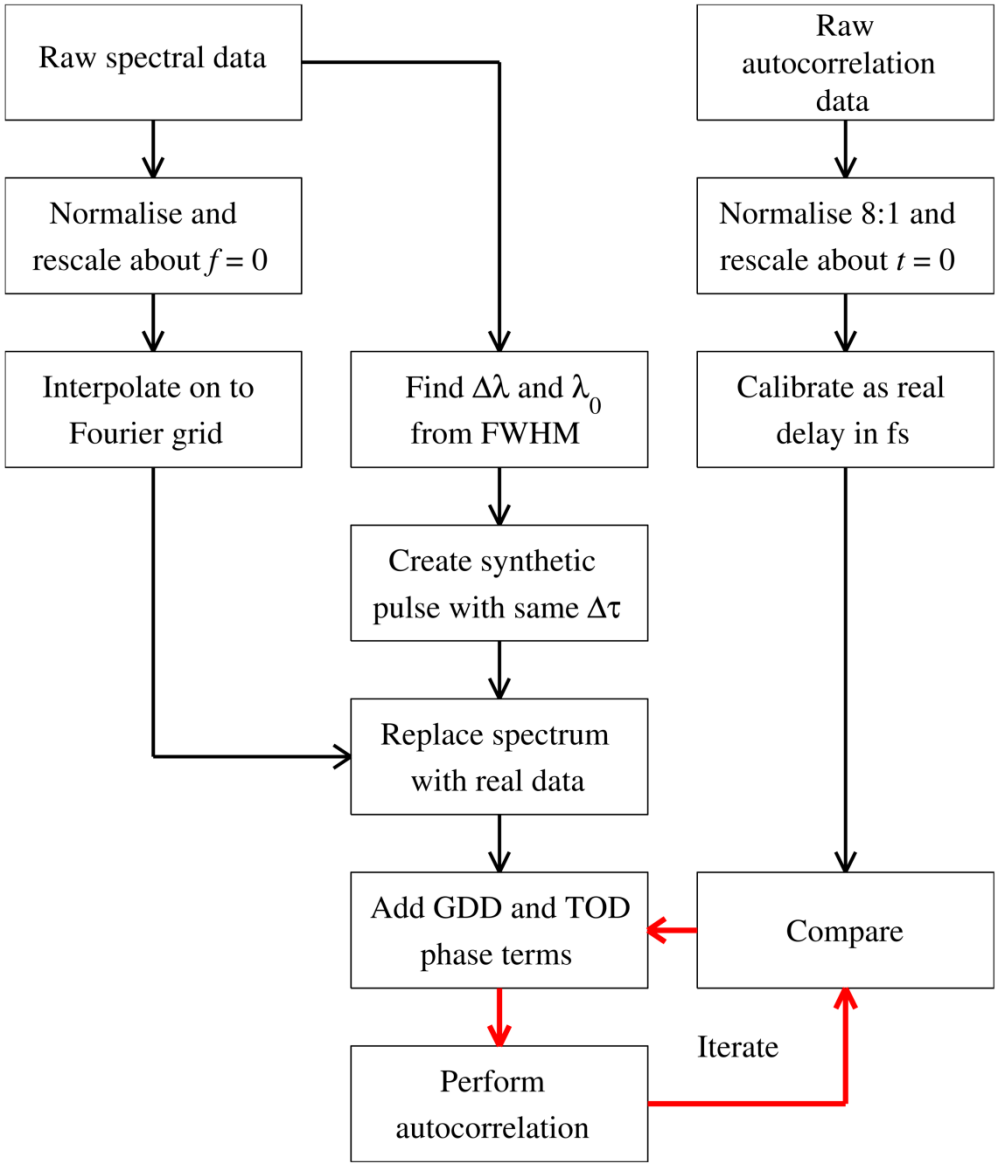

Figure 3.6. Procedure for phase estimation from an interferometric autocorrelation trace.

Using this second method autocorrelation trace envelopes were calculated for 450 fs^2, and are shown as red lines in Figure 3.7. Both methods are inexact, however the second method was preferable as it allowed for the pulse shape to be varied and for additional phase terms to be added to the simulated pulse. It also accounted for spectra that are plainly neither $sech^2$ nor Gaussian in profile.

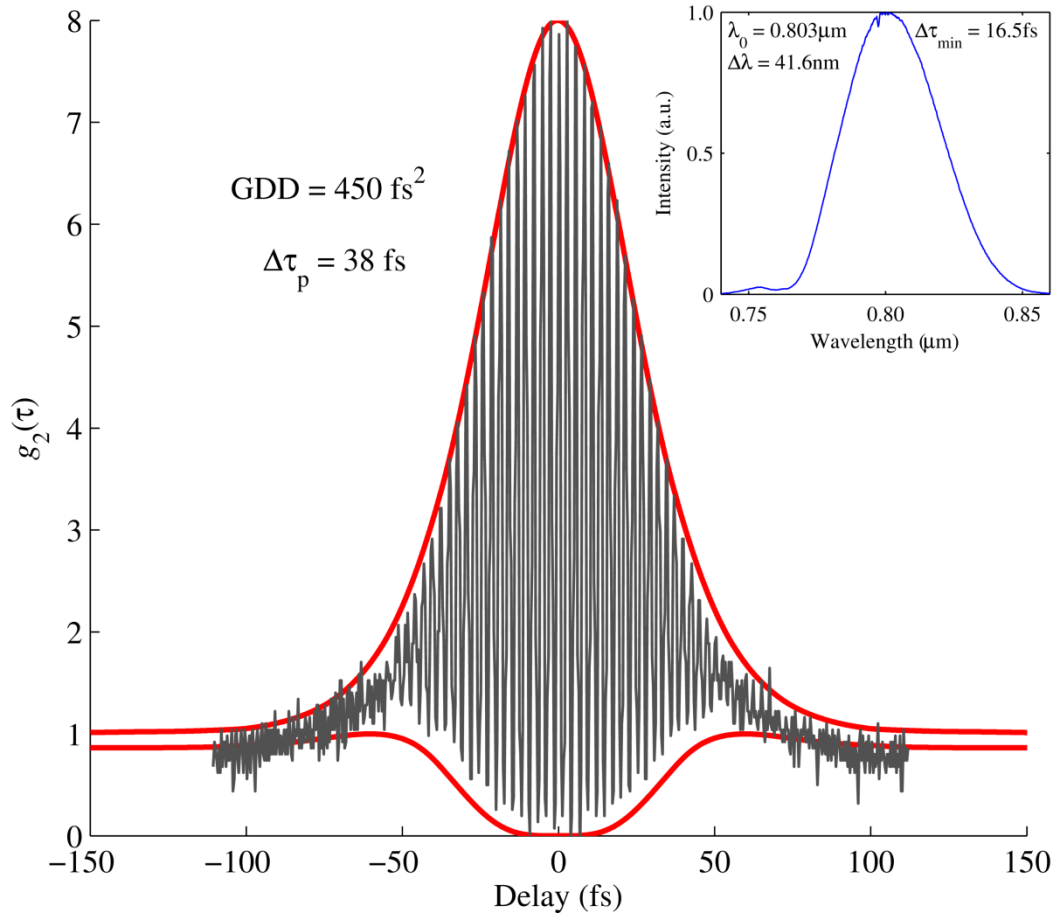

Figure 3.7. Interferometric autocorrelation trace of the Ti:sapphire pump pulses directly after the oscillator. The trace is shown as a black line and the inset shows the corresponding spectrum. The red line indicates the autocorrelation envelopes for a pulse chirped with 450 fs^2 of positive dispersion.

External compression of the chirped pump pulses was achieved through use of a chirped mirror pair provided by Layertec. Each pair of bounces added -150 fs^2 of group delay dispersion over the region of 0.7 – 0.9 µm. The number of bounce pairs was increased and the autocorrelation recorded until a minimum of second-order dispersion was observed. Three pairs of bounces providing -450 fs^2 of GDD were necessary to achieve the shortest pulse duration; the corresponding autocorrelation is shown in Figure 3.8. This is in good agreement with the positive GDD value calculated from the autocorrelation trace.

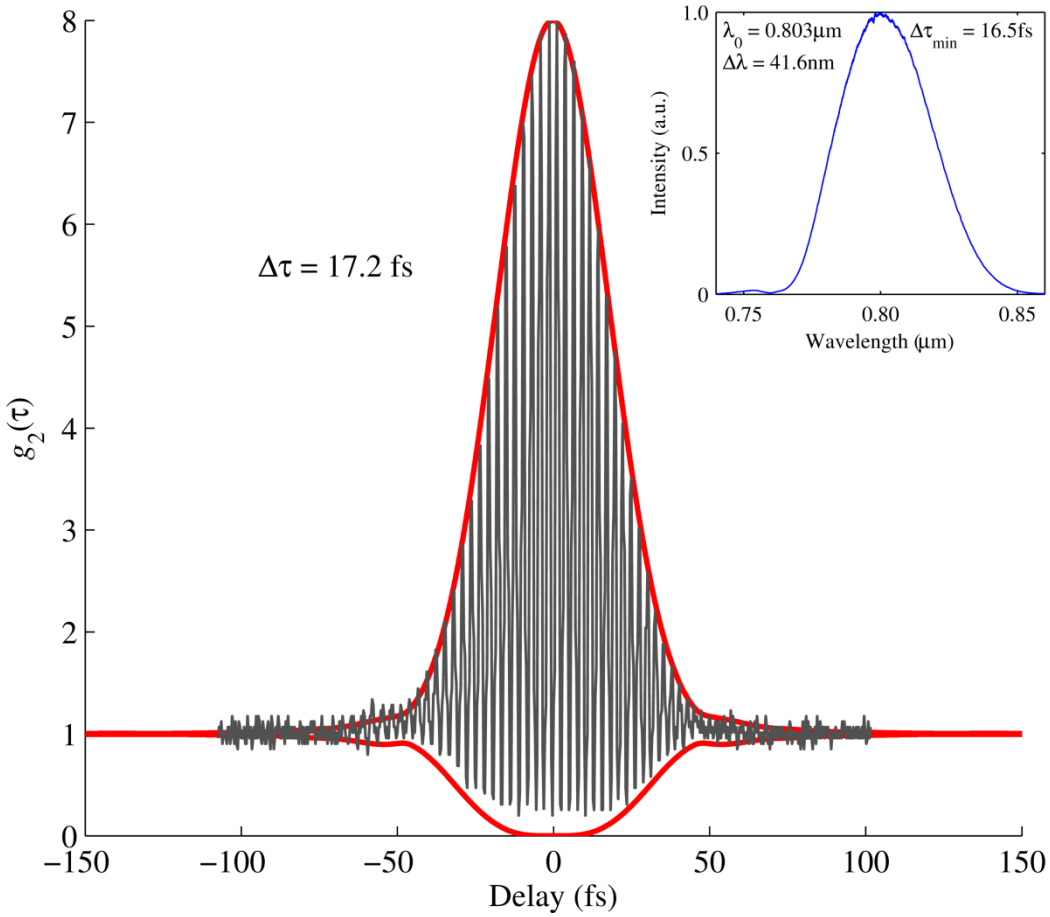

Figure 3.8. Autocorrelation trace after 3 pairs of chirped mirror bounces. The pulse is almost transform-limited.

3.3 First generation PPKTP optical parametric oscillator

Once the characterization of the Ti:sapphire pump laser was complete, construction began on a synchronously-pumped femtosecond OPO for broadband pulse generation in the visible region. The OPO was constructed in two main configurations, a linear cavity and a ring cavity. This section will discuss the design and characterisation of the first generation OPO, a linear cavity. Section 3.4 will detail the motivations for switching to the second generation OPO, a ring cavity design, and will discuss the challenges associated with its construction. A comparison between the two cavities will be made, with the pros and cons of each design considered. In both cases the design parameters were relatively straightforward; the OPO must produce broadband visible pulses of sufficient power for coherent pulse synthesis.

3.3.1 Crystal design & selection

The design of an ultrafast OPO has three main stages. First the wavelength outputs must be chosen for the purpose of the experiment: is the aim is to achieve broadband mid-IR pulses for gas spectroscopy, or to generate tunable visible pulses for biological imaging? The potential wavelength outputs will be constrained by the pump wavelength and the phasematching properties of the nonlinear crystal. The choice of crystal material will depend on the desired centre wavelength and bandwidth(s) of the resonant pulse(s), as well as the type of phasematching employed. The thickness of the crystal will impact the gain of the signal/idler pulses and will also affect the intracavity dispersion of the OPO. Finally the optomechanical design of the cavity will play an important role in the stability and efficiency of the parametric process. Each of these stages must be considered when designing and constructing an OPO.

3.3.1.1 Choice of resonant signal wavelength

For a given pump wavelength there is a finite number of signal and idler wavelengths that can be generated in a parametric process. The sum of the signal and idler frequencies must always equal the pump frequency, such that

$$\omega_p = \omega_s + \omega_i \tag{3.5}$$

or

$$\frac{1}{\lambda_p} = \frac{1}{\lambda_s} + \frac{1}{\lambda_i}. \tag{3.6}$$

This allows the possible signal and idler wavelengths to be determined for a particular pump wavelength. As the aim of this project is the coherent synthesis of ultrashort pulses in the visible, the important wavelengths that must be extracted from the OPO are not the signal or idler but rather the various sum-frequency and second-harmonic visible outputs that will be generated between the pump, signal and idler in the nonlinear crystal. With this in mind, the various visible outputs that would be generated between a 0.800 µm pump source and the signal and idler waves were calculated, with the resultant plot shown in Figure 3.9.

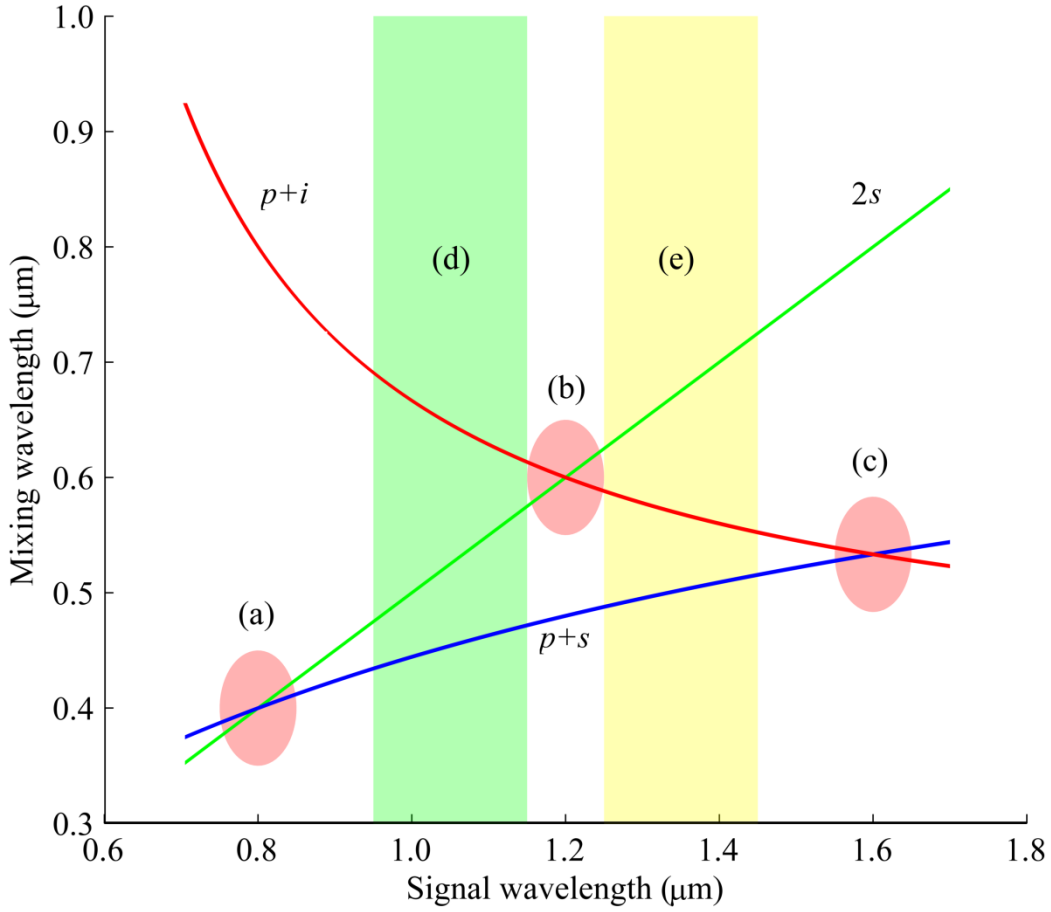

Figure 3.9. Possible sum-frequency and second-harmonic visible outputs between a 0.8 µm pump and increasing signal wavelength. Red line, pump + idler frequency mixing; blue line, pump + signal frequency mixing; green line, signal second-harmonic frequency generation. For details of the shaded areas, see text.

As the signal wavelength is increased in Figure 3.9 three points of intersection are noted. Intersection (a) represents one extreme of the OPO, the case where $\omega_p = \omega_s$. Intersection (b) shows subharmonic signal and idler frequencies, where $3\omega_p = 2\omega_s = \omega_i$, as utilised by Kobayashi et al. for carrier-envelope phase control of an OPO [17]. As this project is concerned with coherent synthesis between *different* visible frequencies from an OPO, the region indicated by (b) is not of interest as two of

the visible frequencies overlap. Intersection (c) indicates degenerate operation of an OPO, where $2\omega_p = \omega_s = \omega_i$, and represents the other bound on possible signal wavelengths. The region near (c) is not desirable as the visible frequencies from the OPO either overlap with each other or with the pump itself. This case was used by Sun et al. to coherently synthesize pulses between a Ti:sapphire oscillator and the second-harmonic signal pulses from a broadband near-degenerate OPO [18], as discussed in Chapter 1.

Using sections (a) and (c) as the possible extremes of signal wavelength choice and eliminating (b), two regions present themselves as likely possibilities. Both (d) and (e) would provide a range of visible wavelengths, however the lower wavelengths available in region (d) allow a wider frequency bandwidth to be accessed for coherent pulse synthesis. For this reason the desired signal wavelength range was chosen to be 0.95 – 1.2 µm.

3.3.1.2 Nonlinear crystal selection

Once a signal wavelength range has been selected, the next step is to select a nonlinear crystal that will allow efficient signal generation over that range. As the Ti:sapphire pump source is broadband and the aim is to produce broadband visible outputs, the resonant signal must also be broadband. This limits the phasematching possibilities as efficient Type I or Type II birefringent phasematching will be limited by bandwidth of the pulses. A common solution is to use Type-0 quasi-phasematching (QPM) to increase the phasematching bandwidth without sacrificing gain. The two common periodically-poled nonlinear crystals used for parametric generation from a Ti:sapphire laser are lithium niobate ($LiNbO_3$, PPLN) and potassium titanyl phosphate ($KTiOPO_4$, PPKTP). Both will be considered before a selection is made. Table 3.1 lists relevant material properties of both crystals.

Table 3.1. Material properties of PPLN and PPKTP [19–24].

Property	PPLN	PPKTP
Transmission range (µm)	0.33 – 5.50	0.35 – 4.50
d_{33} coefficient (pm/V)	25.0	13.7
GDD at 0.5 µm (fs^2/mm)	4673 (o) 3750 (e)	1550
GDD at 0.8 µm (fs^2/mm)	478.8 (o) 383.1 (e)	151.6
GDD at 1.1 µm (fs^2/mm)	74.5 (o) 58.5 (e)	22.1
Requires heating?	Y	N

The transmission range of both crystals is sufficiently broad to allow generation of both the resonant signal wavelength in the near-IR and the doubled pump at 0.4 μm. The d_{33} coefficient of PPLN is almost double that of PPKTP; gain is proportional to $\frac{d_{33}}{n^3}$ however, and so both crystals provide similar gain. PPLN produces significantly more second-order dispersion in the near-IR and visible than PPKTP, and also requires heating to work effectively as PPLN easily suffers from the photorefractive effect. The photorefractive effect is caused by coherent beams of light illuminating a material to produce an interference pattern of dark and light fringes. Electrons within the light fringes are excited into the conduction band of the material, where they flow towards the dark fringes. This leads to an absence of electrons in the light fringes, creating an electric field that in turn causes a change in the refractive index of the material in regions where the field is strongest. A refractive index grating is formed that follows the pattern of the interference fringes, causing the incident light to diffract. Photorefractive effects in PPLN can be alleviated by heating the crystal; PPKTP is much less susceptible to photorefractive effects and so can operate at room temperature.

As both the signal and visible pulses are desired to be broadband (ideally sub-40 fs in duration), the material dispersion of the crystal will have to be compensated intracavity. This limits the length of the crystal to roughly 1 mm for a number of reasons. A longer crystal will have an increased gain and produce more intracavity signal power, however the group delay walk-off between the pump, signal and idler will increase leading to less efficient frequency mixing. Visible frequency mixing is limited to last-coherence-length and fortunate phasematching in the crystal and so a longer crystal will not equate to higher power visible outputs. Longer crystals in a linear cavity will also produce more back conversion into the pump wavelength, reducing the overall efficiency of the visible frequency mixing.

Type-0 quasi-phasematching is the method employed for periodically-poled crystals. Simple phasematching calculations for both PPLN and PPKTP were carried out in order to determine the optimal grating period for a given pump wavelength. A crystal length of 5 mm was chosen in order to quickly visualise the grating periods that would allow phasematching (Figure 3.10, upper panels). The optimal grating period for phasematching to achieve the desired signal output was 27 μm for PPKTP and 18.5 μm for PPLN.

Two considerations must be made when analysing these phasematching diagrams. Firstly, the crystal length will determine the number of grating domains that will generate gain. For PPKTP this number is 18.5 and for PPLN it is 27, both taken over 5 mm. Taking the relative number of grating domains and the relative d_{33} coefficients of each crystal into account, a rough calculation suggests an 8:3 signal gain ratio in favour of PPLN. Secondly, and of more importance to this project, the phasematched output for a broadband input pulse must be determined. The lower panels of Figure 3.10 indicate the possible phasematched output for a single grating in a 0.5 mm crystal for a Ti:sapphire pulse centred at 0.8 μm with a 0.040 μm bandwidth. While PPLN will produce a large gain it can only phasematch to generate a signal with a 0.035 μm bandwidth. With PPKTP a signal pulse with a 0.060 μm bandwidth can be generated, with a potential transform-limited $sech^2$ duration of 20 fs. The phasematching bandwidth available with PPKTP outweighs the gain available from PPLN, and so PPKTP is the OPO crystal used throughout the work presented in this thesis.

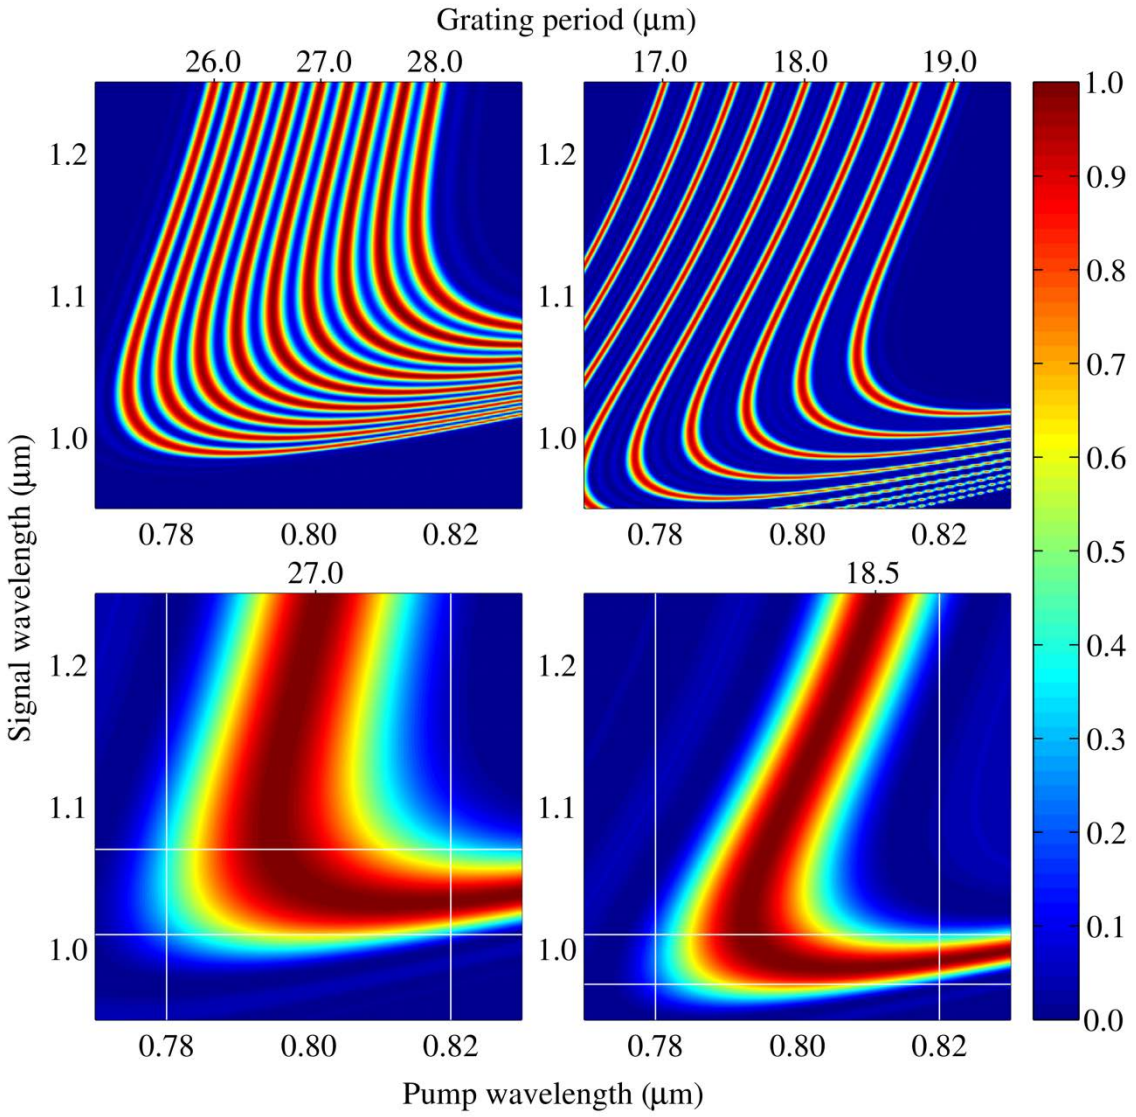

Figure 3.10. Phasematching diagrams for PPKTP (left) and PPLN (right). Upper figures show phasematching curves for a 5 mm crystal length and multiple grating periods. Lower figures show a single grating period that provides broadband phasematching at 0.8 μm for a 0.5 mm crystal.

The design of each PPKTP crystal used is shown in Figure 3.11. Each crystal was grown as a multi-grating, allowing the signal wavelength to be tuned by moving the crystal in the vertical direction. A number of different crystal lengths were purchased and a variety of high-reflective (HR) and anti-reflective (AR) coatings applied. Unfortunately the thinnest crystals (0.3 mm) were damaged during the coating process. These would have minimised the intracavity dispersion while providing sufficient gain to produce parametric oscillation. The next-thinnest crystal size (0.5 mm) was used in all OPO designs in this project.

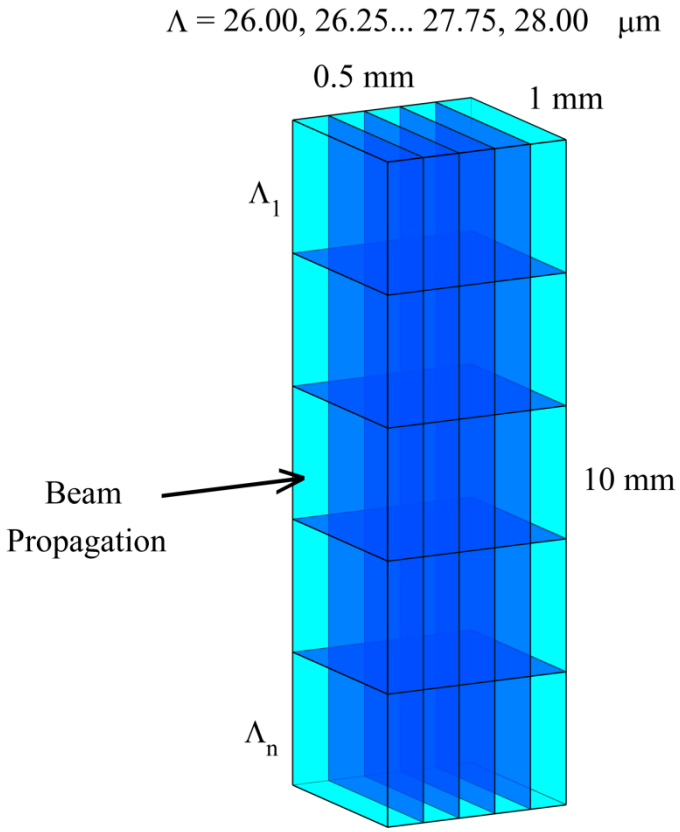

Figure 3.11. Illustration of the PPKTP crystal multi-grating architecture. Each crystal contains nine grating periods allowing for phasematching tunability.

3.3.1.3 Cavity design

The first generation OPO was designed as a linear semi-monolithic cavity, shown in Figure 3.12. The 0.5 mm PPKTP crystal was coated on the input facet with an AR coating from 0.7 – 0.9 μm to reduce pump losses, and with an HR coating from 1.05 – 1.20 μm. The output facet was coated with a broadband visible-near-IR AR coating. These coatings allowed the crystal to act as a cavity end mirror, increasing mechanical stability and reducing the dispersive broadening of the incident pump pulses by removing the need to focus through an intracavity optic.

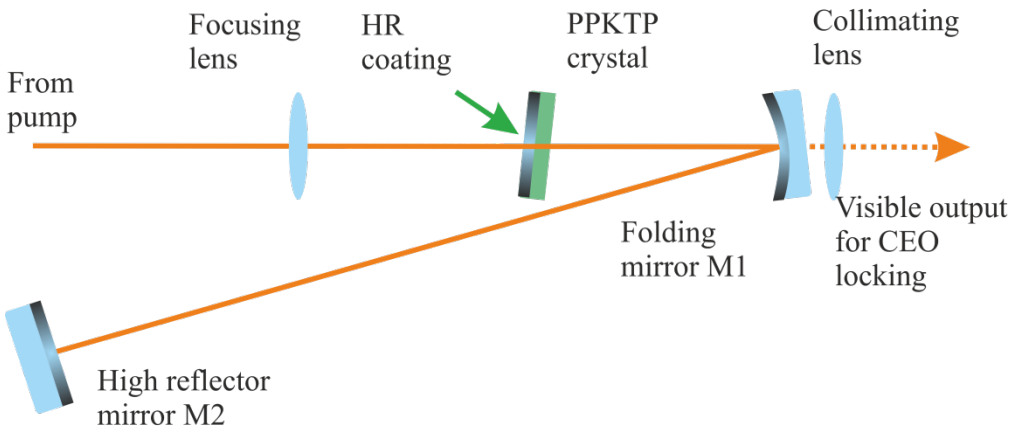

Figure 3.12. Schematic of the first generation OPO. The cavity is a simple dogleg configuration based on a semi-monolithic PPKTP crystal. The crystal is offset from 90° to the pump beam to prevent optical feedback.

An initial attempt was made to use a curved focusing mirror to focus into the crystal, producing a tight focal spot without broadening the pump pulses by propagating through a focusing lens. This proved unsuccessful as the HR coating became damaged under intense focusing. The near-transform-limited pulses from the pump laser had peak powers of almost 0.5 MW, causing irreparable damage to the input surface of the crystal when focused to a spot radius of less than 40 µm. The peak power of the pump pulses is much less than the damage threshold of PPKTP, and so it is more likely that the coating was damaged. In order to find a compromise between pump focal spot and damage considerations the curved focusing mirror was replaced with a lens. The additional chirp introduced by the lens reduced the peak power to a safer level while also allowing for more flexibility in the size of the focal spot.

The intracavity spot size of the resonant signal beam was determined using LCAV, a software package that uses an ABCD matrix algorithm to evaluate the radius of a beam as it propagates through a cavity. LCAV was also used to determine the optimal distance between the output facet of the PPKTP crystal and the curved mirror. A diagram of focal spot size against curved mirror position is shown in Figure 3.13, and is analogous to a stability diagram. The OPO cavity will run most effectively when the curved mirror is placed 50.60 mm after the PPKTP crystal. The intracavity beam profile is shown in Figure 3.14, and the corresponding LCAV program is shown in Table 3.2 for reference.

Table 3.2. LCAV program used to determine the optimal spot size and beam profile of the first generation OPO. All length units are in mm.

```
wavelength = 0.0011;

planemirror;
rod(0.5,1.8);          % 0.5mm plane-plane crystal with n=1.8
space(50.60);
sphermirror(100,3);  % -100mm curvature with 3° folding angle
space(1467);
planemirror;
```

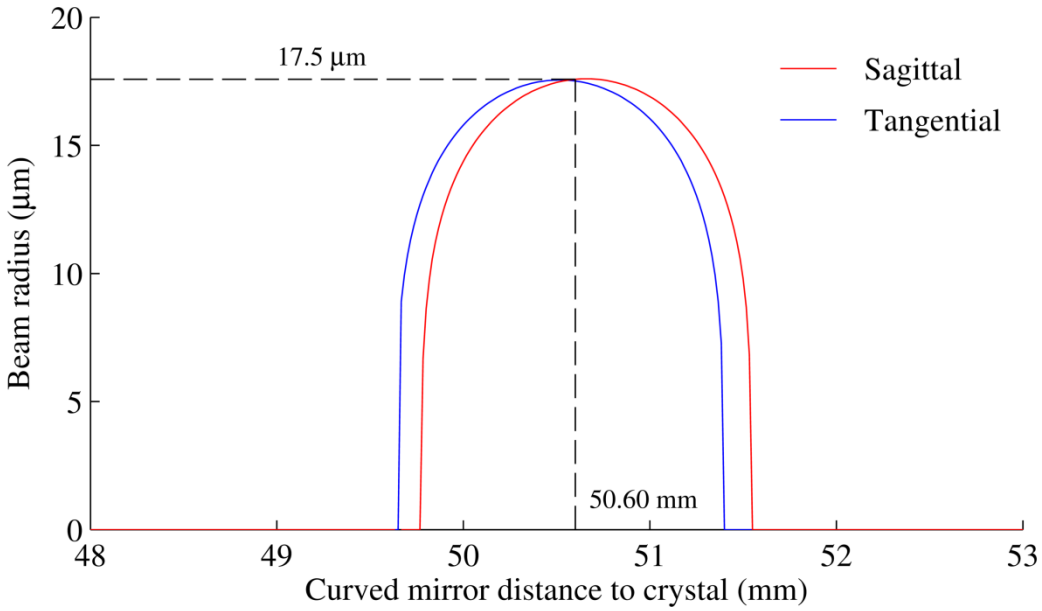

Figure 3.13. LCAV diagram of focal spot size as a function of the curved mirror position for the first generation OPO cavity. The optimal position was 50.60 mm which produced a focal spot of 17.5 µm in the PPKTP crystal. As the cavity will not resonate when the distance between the curved mirror and the crystal is less than 49.79 mm or more than 51.38 mm, this plot can be viewed as a stability diagram.

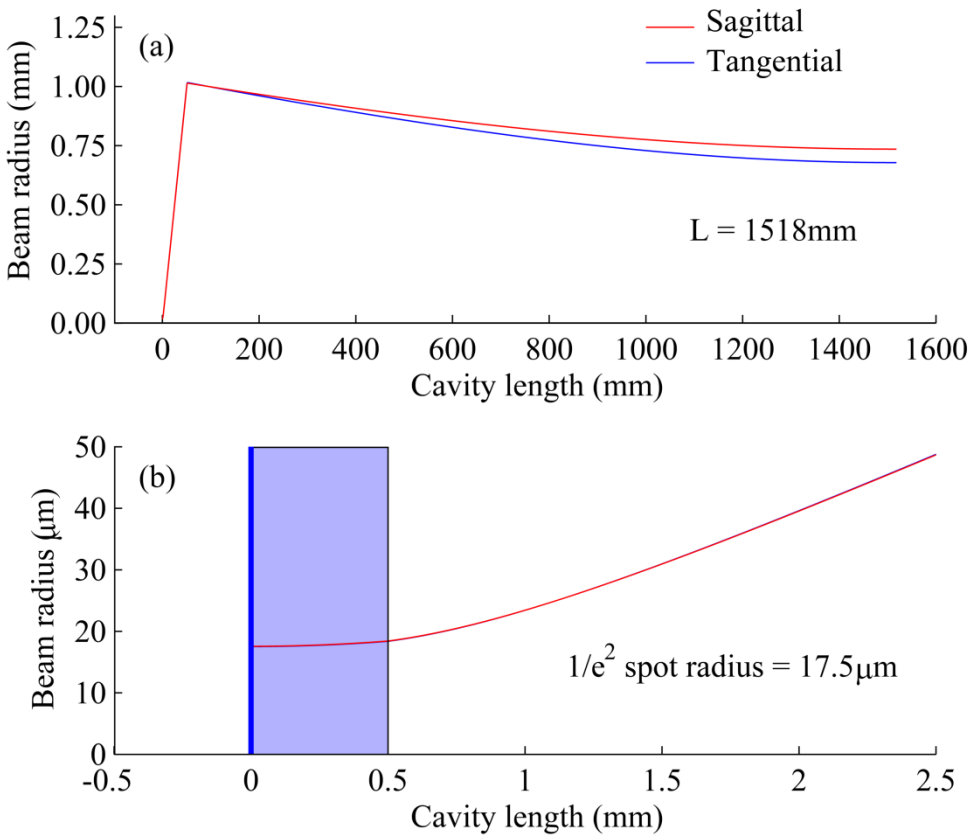

Figure 3.14. Beam profile of the first generation OPO for (a) the full cavity and (b) the crystal focal region.

The PPKTP crystal was glued into a 'Pac-Man' mount, a design that allows for maximum beam clearance on one side of the crystal (Figure 3.15 - right). This design exploits the fact that PPKTP does not require heating and so there is no need to maximise thermal contact with an additional metal plate (Figure 3.15 – left). The crystal mount was located in a lockable Newport 1-inch optical mount that was secured to a vertical translation stage, allowing the grating period to be tuned, and a horizontal translation stage, allowing the crystal to be withdrawn from cavity while maintaining its position relative to the lens and mirror M1. Both the curved mirror M1 and the end mirror M2 were placed in high-stability mirror mounts with differential micrometers (Newport, SL25.4BD). These allowed the cavity to be aligned and optimised with high precision, which is especially important when dealing with a synchronously-pumped cavity. The end mirror mount was attached to a linear motor actuator (Newport, NSA12) allowing the cavity length to be smoothly altered over an 11.6 mm travel range with a 0.1 μm step size.

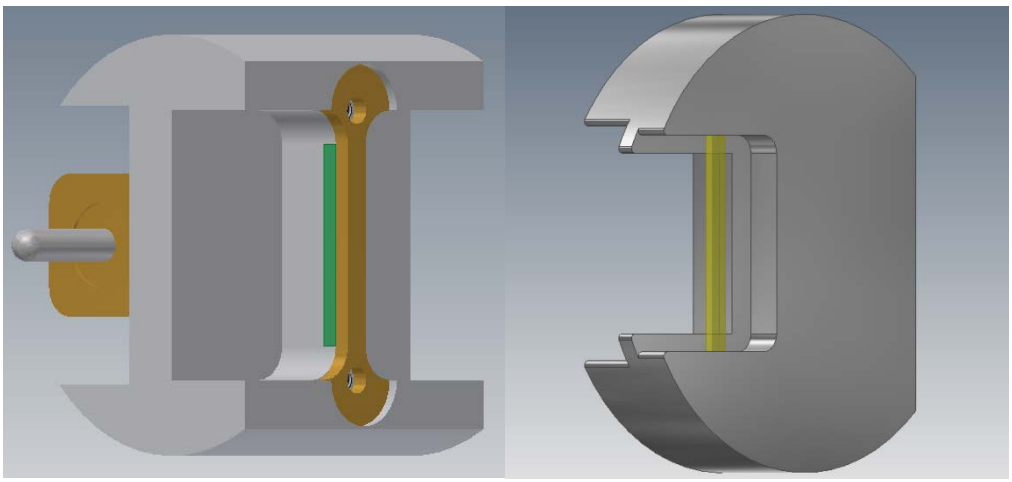

Figure 3.15. Two standard crystal mounts. The design on the left is for PPLN, and includes a resistor to heat the crystal. The design on the right is for PPKTP. The 'Pac-Man' design allows the beam to be steered through the open space without clipping the mount.

3.3.1.4 Cavity alignment procedure

Alignment of the first generation OPO cavity was carried out in a series of rigorous steps that were easily transferable to any OPO cavity design. Having used LCAV to determine the distance between the crystal and mirror M1, and having determined the position of mirror M2 to match the pump cavity length, the alignment procedure itself could now proceed.

As with any alignment procedure it is good practice to ensure the input beam is travelling parallel to the surface of the optical table at a convenient height (here 110 mm), and is also travelling along a set of holes in the table which act as a convenient reference line. This was carried out using a pair of steering mirrors. It is also good practice to attenuate the pump beam to reduce the potential for damaging the optics as they are inserted into the beam line.

The first step was to position the curved mirror M1, which was mounted on a translation stage that moved along the direction of the pump beam. The height of the mirror mount was adjusted so that pump beam was incident with the centre of ½-inch mirror optic. The mirror mount assembly was then secure to the optical bench, leaving sufficient room for the crystal mount and pump lens assembly. A piece of white card was placed over mirror M1 and marked with a cross, so that the pump beam fell on the centre of the cross. A pinhole was placed in the pump beam, reducing the spot on the white card but keeping the beam centred on the cross. This fixed reference was used to monitor how the pump beam deviated as each additional element was added to the beam path.

The second step was to bring the PPKTP crystal into the correct position relative to the pump mirror. The crystal was oriented so that the two end-faces were as normal as possible to the pump beam. The degree of parallelism of the crystal faces will vary for every individual crystal. Using an IR viewer, the angle of the crystal was altered until the pump reflections from the crystal surfaces straddled the pinhole in the pump beam. The translation stage of mirror M1 was then moved until the distance between that mirror and the second crystal surface was 50.60 mm. This was achieved by cutting a piece of card to the correct length and tapering the ends to produce a measurement tool. Gently placing the card against the surface of the crystal, mirror M1 was slowly brought into place until it too touched the card.

The third step was to position the pump focusing lens. The lens had a focal length of 80 mm, was AR coated from 0.7 – 0.9 µm, and was located on a translation stage. The pump beam was blocked while the lens assembly was secured to the optical bench, and the crystal was withdrawn from the beam line to prevent damage during alignment. The pump beam was then unblocked, producing a large spot on the piece of card attached to mirror M1. If the lens was correctly positioned in the horizontal and vertical direction relative to the pump beam then this large spot was centred on the cross on the white card. The angular alignment of the lens was not critical; however it is good practice to align it normal to the pump beam to reduce losses and astigmatism. The pump beam was blocked once more as the PPKTP crystal was brought back into the beam line. When the pump was unblocked, visible frequency doubling of the pump (0.8 µm) to its second harmonic (0.4 µm) was observed. The intensity of the second harmonic beam was increased by moving the focusing lens translation stage, and provided a good visual approximation of the position of the focal spot within the crystal.

The final step was to place the cavity end mirror M2 the correct distance away from mirror M1. LCAV had determined that the correct distance required to match the pump cavity length was 1467 mm. The white card was removed from mirror M1 and the blue SHG beam steered past the crystal and the edge of the focusing lens and through a pair of pinholes. The distance between mirrors M1 and M2 was measured using a tape measure; M2 was placed on a motor-controlled translation stage, and so exact positioning was not necessary for the initial cavity alignment. The blue beam formed a loose waist on mirror M2, and was steered back through the pair of alignment pinholes. The HR coating for the OPO mirrors and crystal input surface was sufficiently reflective at 0.4 µm that the generated SHG was visible for two round trips of the cavity. By

walking the beam with mirrors M1 and M2 the blue beam was overlapped with itself. At this point the attenuation was removed from the pump beam and all pinholes were removed from beam line, noticeably increasing the SHG brightness.

The cavity length could now be altered until oscillation occurred. This was monitored by placing pieces of white card behind mirrors M1 and M2 to monitor the beams that passed through these mirrors. The lab was darkened and the motor-control translation stage attached to M2 slowly scanned. At the point of cavity length synchronicity a flash of colour could be observed on the white cards, a result of various second harmonics and frequency mixing outputs between the pump, idler and resonant signal beams. If the cavity length was scanned too quickly then this flash could be missed; this is why a motorised stage is preferable to a manual version. Once the correct cavity length was established and the visible colours generated, the cavity alignment was optimised by walking the beam using mirrors M1 and M2. Optimisation was performed to maximise the brightness of the generated visible frequencies. The operating threshold for this cavity was 750 mW.

3.3.1.5 Feedback issues and damage considerations

Several practical issues influenced the design and operation of the first generation OPO, and impacted on its optimal performance. The foremost issue, one that also applied to the second generation OPO, was optical feedback. Small back reflections of the pump beam that were propagated back into the laser oscillator would introduce a phase instability, ceasing mode-locked operation. Such back reflections were common from optics that were oriented normal to the input beam; anti-reflection coatings would generally not alleviate the issue, as only a small fraction of the average pump power would be sufficient to break the mode-locking. For lasers operating with longer pulses (more than 100 fs) a practical solution is to use an optical isolator, often composed of a Faraday rotator, a crystal that, when exposed to a strong magnetic field, will rotate the polarisation direction of a beam passing through it. In the near-IR wavelength range the most common Faraday rotator materials are terbium doped borosilicate glass and terbium gallium garnet (TGG) crystal. These materials are highly dispersive, introducing a GDD of 183.5 fs^2 per mm at 0.8 µm, compared to 36.1 fs^2 per mm for fused silica. Second-order dispersion curves for 1 mm of TGG and fused silica are shown in Figure 3.16. The TGG crystal within the available optical isolator (LINOS Photonics, FI-810-5 SV) was 28 mm long, equivalent to 5138 fs^2 of GDD at 0.8 µm.

Using Equation (3.4), and assuming that Gaussian pulses entering the isolator are 22.9 fs as measured in Sub-section 3.2.3, the pulse duration after the isolator would be stretched to 570.4 fs. This amount of GDD is very challenging to remove with mirrors or a prism pair, and so a different approach to removing feedback was implemented.

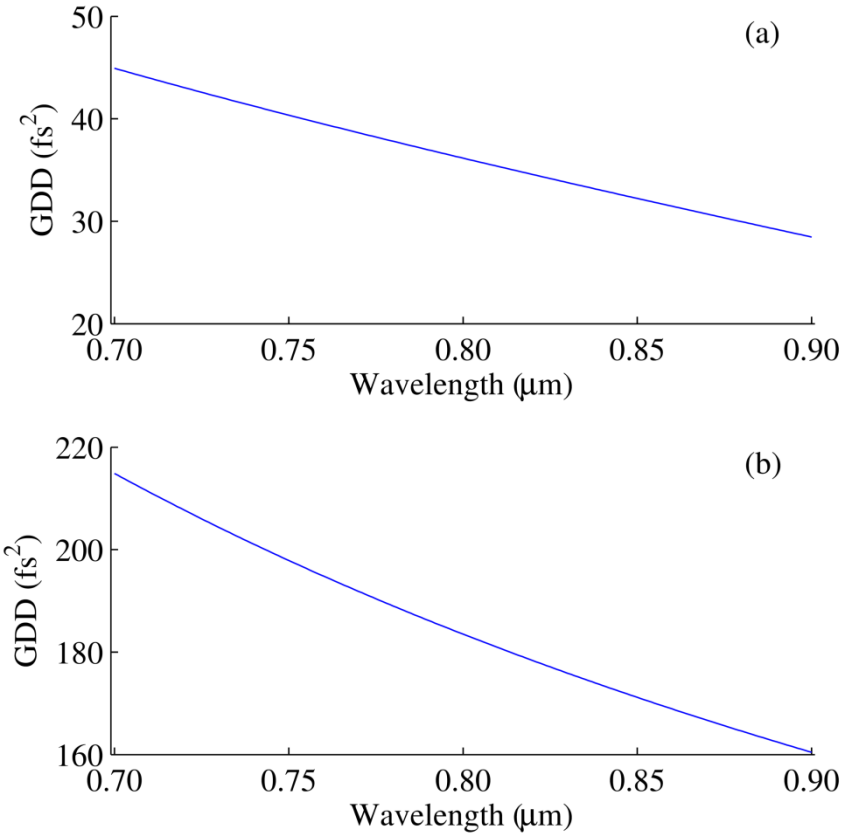

Figure 3.16. Second-order dispersion curves for 1 mm of (a) fused silica and (b) terbium gallium garnet (TGG), highlighting the large GDD produced by the optical isolator.

The alignment procedure described in Sub-section 3.3.1.4 calls for the crystal to be aligned normal to the incoming pump beam. In reality a small horizontal offset was necessary to prevent a reflection of the focused pump beam from returning to the laser oscillator. This reduced the performance of the OPO, an effect that could be observed by increasing the offset and noting a drop in maximum power. The offset was therefore kept to the minimum angle required. This was determined by monitoring the mode-locked pump spectrum as the offset angle was decreased. As the angle was reduced, a number of CW spikes would appear in the spectrum, increasing in bandwidth and intensity until mode-locked operation ceased. The offset angle was adjusted until no distortion was visible in the mode-locked spectrum.

A second practical issue was one of laser-induced damage to the input surface of the crystal. As was discussed in Sub-section 3.3.1.3, the input facet sustained a small amount of damage when a transform-limited pump beam was directly focused onto the surface using a curved mirror. This problem was mitigated by switching to a focusing lens geometry instead, however it was not removed entirely. If the now slightly-chirped pump beam was focused too tightly onto the high-reflective coating then damage would again occur. The minimum focal length that did not induce damage was found empirically to be 80 mm. Further damage was later caused when the intense, focused pump beam was incident on dust particles on the crystal surface. A small spark would appear, occasionally accompanied by a worrying sizzling sound, and the HR coating on the crystal surface would become discoloured and damaged at the focal point. Blocking the input beam, cleaning the crystal surface and moving the crystal so that a clear aperture was available at the focal point allowed OPO operation to continue. Both of these damage issues stem from the thickness of the HR/AR coating on the input face of the crystal.

3.3.2 Characterisation

Several measurements were carried out in order to characterise the performance of the first generation OPO, including temporal and spectral characterisations of the resonant signal pulses, and spectral analysis of the visible mixing pulses as a function of OPO cavity length. At the time of construction no method of measuring the duration of the visible pulses was available; the first temporal characterisation of the visible outputs was carried out for the second generation OPO.

3.3.2.1 Wavelength tuning

Two primary methods were used to tune the wavelength outputs of the OPO; cavity length tuning of the resonant signal pulses was the most direct way of altering the wavelengths of the visible mixing pulses; moving the crystal to select a different grating period provided a much coarser adjustment method. Much of the discussion in this section will deal with cavity length tuning.

The centre frequency of a synchronously-pumped OPO will tune with cavity length [25], and a short derivation of this behaviour will now be given. The synchronous nature of the OPO constrains its operation such that the cavity round trip time (or group delay) is fixed, expressed as

$$\tau(\omega) = \frac{\partial \phi}{\partial \omega} = constant\,. \tag{3.7}$$

A small change in the cavity length L must be reflected by a change in the group delay in order to satisfy Equation (3.7), such that

$$\tau(\omega) + \frac{\Delta L(\omega)}{c} = constant\,. \tag{3.8}$$

Differentiating to find group delay dispersion will give

$$\frac{\partial \tau}{\partial \omega} = \frac{d^2\phi}{d\omega^2} = -\frac{1}{c}\frac{dL}{d\omega}, \tag{3.9}$$

which can be rearranged to give

$$\frac{d\omega}{dL} = \frac{-\frac{1}{c}}{\left(\frac{d^2\phi}{d\omega^2}\right)}. \tag{3.10}$$

This equation describes the gross tuning behaviour of a synchronously-pumped OPO. The centre frequency of an OPO cavity with low GDD will tune quickly with cavity length. Conversely, an OPO with high dispersion (such a cavity with a long nonlinear crystal [26]) will tune slowly with cavity length, making it a good candidate for wavelength stability.

An 8% output coupler was placed in mirror M2 in order to measure the spectrum of the resonant signal; as the signal was not of direct interest to this project it was not routine to have an output coupler in the cavity. Sufficient visible light was reflected from folding mirror M1 and transmitted through the output coupler to measure the visible spectra. The signal pulses were measured using a Heraeus Noblelight E200 laser spectrum analyser that covered wavelengths from 0.8 – 1.5 μm. The visible spectra were measured using two Ocean Optics USB4000 spectrometers; one covered the spectral range from 0.42 – 0.73 μm while one covered 0.66 – 0.93 μm and was used in the second order to measure the SHG pump spectrum at 0.4 μm. Spectra were recorded with each spectrometer as the end mirror was moved by computer controlled stepper motor. The net cavity GDD was calculated as 85 fs^2 at 1.1 μm; such low intracavity dispersion constrained the operating range to only a few tens of microns, making it impractical to change the cavity length by hand. Profiles of output spectra as a function

of increasing cavity length are shown in Figure 3.17 for two different crystal grating periods.

The signal tuned to longer wavelengths with increasing cavity length, as expected from an OPO cavity with positive net GDD and Equation 3.10. Longer grating periods also produced longer signal wavelengths. The signal wavelength was tunable from 1.0 – 1.2 µm, with both lower and upper limits constrained by the mirror reflectivity curves. A number of visible wavelengths were produced as a result of SHG and SFM within the PPKTP crystal, some the result of last-coherence-length phasematching and others due to fortuitous high-order phasematching conditions for a given grating period. The various output wavelengths of the OPO are given in Table 3.3 for a number of resonant signal wavelengths.

Table 3.3. Mixing wavelengths generated by the OPO. The total discontinuous visible frequency bandwidth decreases with increasing signal wavelength.

Primary wavelengths (µm)			**Mixing wavelengths (µm)**				**Visible frequency bandwidth (PHz)**
p	**s**	**i**	**2*p***	***p* + *s***	**2*s***	***p* + *i***	**Δ*f***
0.80	1.00	4.00	0.40	0.44	0.50	0.67	0.29
0.80	1.06	3.26	0.40	0.45	0.53	0.64	0.27
0.80	1.10	2.93	0.40	0.46	0.55	0.63	0.26
0.80	1.14	2.68	0.40	0.47	0.57	0.62	0.25
0.80	1.20	2.40	0.40	0.48	0.60	0.60	0.24

The upper image in Figure 3.17 shows that broader signal bandwidths are achievable for longer grating periods, however the OPO ran more stably at lower wavelengths and so a shorter grating period of 26.50 µm was used for the remainder of the work carried out in this thesis. The centre wavelength of the signal for this grating was 1.06 µm, and the visible pulses produced are listed in Table 3.3. The FWHM of the signal at this centre wavelength and grating period was 0.060 µm.

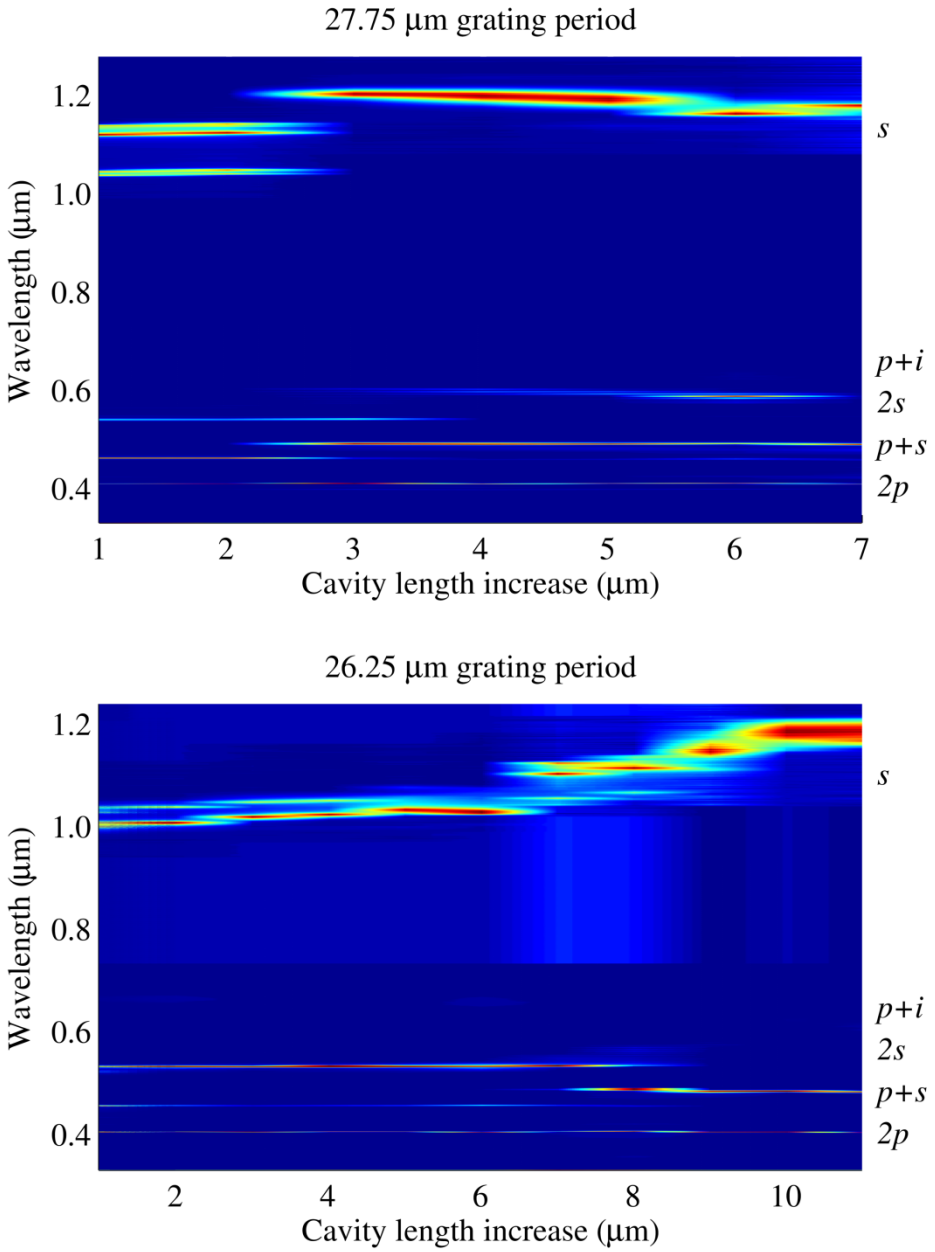

Figure 3.17. Cavity length tuning profiles of the first generation OPO for two separate grating periods. Longer grating periods produced longer signal wavelengths, however the OPO stability was better for shorter grating periods.

3.3.2.2 Dispersion compensation

In order to increase the signal bandwidth, generate shorter signal pulses and hopefully broader bandwidth visible pulses, the intracavity dispersion of the OPO must be minimised, and ideally for stability it should be made very slightly negative. No chirped mirrors or GTI mirrors were available for GDD compensation over 1.0 – 1.2 µm, and so a pair of intracavity prisms was employed, as shown in Figure 3.18. A

pair of SF10 prisms was placed 360 mm apart and the insertion length of each prism varied to try and optimise the signal bandwidth.

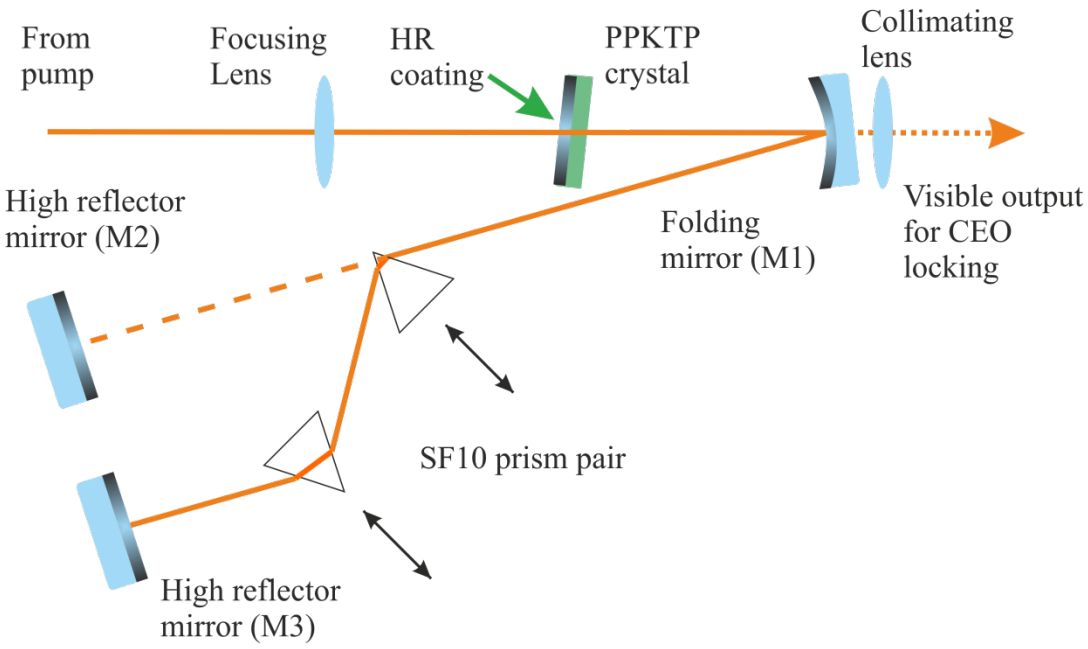

Figure 3.18. Schematic of the first generation OPO with a pair of intracavity SF10 prism for dispersion compensation.

As the OPO cavity must be synchronous in length with pump laser to oscillate, aligning the prisms was not a simple task. The original design was aligned using the second harmonic of the pump light as a beam guide, however the 0.4 µm SHG beam was refracted by a significantly larger angle than the resonant 1.1 µm signal beam, and so a different approach was taken to align the cavity.

The OPO was set to run without prisms in the configuration shown in Figure 3.12. The tip of the first prism was inserted into the cavity to pick off a small fraction of the beam line. The prism was oriented until it reached Brewster's angle for the signal wavelength. This was determined by observing the refracted signal beam from the prism in the far field. As the prism was turned the wavelengths in the pick off beam would move in one direction before stopping and moving in the other direction. The angle at which any wavelength stopped moving was Brewster's angle for that wavelength. Once Brewster's angle had been determined the second prism was placed 360 mm away at a point where it captured the refracted signal beam. This second prism was also oriented at Brewster's angle but with an extra 180° displacement, sending the signal light along a path that was parallel to the original OPO arm. A second high-reflector mirror (M3) was placed after the second prism at a distance that matched the second arm to the pump cavity length and was secured to a motor-controlled stage.

At this point the first prism was fully inserted into the OPO beam line so that oscillation in the first arm ceased. Inserting this prism caused the 0.4 µm SHG light to be refracted, along with any remaining pump light that was not transmitted through mirror M1. The second prism was moved perpendicular to its base so that it intersected the refracted pump beam, sending it towards mirror M3 along a parallel path to the original picked off signal beam. A pinhole was placed in the pump beam between mirror M1 and the first prism, and mirror M3 an angled so that the pump beam returned through this pinhole.

Glass from the first prism was withdrawn until no signal was detected when a power meter was placed between the first prism and mirror M2. Glass was also withdrawn from the second prism until the pump beam missed the optic. The motor-controlled stage attached to mirror M3 was moved until oscillation occurred. The cavity alignment was then optimised using mirrors M1 and M3, a process that was more sensitive than for the uncompensated cavity due to small beam deviations introduced by the prism pair.

By altering the insertion length of the prisms the net dispersion in the cavity could be altered. Each movement of the prism was accompanied by a corresponding change in position of mirror M3, as the cavity length would also be altered. The high reflector in mirror M3 was replaced by an 8% output coupler in order to measure the signal bandwidth. No noticeable improvement in bandwidth was measured as the prism insertion or separation was changed, suggesting that the cavity was operating near zero dispersion to begin with. The GDD curves of the high reflecting mirrors may have been slightly negative near 1.06 µm, compensating for the material dispersion of the crystal; GDD data for the mirror coating was not available from the supplier.

3.3.2.3 Visible light extraction

As the purpose of the project was to coherently combine the visible pulses generated in the OPO it was important to efficiently extract them from the cavity. This was achieved by using different methods for the uncompensated and prism-based cavity designs. In the uncompensated cavity the majority of the visible light exited the cavity through curved mirror M1 due to the transmission properties of the HR infrared coating. A lens was placed immediately after the curved mirror to collimate the visible light into a single white light beam along with the depleted pump. This method was inefficient as much of the visible light was reflected by mirror M1 and exited the cavity via mirror

M2. The large mount holding mirror M1 was chosen for its stability and the differential tuning knobs that allowed for fine control of the cavity alignment, however the size of the mount limited the distance at which the collimating lens could be placed. This limitation produced a visible beam with a diameter of several millimetres, a size which can be cumbersome to work with further down the optical line.

A different approach became available when the prism pair was inserted into the cavity. The material dispersion introduced by the first prism caused the visible and infrared beams to travel along different paths, with only the signal beam being collected by the second prism. An achromatic lens was placed one focal length away from the first prism, creating a line focus at the back focal plane of the lens. A silver mirror was placed at this line focus to reflect the visible pulses back through the lens and prism at a slightly lower height. A silver pick-off mirror (M5) was then used to steer the visible beam further down the optical line. This is illustrated in Figure 3.19.

The dielectric mirror M1 was replaced with a silver mirror of the same curvature in order to steer all the visible light towards the prism pair. The silver mirror could only be employed in the prism-based cavity because the first prism prevented the pump from returning along the beam line. When a silver mirror was placed in the uncompensated cavity sufficient pump power was reflected by mirror M2 to break the mode-locking in the pump laser.

The measured visible powers for each method are shown in Table 3.4 for the most often used wavelength configuration. More power was available using the line focus method as no light was lost through mirror M1, however in both cases the pump + idler power at 0.64 μm was not measureable.

The CEO frequencies of the pump and OPO were locked using the first output coupling method, detailed in Chapter 4, and the cavity was folded in order to box it for additional stability. An attempt was also made to characterise the visible pulse durations using XFROG, however it was unsuccessful due to the low average powers available using this output coupling method. The OPO was reconfigured to the prism-based cavity in order output couple more visible light, however at this point the damage to the input surface of the PPKTP crystal was substantial and oscillation of the OPO was sporadic at best. The decision was made to change the crystal, however as no further HR/AR coated crystals were available the cavity would have to be completely redesigned.

Table 3.4. Visible output powers from the first generation OPO for each method of light collection. In method 1 a lens was placed after mirror M1 to collimate the visible light. In method 2 mirror M1 was replaced by a silver mirror and the visible light was extracted through the prism arrangement shown in Figure 3.19.

Wavelength (μm)	Method 1 power (mW)	Method 2 power (mW)
400	not measureable	3
450	3	8
530	20	45
640	not measureable	not measureable

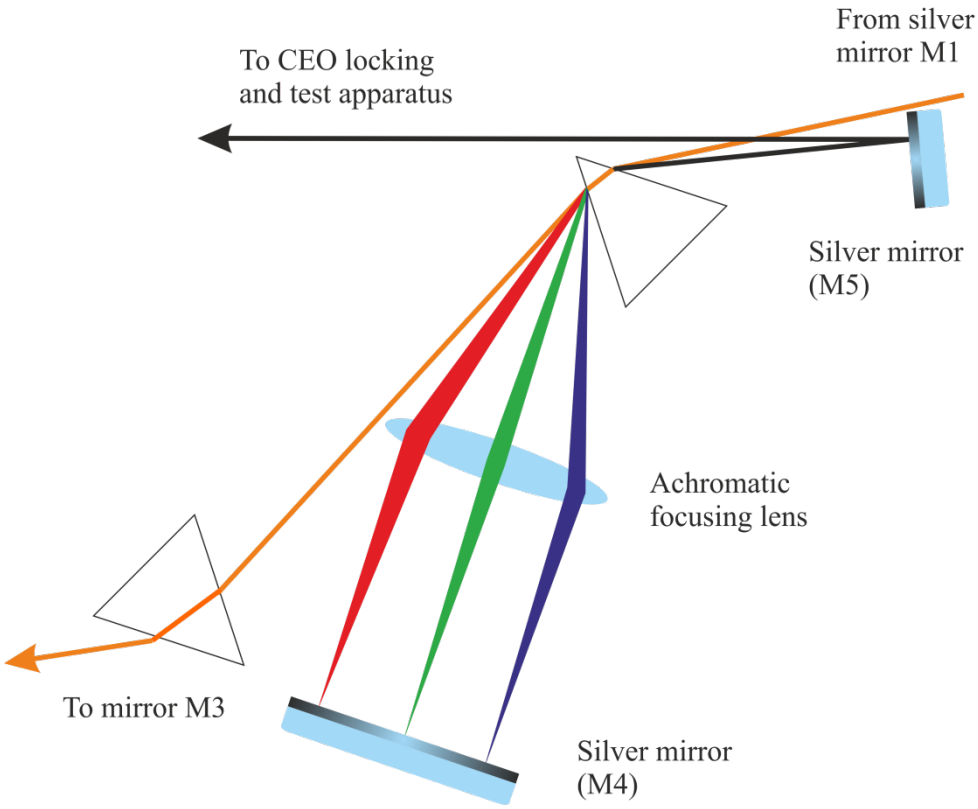

Figure 3.19. Line focus apparatus used to create a collimated beam of visible light from the OPO.

3.4 Second generation PPKTP optical parametric oscillator

As no HR/AR coated crystals were available, the semi-monolithic design used in the first generation cavity could not be repeated. Several AR/AR coated crystals with the same length and grating periods were available, and so it was possible to replace the semi-monolithic cavity with a standard X-cavity design. A number of intermediate OPO cavities were constructed with different configurations in an attempt to maximise the visible output power. Visible light was generated in both propagation directions in all linear configurations, reducing the output power available in the primary beam line.

A ring cavity offers several advantages over a linear cavity for the purpose of visible frequency generation. As the signal pulses only propagate in one direction, all SHG and SFM processes occur in a single pass; in the linear configuration visible light was seen to exit the cavity through the input surface of the crystal and pass through the input lens. The intracavity dispersion in this configuration is reduced to 45 fs^2 as the result of only a single pass through the crystal. Disadvantages in the design are also evident. Operating close to zero dispersion reduces the operational cavity length. Placing prisms into a ring cavity is no simple task, and was not attempted due to time constraints. The cavity length required for synchronous operation at 100 MHz is 3 metres, making it difficult to fold and box in a compact way.

The potential gains from a ring cavity outweighed the immediately identifiable problems, and so construction went ahead.

3.4.1 Cavity Design

The relevant details of the PPKTP grating periods, AR coating and cavity mirror coatings were given in Sub-section 3.3.1, and as they remain the same here they will not be repeated. The cavity design before folding and boxing is shown in Figure 3.20. As two curved mirrors were now in place the focusing geometry was altered and the pump beam was focused through mirror M4 using a 60 mm focal length lens. As with the first generation OPO the cavity stability and beam size were modelled in LCAV, details of which are given in Table 3.5, Figure 3.21 and Figure 3.22.

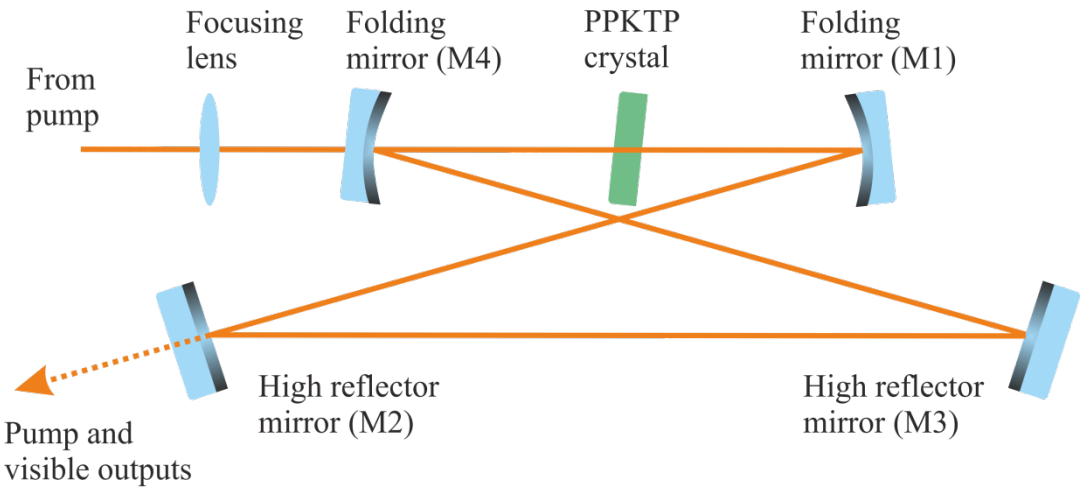

Figure 3.20. Schematic of the second generation OPO in a ring configuration.

Table 3.5. LCAV program used to determine the optimal spot size and beam profile of the second generation ring OPO. All length units are in mm.

```
wavelength = 0.0011;
type       = ring;

sphermirror(100,3);  % -100mm curvature with 3° folding angle
space(50.97);
rod(0.5,1.8);        % 0.5mm plane-plane crystal with n=1.8
space(50.50);
sphermirror(100,3);  % -100mm curvature with 3° folding angle
space(770);
space(1423);         % Spaces reflect X-geometry of ring cavity
space(770);
```

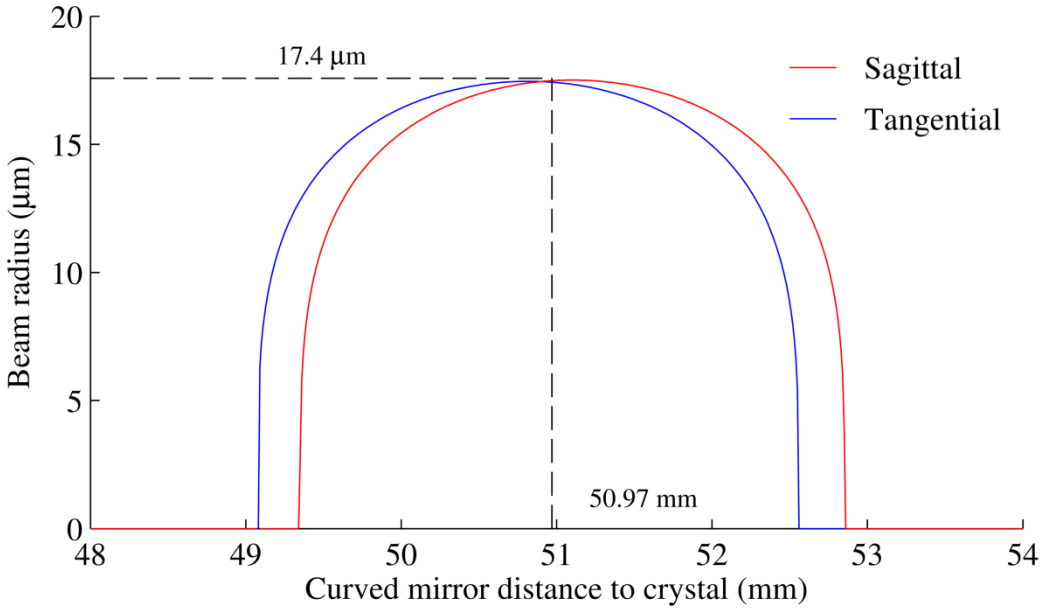

Figure 3.21. LCAV diagram of focal spot size as a function of the curved mirror position for the second generation ring OPO. The optimal position was 50.97 mm which produced a focal spot of 17.4 µm in the PPKTP crystal.

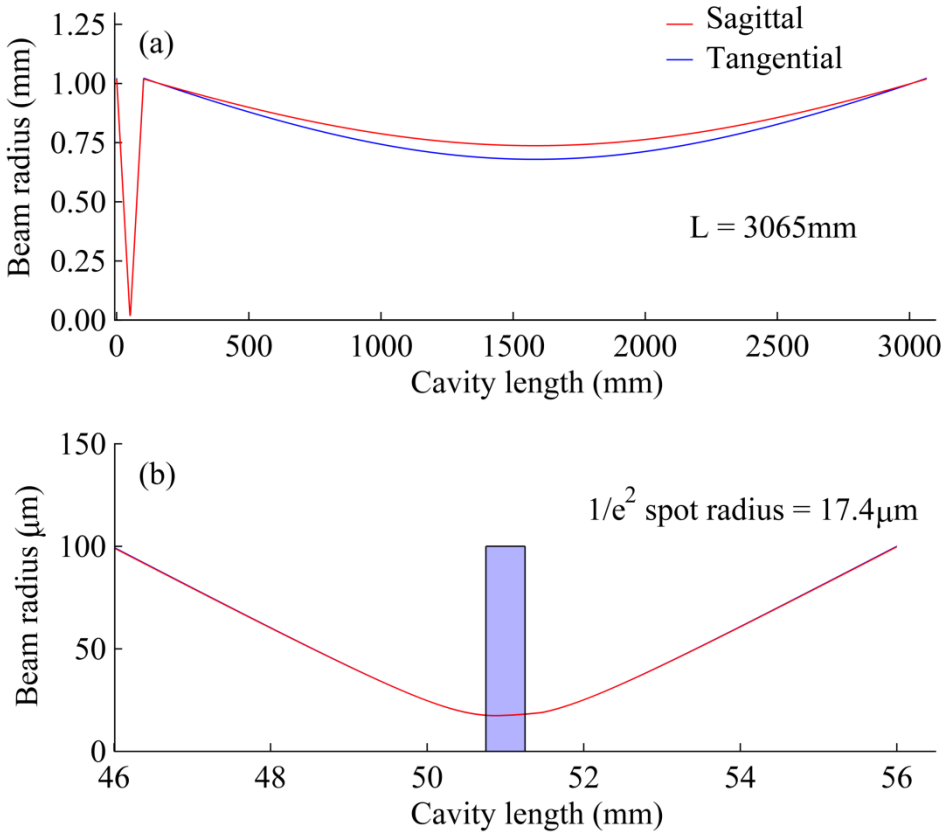

Figure 3.22. Beam profile of the second generation ring OPO for (a) the full cavity and (b) the crystal focal region.

Looking at Figure 3.21 and Figure 3.22 in combination with Table 3.5 it is evident that for the crystal to be situated at the focal point of the intracavity beam, the distance between the crystal faces and each of the curved mirrors should not be equal, which is not the case for a linear cavity. The AR/AR coated crystal was mounted in a similar 'Pac-man' style mount as the HR/AR coated crystal as it provided good clearance for beam steering. The crystal mount was secured in a lockable Newport 1-inch optical mount attached to an XYZ translation stage. Mirror M4 was a ½-inch diameter, -100 mm radius of curvature high reflector held in a ½-inch Newport mirror mount that provided a clear edge for beam clearance. Mirror M1 was an identical optic, however it was held in a 1-inch adaptor and secured in a lockable 1-inch Newport mount with a clear edge for beam clearance. This allowed the dielectric mirror to be replaced by a silver mirror for visible light collimation in the manner of the first generation OPO; the cavity was aligned using only dielectric mirrors to keep intracavity losses to a minimum. Mirrors M2 and M3 were mounted in high-stability mirror mounts with differential micrometers to allow the beam to be walked with accuracy, with the mount of mirror M2 attached to a motor controlled stage for cavity length adjustment.

3.4.2 Cavity alignment procedure

The alignment procedure for a ring OPO is different to that of a semi-monolithic design, however many principles are the same. The distances between the crystal faces and the curved mirrors must be determined beforehand, as must the total cavity length necessary to achieve synchronous operation.

The first step was to attenuate the pump beam and align it so that it travelled parallel to set of holes in the optical bench. The mount for curved mirror M4 was then secured to the bench, leaving enough room for the pump lens on a translation stage behind the optic. The position of this mirror mount remained fixed, with the crystal and curved mirror M1 positioned relative to the front surface of mirror M4. After checking that the pump beam passed clearly through mirror M4 without clipping, the optic was removed.

Mirror M1, secured to a translation stage, was then placed roughly 100 mm away from the position of mirror M4. The height of mirror M1 was adjusted so that the attenuated pump beam was incident on the centre of the optic. A piece of white card with a cross on it was attached to the mirror so that the pump beam fell on the centre of the cross. A pinhole was placed in the pump beam to ensure that the pump spot on the white card hit the exact centre of the cross.

The PPKTP crystal was then placed into the path of the pump beam, ensuring that the XYZ stage had sufficient travel to remove the crystal from the beam path, travel the full vertical height of the crystal and move through a region that included a distance roughly 50 mm away from mirrors M1 and M4. The crystal was oriented until the pump reflections from the front and back surfaces of the crystal straddled the pinhole in the pump beam.

The next stage was to insert the focusing lens assembly into the beam. The pump beam was blocked and the crystal moved out of the beam path. The 60 mm focal length lens was placed in a 1-inch Newport mount that provided a clear aperture on one edge for beam clearance. The lens mount was placed roughly 10 mm behind the mount for mirror M4 and attached to a translation stage to optimise the focal position. The pump beam was unblocked, producing a large spot on the cross attached to mirror M1. The horizontal and vertical positions of the lens were adjusted until the large spot was centred on the cross, at which point mirror M4 could be reinserted into its mount. Mirror 4 was angled slightly so that the beam expected to come from mirror M3 would be returned towards mirror M1. This angle introduces a horizontal deviation in the path of the pump beam as it is focused through the mirror, causing the spot on the card to shift to one side. If the pump beam is travelling parallel to the surface of the optical table and passes through the centre of the lens and mirror M4, no vertical displacement should be observed.

With the optical head of the OPO in place, the distances between the crystal faces and the two curved mirror could now be set. Two thin strips of card were cut to lengths of 50.97 mm and 50.50 mm and their ends tapered to a point. The pump beam was blocked, the crystal brought back into the beam line and the white card on mirror M1 removed. The 50.97 mm length piece of card was placed against the front surface of mirror M4 and the position of the crystal adjusted until its input face touched the other end of the card. The 50.50 mm length of card was then placed against the output surface of the crystal and the position of mirror M1 adjusted until it touched the other end of the card. This procedure placed the crystal at the centre of the intracavity beam. The pump beam was unblocked and the pinhole was removed. Second harmonic generation at 0.4 μm was immediately visible from the crystal. The position of the focusing lens was altered until maximum SHG was observed.

The blue SHG beam was steered with mirror M4 so that it passed mirror M1 and the focusing lens without clipping the edge of either optic. Mirror M2 was attached to a motorised translation stage and was placed 770 mm away from mirror M1. Mirror M2 was angled to steer the beam parallel to the pump beam. Mirror M3 was positioned 1423 mm from mirror M2, placing it 770 mm away from mirror M4. This symmetric cavity configuration kept the folding angles M1-M2 and M3-M4 equal. The blue beam was steered with mirror M3 to place it on the centre of mirror M4.

Once the cavity optics were positioned the beam alignment procedure could begin. As with the first generation OPO, feedback was observed when pump light was reflected from the surface of the crystal and so the crystal needed to be angled. A brief discussion will first be given to aligning a ring OPO for a cavity where feedback is not an issue, and then a full description of the alignment procedure for a ring OPO with an angled crystal will be given.

For a crystal angled normal to the pump beam, cavity alignment is a simple matter. SHG light generated in the crystal is aligned in the forward direction to return to the crystal. SHG light is also observed in the backwards direction, reflected from the output face of the crystal. The beam can be walked in forwards direction with mirrors M1 and M2 and in the backwards direction with mirror M3 and M4. The two beams are walked until they overlap at various points in the cavity, typically at pinholes places in the 770 mm length arms. Care must be taken to ensure that the overlapping beams come from the front surface reflections of each cavity mirror. The cavity length can then be adjusted until oscillation is observed.

For a crystal oriented at an angle to prevent optical feedback the alignment is more difficult. The SHG light that is reflected off the output face of the crystal does not return along the input path of the pump beam and so does not follow a path that would allow the generated signal pulses overlap. As no beam is available for overlap in the backwards propagating direction, the forwards direction only must be used. The blue SHG light is passed once round the cavity until it is returned to the crystal. The light is then looked for again on mirrors M1 and M2, appearing as a very dim spot near the bright SHG light from the first trip through the crystal. It is often necessary to turn off all bright lights in the lab to see this second spot. The second spot must be overlapped with the beam from the first crystal pass, and is steered using mirrors M3 and M4. As the second pass beam is walked to overlap with first pass one, dim spots of SHG light

may by observed on mirrors M3 and M4. The beam is walked until no secondary reflections can be observed, at which point the SHG light from each crystal pass is overlapped. Finally, the cavity length can be altered using the motorised translation stage until oscillation occurs. While it is true that changing the position of mirror M2 will alter the position of the SHG beam on all subsequent optics, for small movements the change in beam overlap is minimal. Once oscillation occurs the alignment can be optimised by walking the beam with mirror M2 and M3.

3.4.3 Characterisation

In a ring configuration the OPO cavity length was a little over 3 metres, with the largest mirror separation almost 1.5 metres. This large footprint led to instability due to cavity length drift and air currents, and so it was necessary to fold and box the OPO. An image of the folded OPO is shown in Figure 3.23. The folding procedure was simple and efficient provided the folding angles were kept small and enough SHG pump light was visible after several mirror bounces to allow the cavity to be aligned. The folded cavity contained eight plane mirrors, reducing the footprint to the size of an A3 sheet of paper. Six of the mirrors were secured in Newport VIZIX single-axis lockable mirror mounts. These compact mounts only provide beam steering in the vertical direction, with horizontal steering achieved by rotating the mount and post assembly. The mounts also provide large clearance apertures for beam steering, making them ideal for use in a compact folding geometry. The remaining two mirror mounts were the same high-stability Newport mounts with differential micrometers that were previously used in the first generation OPO. It was necessary to place these mounts at different ends of the cavity instead of in sequence, as this provided more sensitivity when walking the beam. The threshold of the ring OPO was 750 mW, the same as the semi-monolithic design. While the ring cavity should offer a decrease in operating threshold compared to the semi-monolithic design, the increased number of bounces in the folded configuration leads to higher intracavity loss.

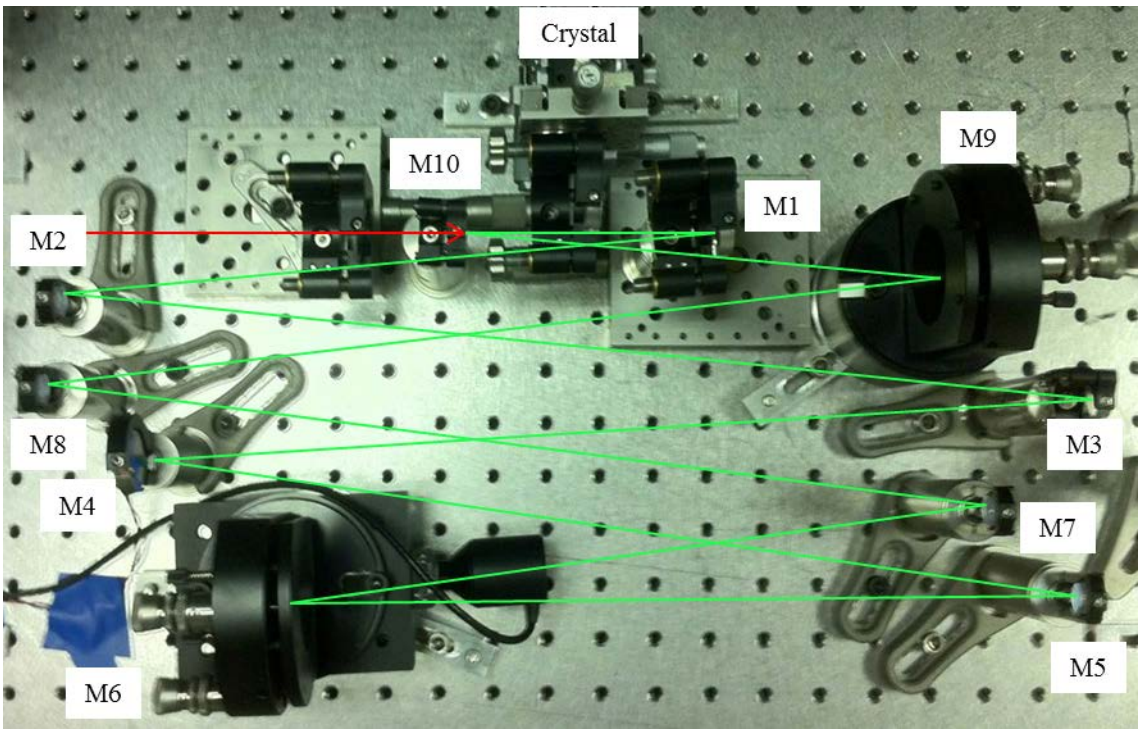

Figure 3.23. Picture of the second generation OPO in a folded configuration. The primary output coupler for the pump and visible light was mirror M2. Mirror M4 was attached to a piezoelectric transducer for CEO control.

3.4.3.1 Pump depletion

Most of the pump light exited the cavity through mirror M2 where it could be redirected for use in further experiments. The depleted pump spectrum is shown in Figure 3.24.

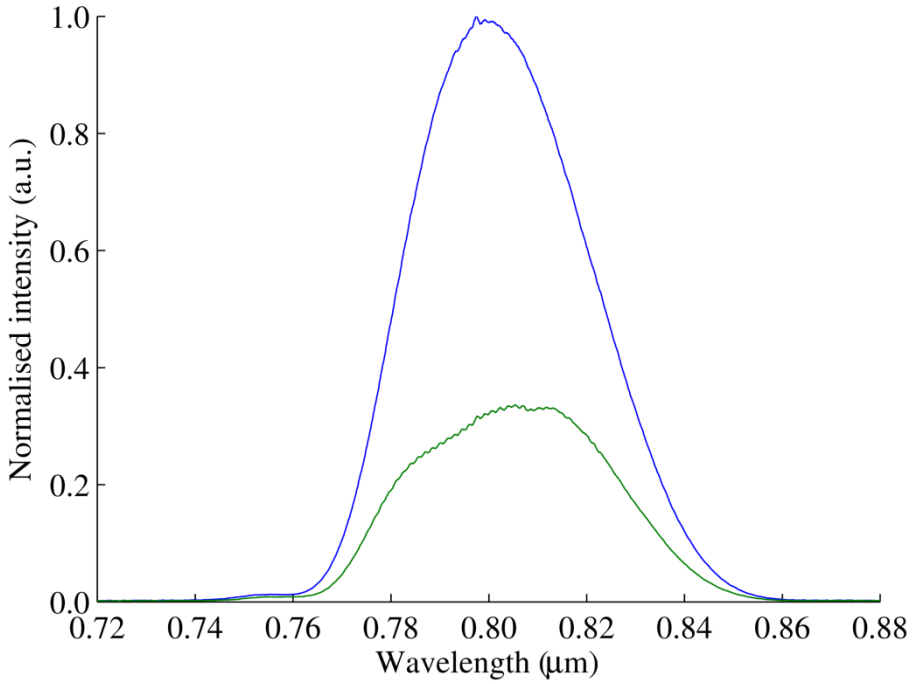

Figure 3.24. Undepleted (blue) and depleted (green) pump spectrum.

3.4.3.2 Cavity length tuning

The visible spectra were recorded as the cavity length was tuned. The low intracavity dispersion yielded a short and rapid tuning range, with high sensitivity to alignment and pump cavity length changes. A map of the visible outputs as a function of cavity length is shown in Figure 3.25.

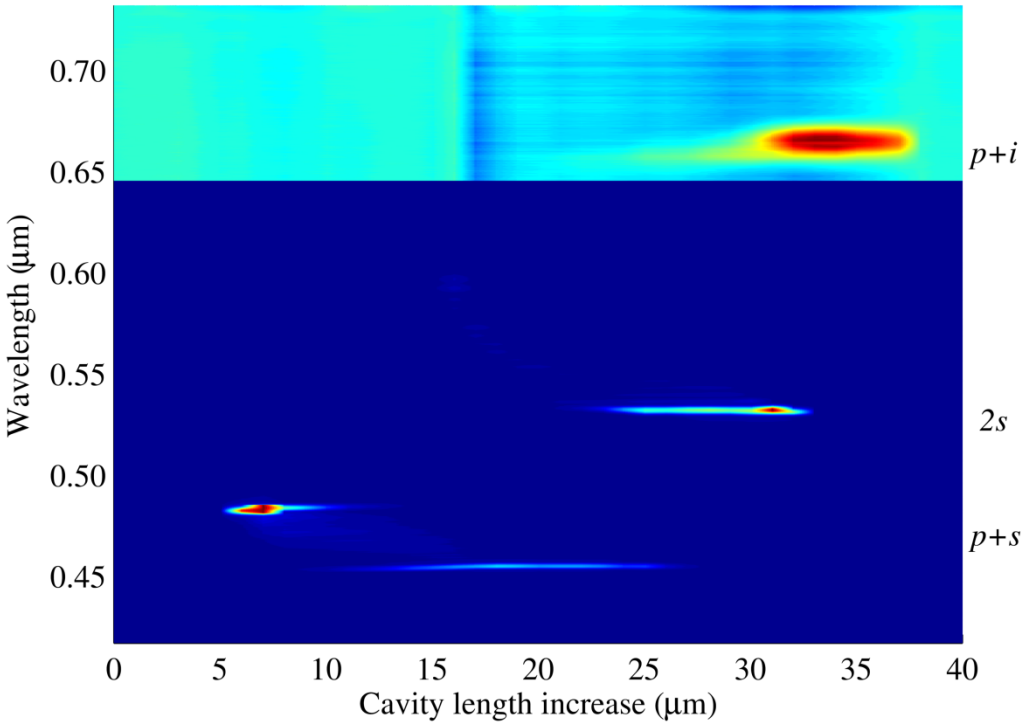

Figure 3.25. Visible spectral outputs from the second generation ring OPO as the cavity length was increased. Recording the weak pump + idler beam required a longer integration time, producing a poorer signal-to-noise ratio.

3.4.3.3 Visible power performance

The primary motivation for switching to a ring cavity was the potential for an increase in the SHG and SFG efficiencies within the PPKTP crystal. The visible light was output coupled through mirror M2 where it could be characterised and used in further experiments. The power levels of each of the visible pulses were measured after they were spatially dispersed with an SF10 prism. The results are shown in Table 3.6 along with the power levels recorded for the first generation linear cavity.

Table 3.6. Comparison of the visible output power levels from the OPO in a ring and linear configuration.

Wavelength (μm)	Ring cavity power (mW)	Linear cavity power (mW)
400	10	3
450	30	8
530	100	45
650	4	not measureable

The single-pass constraint of the ring cavity allowed for efficient frequency mixing between the pump, signal and idler waves, reducing back conversion and confining the visible mixing pulses into a single beam. Most of the visible power was extracted as leakage light through mirror M2, however sufficient power was reflected and output coupled through mirror M3 to be used for CEO locking, as discussed in Chapter 4.

3.5 Visible pulse characterisation using XFROG

Characterisation of the visible pulses was non-trivial as the wavelengths involved precluded a single- or two-photon autocorrelation measurement. With help from our collaborators at Oxford University an SFG-XFROG experiment was designed to measure the visible pulse durations.

3.5.1 Design

The FROG experiment was centred on a 30 µm BBO crystal cut at 45.6° for Type II SFG between the 0.8 µm pump and the visible pulses. Both the pump and visible beams were horizontally polarized, however Type I phasematching was not possible over such wide frequency bandwidths. A schematic and photograph of the FROG set-up are shown in Figure 3.26. The pump beam acted as the gating arm and was separated from the visible pulses using a visible-NIR beam splitter. The visible pulses were recombined with the pump pulses using a polarizing beam cube after their polarization was rotated to vertical. The two beams were overlapped so that they were collinear and then focused into the BBO crystal using an off-axis parabolic mirror. A curved aluminium mirror was placed off-focus after the crystal to allow the generated UV light to be focused and steered into an Ocean Optics UV spectrometer that covered the spectral range 0.227 – 0.450 µm. A UV-pass filter was placed before the spectrometer to block the pump and visible wavelengths.

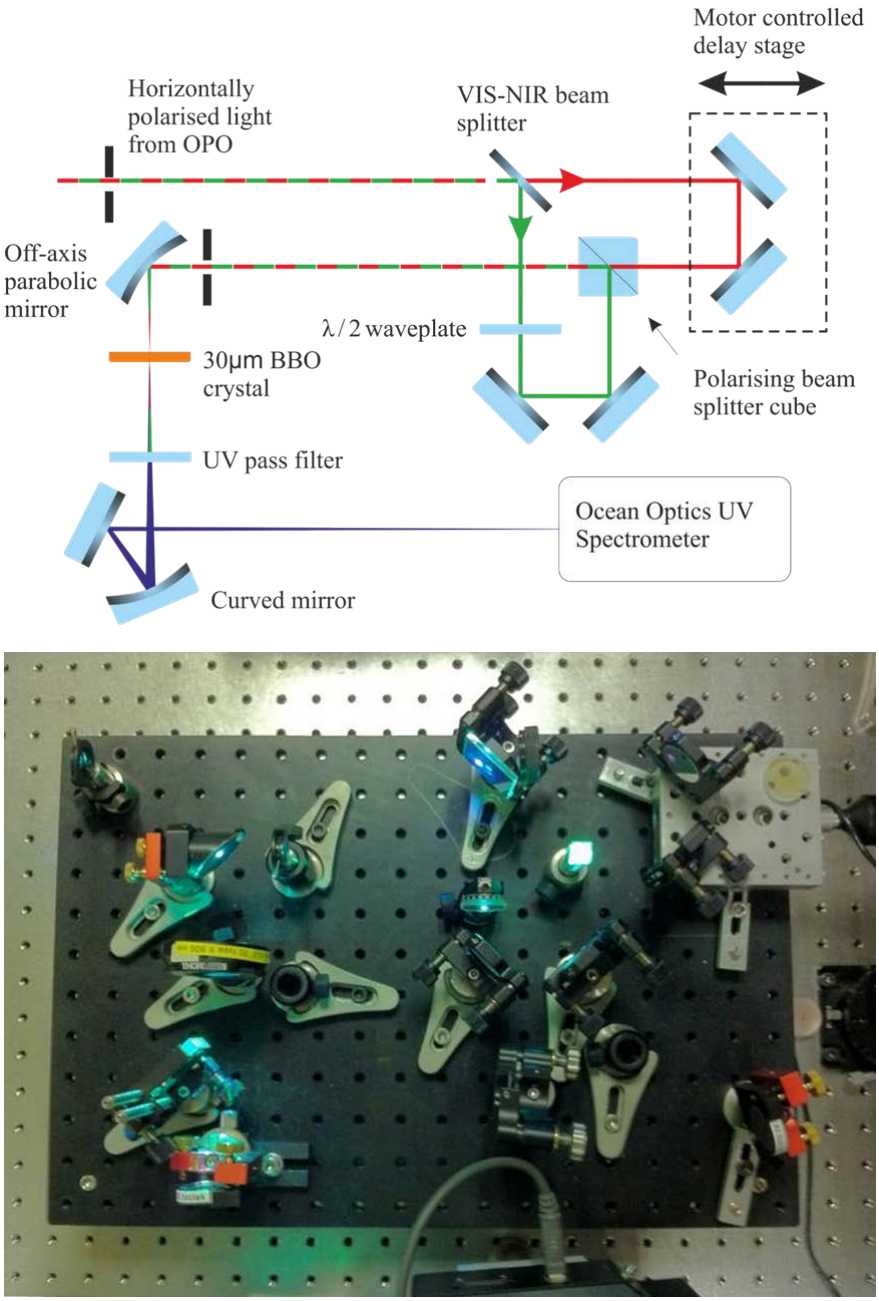

Figure 3.26. Schematic (top) and photograph (bottom) of the XFROG experiment.

3.5.2 Alignment

It was critical that pump and visible beams were perfectly collinear to maximise the SFG signal. Checking for spatial overlap by eye was found to be insufficient, and so a webcam was used to check for overlap in the focal plane of the off-axis parabolic mirror and in the far field. A variable ND filter was placed in the beam path before the FROG breadboard to reduce the power incident on the webcam. When the webcam was placed in the focal plane of the parabolic mirror two spots were visible, shown in the left hand image of Figure 3.27. The spot from the pump beam was assumed to be fixed; the pump steering mirrors attached to the motor control stage were aligned to send the

outgoing beam back parallel to the incoming beam. The visible beam was steered until the focal spots overlapped. The webcam was then moved to a point between the off-focus collimating mirror and the spectrometer and the beam overlap checked again, as shown in the right hand image of Figure 3.27. If the beams were not overlapped the visible beam was walked until far field overlap was achieved. This whole process was repeated until the beams overlapped in both positions.

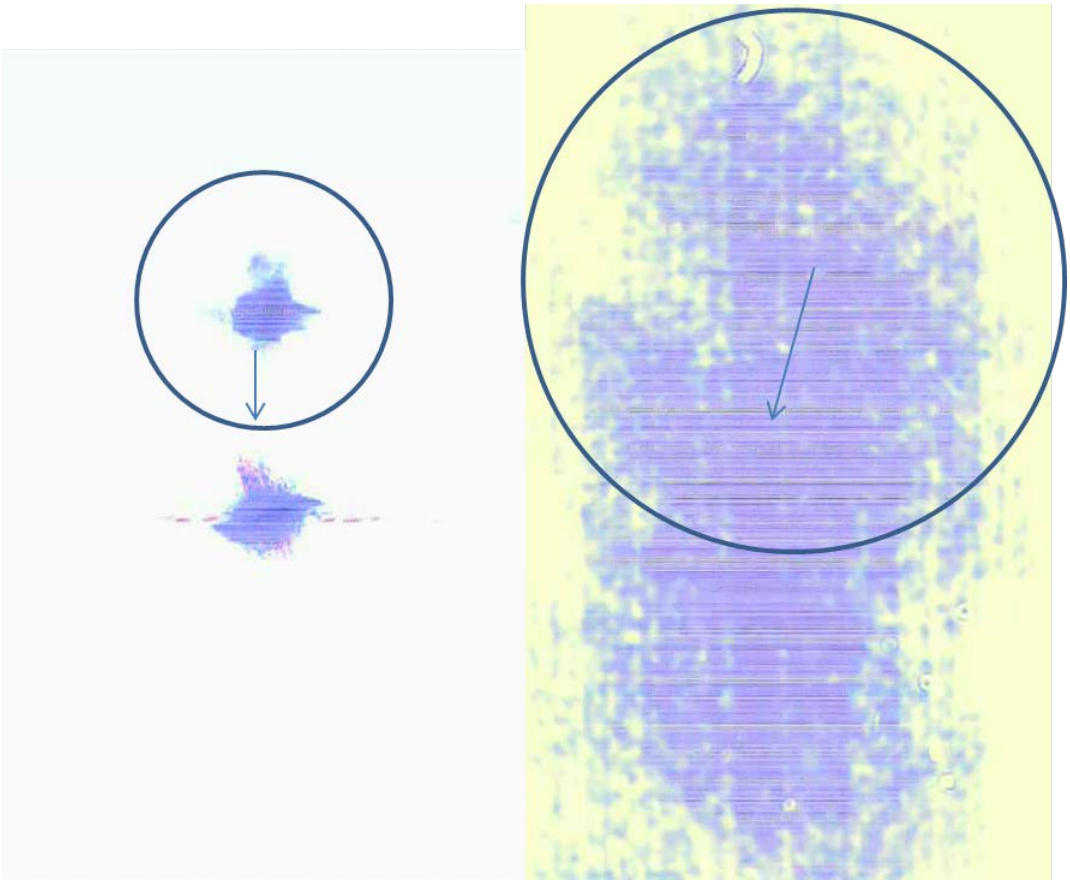

Figure 3.27. Webcam images of the pump and visible beams in the focal plane of the parabolic mirror (left) and the far field (right). The visible beam was walked until it overlapped with the pump beam in both positions.

Once collinearity had been established, the SFG crystal was located in the focal plane of the parabolic mirror. This was located by first placing a near-infrared pass filter after the polarizing beam splitter cube to remove the visible light and allow through the strong pump beam. By rotating the crystal around the axis normal to the input beam the phasematching condition was changed from Type II SFG to Type I SHG for a 0.8 μm input. The crystal was positioned by maximising the SHG signal before rotating the crystal once more to minimise the SHG. This crystal angle was optimal for SFG between the pump and visible beams. The SFG beam was focused into the Ocean Optics spectrometer using plane and curved mirrors with an aluminium coating with

extended reflectivity in the UV region. A UV bandpass filter was placed before the spectrometer to remove the intense pump and visible light, providing 80% transmission from 0.24 – 0.395 µm.

3.5.3 Pulse measurement and retrieval

The collinear pump and visible beams were steered into the FROG apparatus. A half-wave plate with a centre wavelength of 0.5 µm was placed in the path of the visible beam to rotate its polarization to vertical. Scanning the gate arm over a 1 mm travel range allowed the XFROG traces between the pump and the visible OPO outputs to be recorded simultaneously, as shown in Figure 3.28. This trace allowed the delay between the visible pulses to be clearly observed.

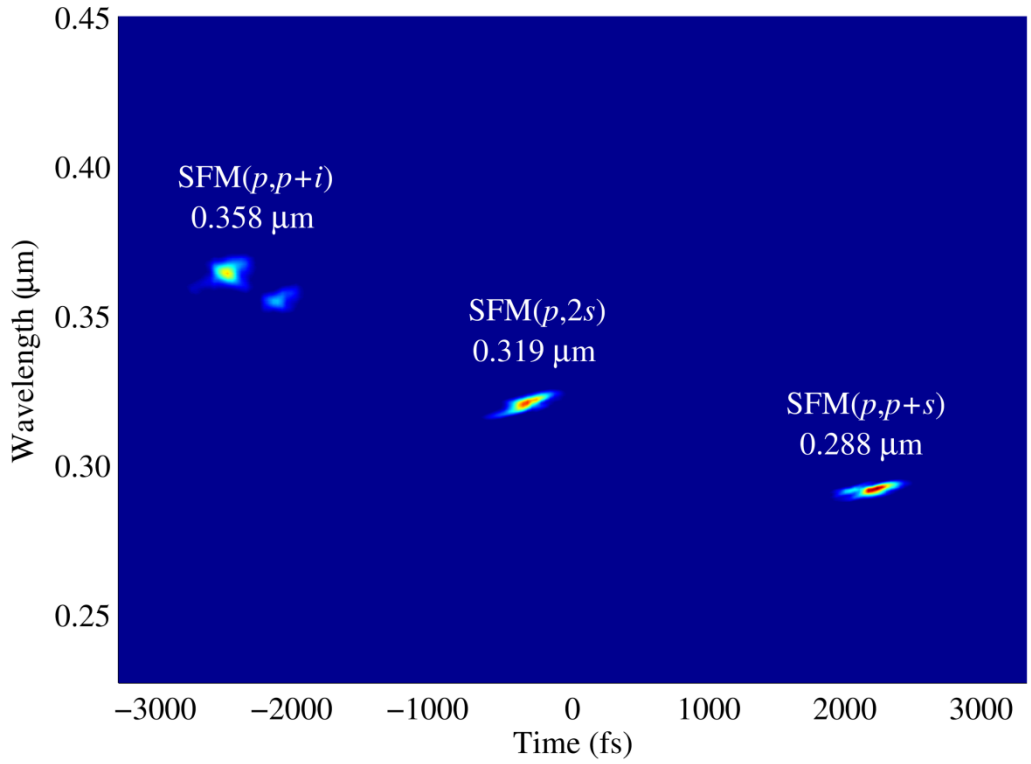

Figure 3.28. Simultaneous XFROG traces between the pump and visible OPO outputs. The two smaller traces evident for the pump-idler pulses arise from the spectral shape of this pulse.

Data retrieval was carried out using the techniques discussed in Chapter 2, Subsection 2.5.2 of this thesis. As the simultaneous recovery of several pulses is non-trivial however, several steps will be discussed in detail.

Firstly the interacting pulses were simulated to produce an ideal FROG trace with no relative delay between the pulses. Each pulse was assumed to be transform-limited with duration of 25 fs, as shown in the top panel of Figure 3.29. XFROG traces for these

ideal pulses were modelled and are shown in the lower panel of Figure 3.29. The 0.4 µm pulse and its corresponding XFROG trace are shown for completeness; however its XFROG trace was not measured experimentally due to the limiting bandwidth of the UV bandpass filter.

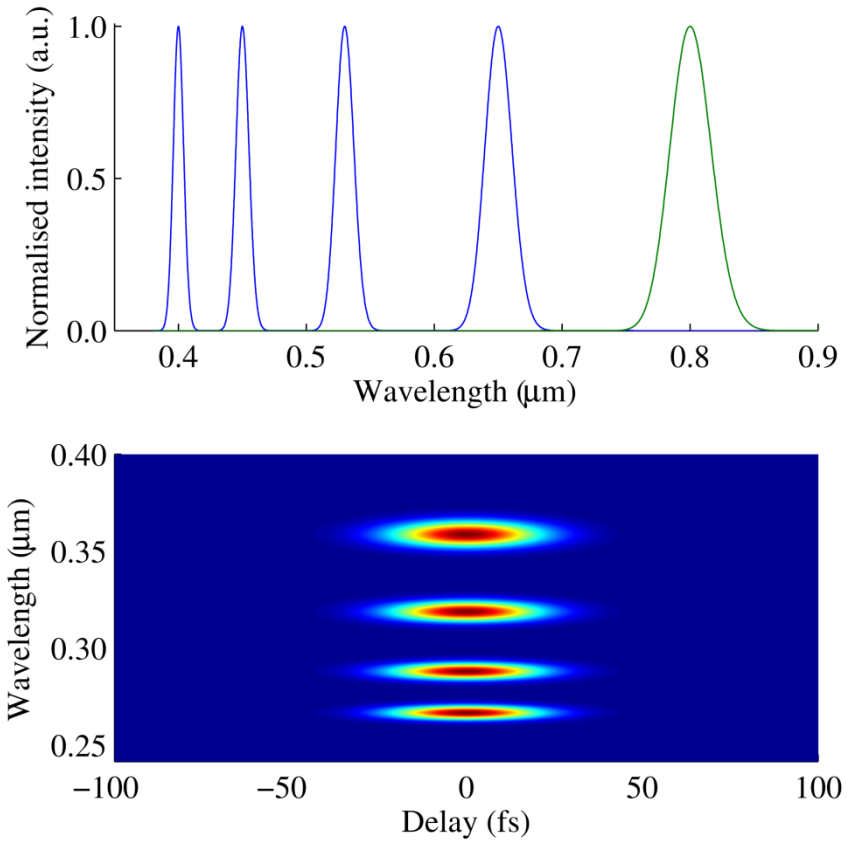

Figure 3.29. Upper figure: spectral bandwidths supporting 25 fs pulses for the simulated pump and visible pulses. Lower figure: simulated XFROG traces for the pump and visible pulses.

A starting guess for the FROG retrieval algorithm was calculated by using the spectral bandwidth of the simulated XFROG trace. A synthetic broadband pulse was generated and a duplicate made. XFROG was carried out between these two pulses, and the initial pulse bandwidth modified until the bandwidth of its XFROG trace exceeded that of the trace shown in Figure 3.29. The synthetic guess pulse and the corresponding XFROG trace are shown in Figure 3.30.

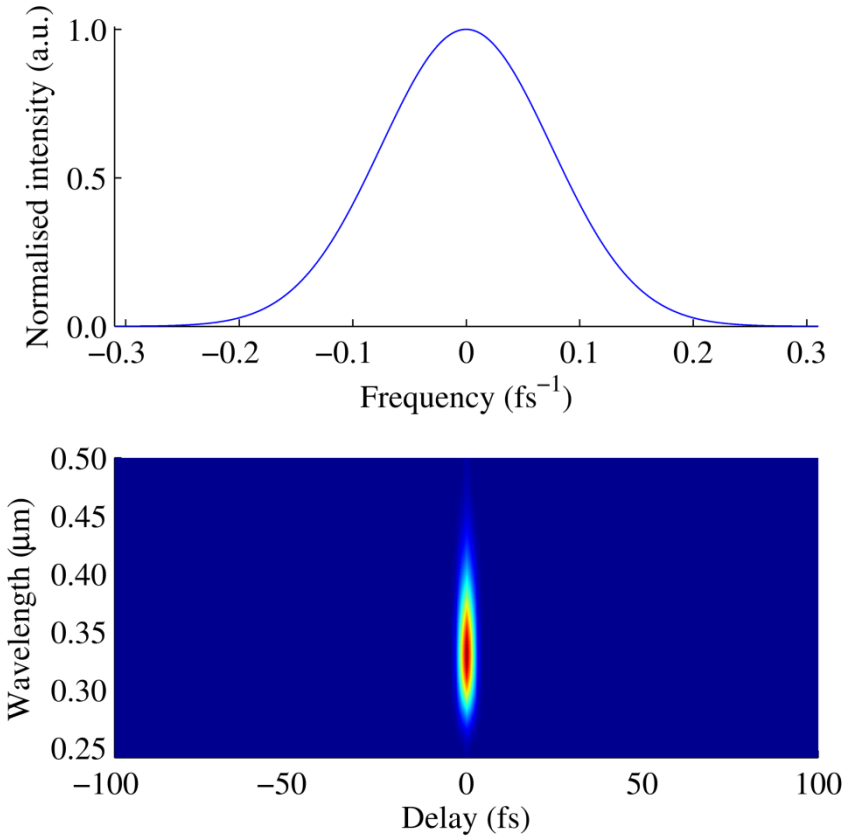

Figure 3.30. Upper image: frequency bandwidth of the synthesised guess pulse. Lower image: XFROG trace between the synthesised pulse and an identical copy.

The experimental data had to be processed in several ways before the FROG retrieval algorithm could be used. Individual background traces from both the test and gate arms were recorded and subtracted from the full trace. This was particularly important when measuring FROG traces near 0.4 µm as residual Type I SHG from the pump beam raised the noise floor of the spectrometer. The experimental data was then interpreted in time and frequency onto a 4096 x 4096 Fourier grid. The FROG retrieval algorithm was processed for 10 iterations, with the algorithm converging to a result that closely matched the measured trace after 6 iterations. The calculated XFROG traces are shown for 2, 4 and 6 iterations in Figure 3.31. The standard PCGP algorithm is robust enough to account for the relative delay between the different pulses, which it treats as separate spectral components of the same pulse.

The pump and visible pulse durations were calculated from the retrieved FROG trace after 40 iterations of the PCGP algorithm. To improve and constrain the retrieved results the calculated pump spectrum was replaced with the experimental pump spectrum before each iteration. The final retrieved FROG trace is compared alongside

the experimental FROG trace in Figure 3.32. The retrieved visible spectra are shown beside the experimentally measured spectra in Figure 3.33, along with their calculated pulse durations. Both images show good agreement between the experimental and calculated data, indicating that the FROG algorithm is suited to retrieve information on multiple pulses at once.

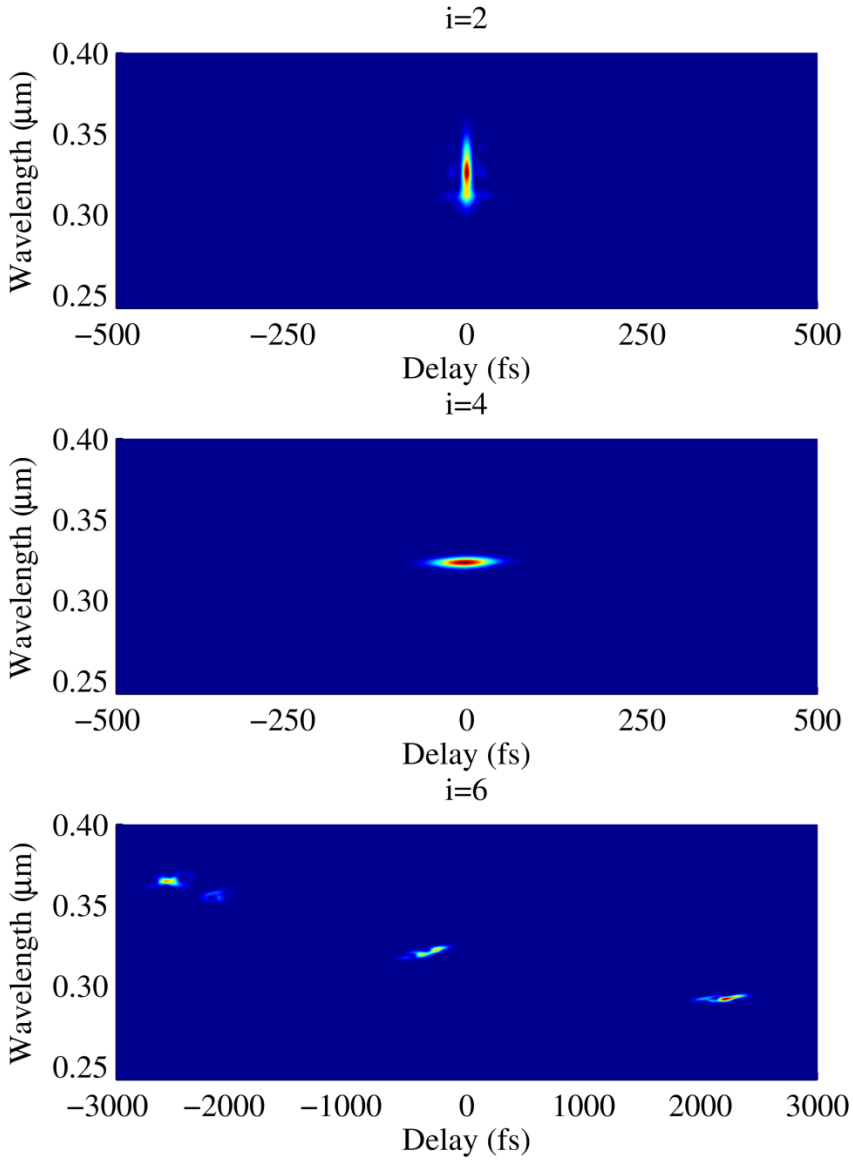

Figure 3.31. Calculated XFROG traces for i=2, 4 and 6 iterations of the FROG retrieval algorithm. After 6 iterations the algorithm has correctly retrieved the relative delay between the SFG pulses.

The quality of the experimental trace obviously affects the retrieved results; the intensity of the generated UV pulses depends strongly on alignment and crystal phasematching. Slight spatial walk-off between the visible pulses meant that, in a collinear configuration, it was not possible to ensure simultaneous overlap between the pump beam and each visible beam, and so a compromise was made. Additionally,

while the 30 µm BBO crystal allows for broadband phasematching, the phasematching efficiency will not be the same for each SFG interaction. Spatially separating the visible pulses to allow their relative delays to be controlled would improve their spatial overlap in the FROG apparatus; this is discussed in Chapter 6.

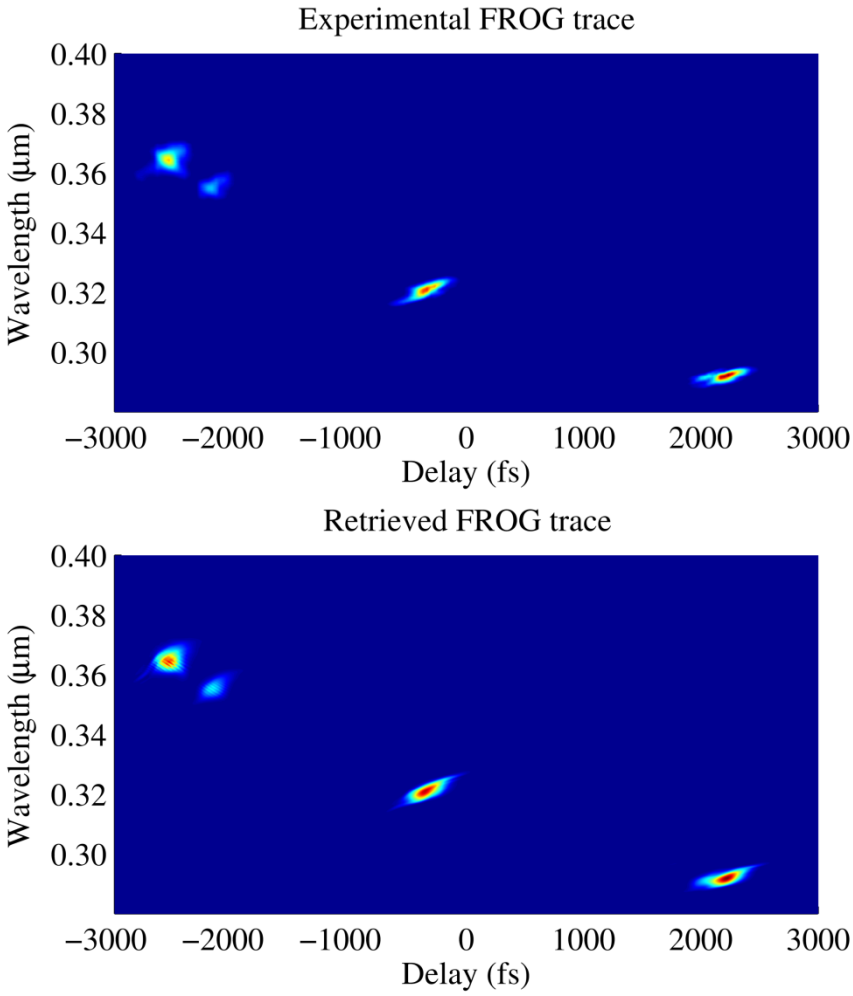

Figure 3.32. Top: experimental FROG trace. Bottom: calculated FROG trace using the PCGP retrieval algorithm.

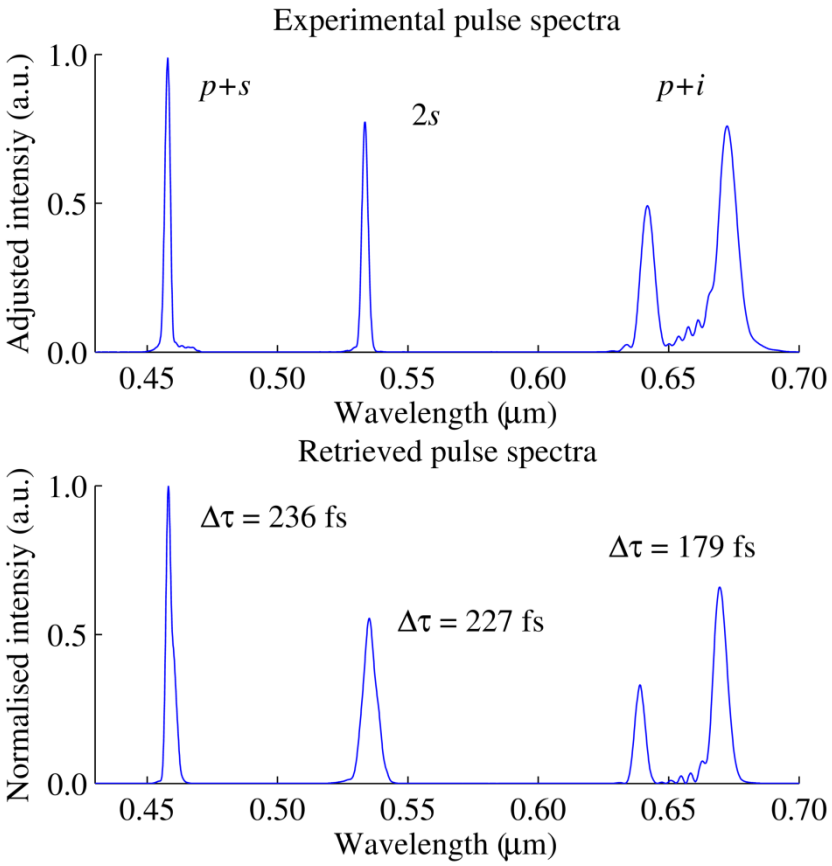

Figure 3.33. Top: Experimentally measured visible spectra. The relative intensities of the measured spectra have been adjusted for ease of comparison. Bottom: Calculated visible spectra. The pulse durations were calculated from the retrieved FROG data.

The calculated durations of the *p*+*i*, 2*s* and *p*+*s* pulses were 179 fs, 227 fs and 236 fs respectively. These pulse durations were not unexpected; the spectral bandwidths of the individual pulses support shorter pulse durations, however the pulses were chirped as they were output coupled from the cavity. The FROG apparatus also introduces some intrinsic chirp in the form of a beam splitter, polarising beam cube and half-wave plate. Pulse compression is further discussed in Chapter 6.

3.6 Conclusions

A 20-fs-femtosecond Ti:sapphire laser was characterised and used to synchronously-pump a PPKTP-based femtosecond OPO. Two OPOs were designed, constructed and characterised, with a ring cavity proving more efficient for visible pulse generation than a linear cavity. An XFROG apparatus was constructed to measure the duration of the visible pulses.

3.7 References

1. D. E. Spence, P. N. Kean, and W. Sibbett, "60-fsec pulse generation from a self-mode-locked Ti:sapphire laser," Opt. Lett. **16**, 42–4 (1991).

2. T. Brabec, C. Spielmann, P. F. Curley, and F. Krausz, "Kerr lens mode locking," Opt. Lett. **17**, 1292–4 (1992).

3. H. A. Haus, J. G. Fujimoto, and E. P. Ippen, "Analytic theory of additive pulse and Kerr lens mode locking," IEEE J. Quantum Electron. **28**, 2086–2096 (1992).

4. U. Morgner, F. X. Kärtner, S. H. Cho, Y. Chen, H. A. Haus, J. G. Fujimoto, and E. P. Ippen, "Sub-two-cycle pulses from a Kerr-lens mode-locked Ti:sapphire laser," Opt. Lett. **24**, 411–3 (1999).

5. D. H. Sutter, G. Steinmeyer, L. Gallmann, N. Matuschek, F. Morier-Genoud, U. Keller, V. Scheuer, G. Angelow, and T. Tschudi, "Semiconductor saturable-absorber mirror assisted Kerr-lens mode-locked Ti:sapphire laser producing pulses in the two-cycle regime," Opt. Lett. **24**, 631–3 (1999).

6. A. E. Siegman, *Lasers*, First edition (University Science Books, 1986).

7. O. Svelto, *Principles of Lasers*, Fourth edition (Plenum Press, 1998).

8. D. T. Reid, "Few cycle EM pulses," Contemp. Phys. **40**, 193–204 (1999).

9. J. D. Simon, "Ultrashort light pulses," Rev. Sci. Instrum. **60**, 3597–3624 (1989).

10. C. Spielmann, F. Krausz, T. Brabec, E. Wintner, and A. J. Schmidt, "Femtosecond pulse generation from a synchronously pumped Ti:sapphire laser," Opt. Lett. **16**, 1180–2 (1991).

11. N. Sarukura and Y. Ishida, "Pulse evolution dynamics of a femtosecond passively mode-locked Ti:sapphire laser," Opt. Lett. **17**, 61–3 (1992).

12. U. Keller, K. J. Weingarten, X. K. Franz, D. Kopf, B. Braun, I. D. Jung, R. Fluck, H. Clemens, N. Matuschek, and J. A. Der Au, "Semiconductor saturable absorber mirrors (SESAM's) for femtosecond to nanosecond pulse generation in solid-state lasers," IEEE J. Quantum Electron. **2**, 435–453 (1996).

13. D. V Khudyakov, A. S. Lobach, and V. A. Nadtochenko, "Passive mode locking in a Ti:sapphire laser using a single-walled carbon nanotube saturable absorber at a wavelength of 810 nm," Opt. Lett. **35**, 2675–7 (2010).

14. I. H. Baek, H. W. Lee, S. Bae, B. H. Hong, Y. H. Ahn, D.-I. Yeom, and F. Rotermund, "Efficient mode-locking of sub-70-fs Ti:sapphire laser by graphene saturable absorber," Appl. Phys. Express **5**, 032701 (2012).

15. "Timewarp autocorrelator," http://www.light-speed-tech.com/TimeWarp.htm.

16. U. Keller, "Ultrashort Pulse Generation," in *Proceedings of SUSSP 52 - Advances in Lasers and Applications*, First edition (Institute of Physics Publishing, 1999), pp. 83–115.

17. Y. Kobayashi, H. Takada, M. Kakehata, and K. Torizuka, "Optical phase locking among femtosecond subharmonic pulses," Opt. Lett. **28**, 1377–9 (2003).

18. J. Sun and D. T. Reid, "Coherent ultrafast pulse synthesis between an optical parametric oscillator and a laser," Opt. Lett. **34**, 854–6 (2009).

19. K. Fradkin, A. Arie, A. Skliar, and G. Rosenman, "Tunable midinfrared source by difference frequency generation in bulk periodically poled KTiOPO[sub 4]," Appl. Phys. Lett. **74**, 914–6 (1999).

20. H. Zhao, I. T. Lima, and A. Major, "Near-infrared properties of periodically poled KTiOPO4 and stoichiometric MgO-doped LiTaO3 crystals for high power optical parametric oscillation with femtosecond pulses," Laser Phys. **20**, 1404–09 (2010).

21. J. D. Bierlein and H. Vanherzeele, "Potassium titanyl phosphate: properties and new applications," J. Opt. Soc. Am. B **6**, 622–633 (1989).

22. I. Shoji, T. Kondo, A. Kitamoto, M. Shirane, and R. Ito, "Absolute scale of second-order nonlinear-optical coefficients," J. Opt. Soc. Am. B **14**, 2268–94 (1997).

23. O. Gayer, Z. Sacks, E. Galun, and A. Arie, "Temperature and wavelength dependent refractive index equations for MgO-doped congruent and stoichiometric LiNbO3," Appl. Phys. B **91**, 343–8 (2008).

24. O. Gayer, Z. Sacks, E. Galun, and A. Arie, "Erratum to: Temperature and wavelength dependent refractive index equations for MgO-doped congruent and stoichiometric LiNbO3," Appl. Phys. B **101**, 481 (2010).

25. D. C. Edelstein, E. S. Wachman, and C. L. Tang, "Broadly tunable high repetition rate femtosecond optical parametric oscillator," Appl. Phys. Lett. **54**, 1728–30 (1989).

26. Z. Zhang, J. Sun, T. Gardiner, and D. T. Reid, "Broadband conversion in an Yb:KYW-pumped ultrafast optical parametric oscillator with a long nonlinear crystal," Opt. Express **19**, 17127–32 (2011).

Chapter 4 - Locking the CEO frequencies of all pulses in a synchronously-pumped femtosecond optical parametric oscillator to zero Hz

4.1 Introduction

This chapter describes the development and implementation of a formal approach concerning how to lock to zero the carrier-envelope offset (CEO) frequencies of a pump laser and a synchronously-pumped femtosecond optical parametric oscillator (OPO). The concept of the CEO frequency is introduced, and some discussion is given to its detection and multiple methods for its control. Next, the electronic locking scheme necessary for CEO frequency control is outlined. The final part of this chapter discusses the motivation for, requirements and realization of a broadband zero-offset frequency comb, a critical prerequisite for broadband coherent pulse synthesis.

4.1.1 The carrier-envelope offset frequency

A pulse generated by a mode-locked laser can generally be described as a time-dependent electric field composed of a slowly varying envelope and a faster underlying carrier wave. In many areas of optics the carrier-envelope phase can be ignored, however for laser sources producing mode-locked bandwidths that can support few-cycle pulses [1,2] the relative phase between the carrier and the envelope becomes important. As the pulse makes one round trip through a dispersive laser cavity the relative position between the peak of the envelope and the carrier wave changes due to the difference in the group and phase velocities of the pulse. This carrier-envelope phase change is denoted by $\Delta\varphi_{CEP}$. In the time domain this is referred to as the carrier-envelope phase slip (CEPS), and is illustrated in Figure 4.1 (a).

If the pulses from a mode-locked laser are considered in the frequency domain then the output spectrum can be viewed as a comb of frequencies each separated from their nearest neighbour by the pulse repetition frequency. Each comb line is not necessarily an integer multiple of the repetition frequency; generally, each comb frequency has an offset f_{CEO}, again due to the intracavity difference in the phase and group velocities. This offset is the frequency-domain equivalent of the CEPS, known as the carrier-envelope offset (CEO frequency) and is given by

$$f_{CEO} = \frac{\varphi_{CEP}}{2\pi} f_{REP}. \tag{4.1}$$

This concept is illustrated in Figure 4.1 (b).

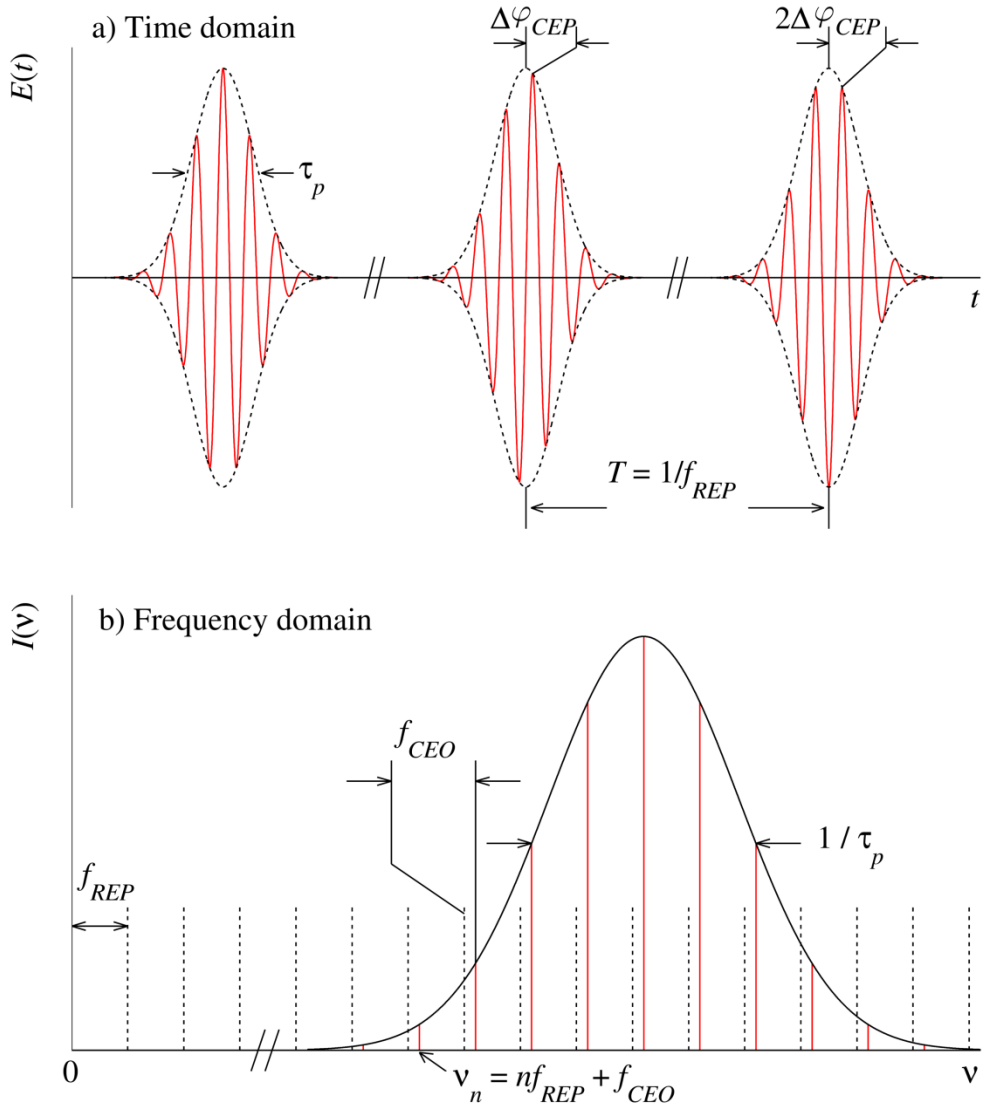

Figure 4.1. Illustration of the evolution of the carrier-envelope phase of a pulse train in the time and frequency domains.

4.1.2 Detecting the CEO frequency

In order to control the CEO frequency of an oscillator it is first necessary to measure it. The standard method of detecting this frequency is known as *f*-to-2*f* self-referencing [3]. This technique requires the frequency bandwidth of the pulses exiting the oscillator to be octave spanning; that is, the oscillator must produce frequencies *f* and 2*f* which have the same CEO frequency. An octave spanning bandwidth can be achieved from a laser source with sufficient gain bandwidth, such as a Ti:sapphire oscillator [1], or by external spectral broadening in a nonlinear medium, typically a

hollow-core or photonic-crystal fiber (PCF), which often leads to supercontinuum generation. When employing the second method it is important to maintain the phase relationship between the lower and higher frequencies generated during broadening. Limiting the supercontinuum to lower-order soliton generation ensures a fixed phase-relationship, which maintains the integrity of the CEO frequency at both ends of the spectrum [4].

The principle behind *f*-to-2*f* self-referencing is illustrated in Figure 4.2 and is detailed here. Frequencies ν_1 at the low-frequency (red) end and ν_2 at the high-frequency (blue) end of an octave-spanning bandwidth are given by

$$\nu_1 = f_{CEO} + nf_{REP} \tag{4.2}$$

and

$$\nu_2 = f_{CEO} + 2nf_{REP} \tag{4.3}$$

respectively, where *n* is an integer. If ν_1 is frequency-doubled though second-harmonic generation in a nonlinear crystal then the resultant frequency will be

$$2\nu_1 = 2f_{CEO} + 2nf_{REP}. \tag{4.4}$$

Two frequencies now exist at $2nf_{REP}$ with different CEO frequencies. Heterodyne mixing between these two frequencies on an avalanche photodiode (APD) or fast silicon photodiode will yield a beat frequency at f_{CEO}. In practice many thousands of modes are frequency-doubled, leading to a beat note at

$$f_{BEAT} = f_{CEO} + pf_{REP} \tag{4.5}$$

where *p* is a positive integer. Frequency bandwidth filtering and the bandwidth of the photodiode limit how big *p* can be, and will alter the signal-to-noise level of the beat note. The detected f_{CEO} beat will lie in the RF domain and will have a value somewhere between 0 Hz and $\frac{f_{REP}}{2}$. An example can be seen in Figure 4.4.

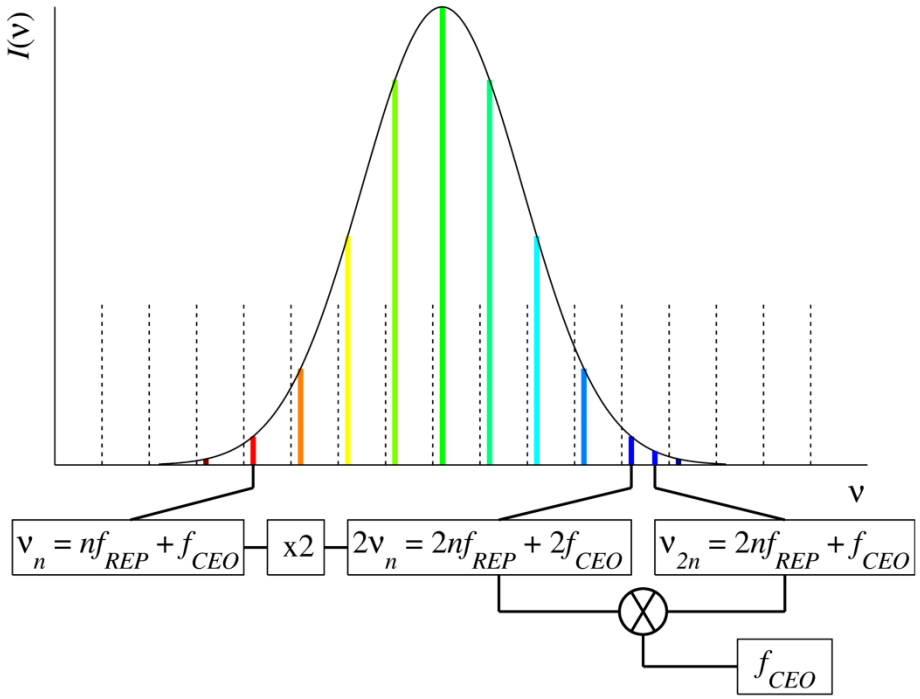

Figure 4.2. Schematic of the *f*-to-2*f* referencing technique using an octave-spanning coherent supercontinuum.

The *f*-to-*2f* interferometer is not the only method by which the CEO frequency of a laser can be detected. Other self-referencing schemes such as 2*f*-to-3*f* can be employed if the laser pulses lack sufficient peak power to generate an octave-spanning supercontinuum [5]. In the case of the femtosecond OPO one can exploit the fact that such OPOs typically generate multiple visible wavelengths through non-phasematched second harmonic generation (SHG) and sum-frequency mixing (SFM) processes in the OPO crystal [6]. These visible OPO outputs can then be heterodyned against visible wavelengths generated from a pump supercontinuum source to obtain a CEO beat frequency, as discussed in Section 4.2.

4.1.3 Methods of controlling the CEO frequency

Once the CEO frequency of the laser oscillator has been detected, a number of methods exist that can be employed to control and lock it. These methods will now be discussed. The electronics required to convert the detected beat frequency into a signal that can be used for locking will be discussed in Sub-section 4.1.4.

For many years the traditional method of controlling the CEO frequency of a laser source relied on the manipulation of the intracavity laser power through modulation of the pump power [7–9]. An acousto-optic modulator (AOM) placed in the beam of a

laser pump source (such as the 0.532 µm pump source of a Ti:sapphire oscillator) was used to alter the phase of the oscillator pulses through changes in the non-linear refractive index of the laser gain medium. This technique maintains the pulse duration and beam quality of the laser, but is limited to DPSS-pumped systems. Pump power is also lost due to diffraction effects in the AOM. In the case of a laser directly pumped by a diode source the feedback signal can instead be used to modulate the driving current of the diode [10]. This enables CEO frequency control of diode-pumped, SESAM mode-locked systems.

A relatively new method known as feed-forward uses an AOM to change the CEO frequency of the laser downstream of the resonator itself [11–13]. This technique first requires the CEO frequency to be measured, often by using a technique similar to that described in Sub-section 4.1.2. This signal is then amplified and used to drive an AOM. The AOM is aligned for Bragg diffraction at the -1^{st} order, and so the emerging laser pulses have a CEO frequency of 0 Hz. This method is robust as it requires no feedback loop or locking electronics and can be used with any laser source. Typical diffraction efficiencies are 60-70%, with the remaining zero-order output used for CEO frequency detection. The frequency-shifted pulses are both spatially and temporally chirped.

The two methods discussed require high-bandwidth frequency modulation of the laser beam, however both can be intrinsically inefficient as laser power is lost due to diffraction effects. Another approach is to manipulate the intracavity dispersion, directly controlling the CEO frequency of the emerging mode-locked pulses. Two relevant examples are discussed below.

By definition the carrier-envelope offset phase arises from a difference between the intracavity group velocity and phase velocity. Modulation of the intracavity dispersion allows control of this CEO phase [14]. A common example is a resonator with an intracavity prism pair for dispersion compensation, as described in Chapter 2. After the beam passes through the prism pair it is spatially dispersed across the face of a mirror before being reflected back. Detection of the CEO frequency occurs as described in Sub-section 4.1.2, and the resulting error signal is used to modulate the pointing of the mirror through the use of a piezoelectric transducer (PZT) [15]. The tilt of the end mirror must be kept small compared to the end mirror angular aperture in order to prevent intracavity power fluctuations. Mirror excursions are therefore kept to the microradian range. This method allows for maximum efficiency over a control

bandwidth comparable to that of an acousto-optic method, provided that mirror inertia is taken into account when selecting the PZT and mirror. There are two drawbacks however; this technique cannot be applied to a mirror-dispersion-compensated oscillator as the intracavity dispersion is essentially 'fixed' by the mirror coatings and is insensitive to beam pointing. The use of intracavity prisms also leads to a faster-varying CEO frequency than would be measured from a mirror-based system, as minute changes in the intracavity pointing produce a large phase change due to the higher refractive index of the prisms [16–19].

A second dispersion-based approach to CEO frequency control exists in the case of synchronously-pumped OPOs. The intracavity pulses in a femtosecond OPO have lower peak powers than in a laser oscillator, and so the contribution to the nonlinear dispersion of the crystal is much lower in an OPO. A much larger pump power modulation would be required to alter the CEO frequency of an femtosecond OPO than is needed for a Ti:sapphire laser, which is typically of the order of 10^{-3}, and so pump power modulation is not a viable option. For OPOs that produce multiple outputs from SHG and SFM processes, an AOM-based feed-forward approach is also unsuitable as each output has its own unique CEO frequency, as discussed in Sub-section 4.2.1, and so a dispersion-modulation approach must be employed instead.

In Chapter 3 the cavity length tuning behaviour of an OPO was examined. The same derivation will now be carried out for the CEO phase of the OPO. In terms of the absolute phase of an OPO pulse, the CEO phase is given by

$$\phi_{CEO} = \phi(\omega) - \omega \frac{\partial \phi(\omega)}{\partial \omega}. \tag{4.6}$$

The change in CEO phase with frequency is therefore given by

$$\frac{\partial \phi_{CEO}}{\partial \omega} = -\omega \frac{\partial^2 \phi}{\partial \omega^2}, \tag{4.7}$$

which can be compared with Equation (3.9) to give

$$\frac{\partial \phi_{CEO}}{\partial \omega} = \frac{\omega}{c} \frac{dL}{d\omega} = k \frac{\partial L}{\partial \omega}. \tag{4.8}$$

This equation can be rearranged to provide

$$\frac{\partial \phi_{CEO}}{\partial \omega} = \frac{\left(\frac{\omega}{c}\right)}{\left(\frac{\partial \omega}{\partial L}\right)} \tag{4.9}$$

which in turn can be rearranged to show that

$$\frac{\partial \phi_{CEO}}{\partial L} = \frac{2\pi}{\lambda}. \tag{4.10}$$

Comparing Equations (3.10) and (4.9) it can be shown that the OPO CEO phase will tune much quicker with cavity length than the centre frequency. This allows the CEO phase to be controlled by making small changes to the cavity length without affecting the centre frequency of the resonant pulse. Equation (4.10) shows that a cavity round-trip change of λ will provide a CEPS change of 2π. Cavity length changes of this magnitude can easily be achieved by mounting a cavity mirror on a piezoelectric actuator, which can be controlled through a feedback loop for signal or idler CEO phase control [20–23].

4.1.4 An electronic feedback loop for CEO frequency control

The feedback circuitry required to successfully control the CEO frequency of a laser is non-trivial, and contains a number of key components which will be described below. For the purposes of this discussion it will be assumed that the system used for CEO frequency detection is an APD in some form of nonlinear interferometer, as previously discussed in Sub-section 4.1.2. The general outline of the electronic feedback loop is shown in Figure 4.3.

4.1.4.1 CEO frequency detection and amplification

The APDs employed during the course of this work were Hamamatsu C5331-11s, which have peak sensitivity at visible wavelengths and a high-band cut-off frequency of 100 MHz, making them ideal for the experiments described in Section 4.2. An APD is necessary due to its ability to detect and amplify extremely weak light signals. The output from the APD (Figure 4.4 (a)) is an analogue RF signal containing strong frequencies at f_{REP} and $f_{REP} \pm f_{CEO}$. The absolute strength of the CEO signal is typically less than -35 dBm and so requires amplification; a minimum of -20 dBm signal strength is required for locking, with at least a 30-dB signal-to-noise ratio, typically 45-dB. An amplifier (Mini-Circuits, model ZFL-500LN+) was used to supply

a minimum of 25-dB gain. A 50-MHz low-pass filter was used to attenuate the strong f_{REP} signal and prevent damage to the amplifier (Figure 4.4 (b) and (c)). A final 32-MHz low-pass filter was placed after the amplifier to provide a clean CEO signal (Figure 4.4 (d)). The desired cut-off frequency of the low-pass filter depends on the frequency to which the CEO frequency will be referenced.

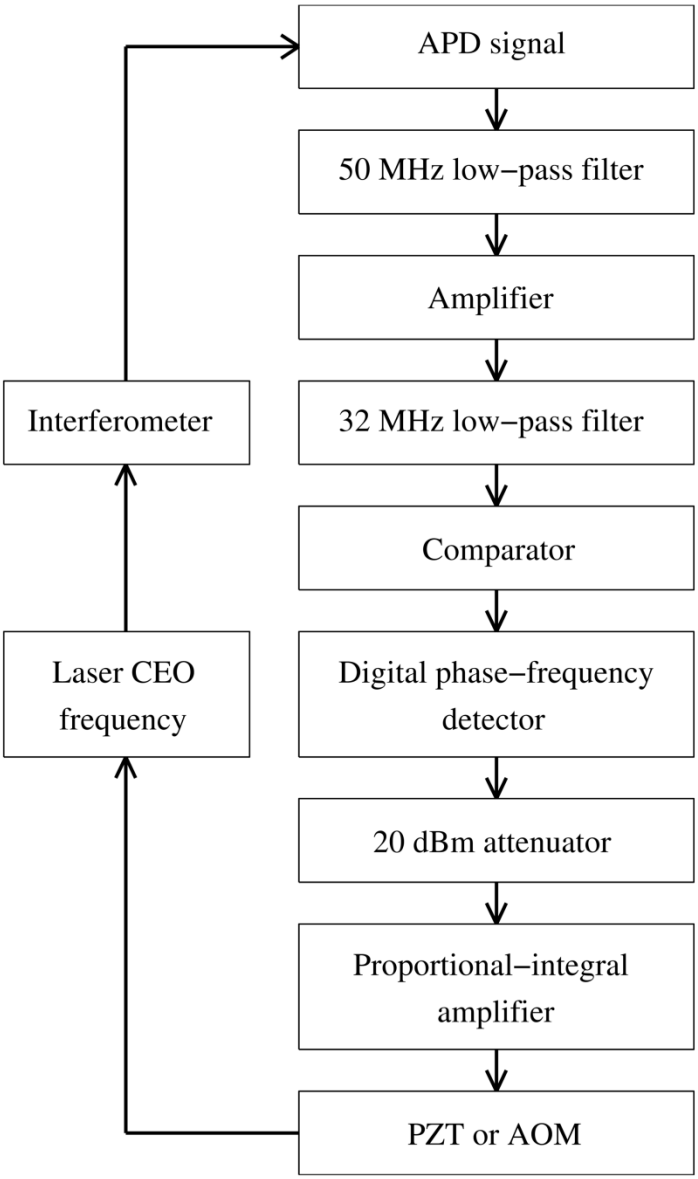

Figure 4.3. Outline of an electronic feedback loop for CEO frequency control.

The low-pass RF signal is passed into a comparator (Pulse Research Lab, PRL-350TTL) to generate a low-noise TTL signal (Figure 4.5). The output voltage from the comparator is either 0 or 2.7 V. The TTL CEO signal, along with a similarly formatted reference signal, was passed into a digital phase-frequency detector to generate an error signal for CEO frequency locking.

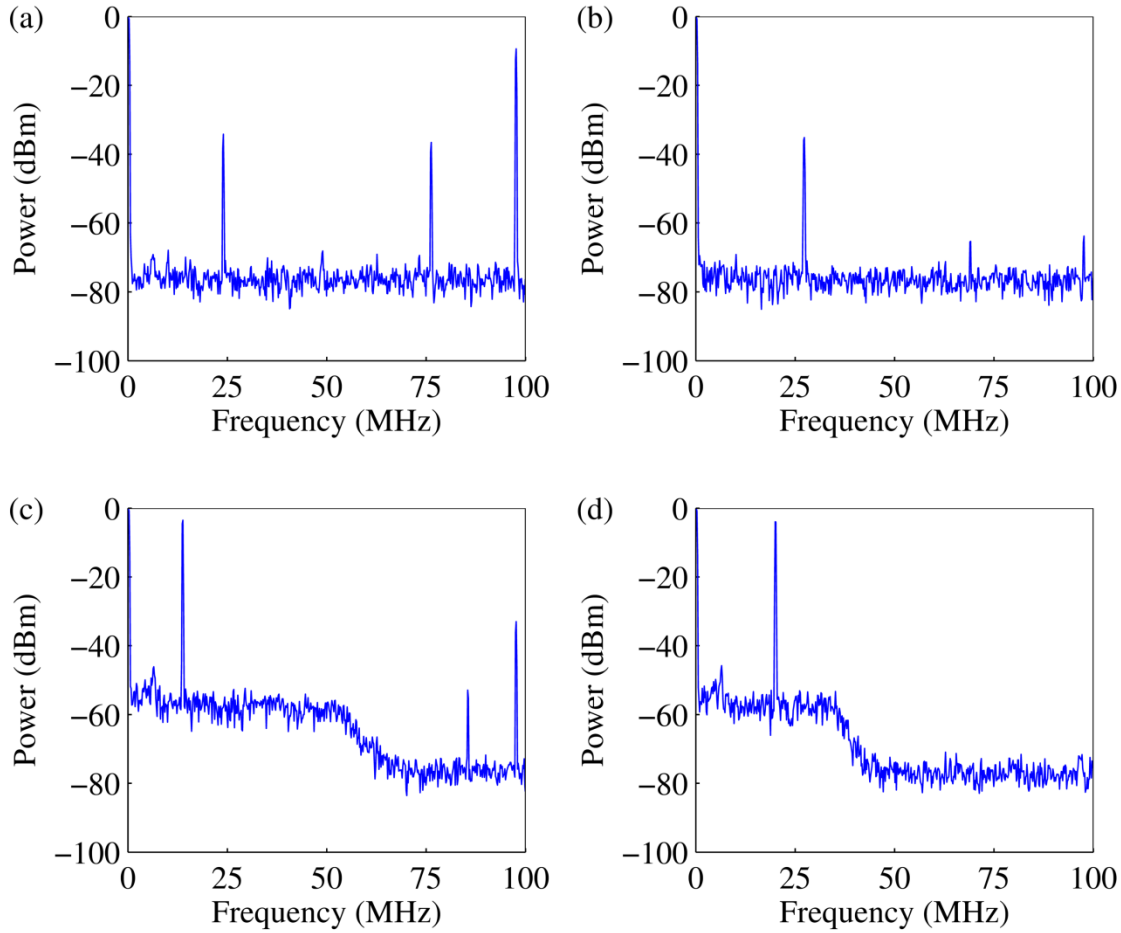

Figure 4.4. CEO beat signals at different stages of formatting. (a) The raw beat signal as detected by the APD; (b) after a 50 MHz low-pass filter; (c) after amplification; (d) after a 32 MHz low-pass filter. Figures (c) and (d) were recorded with 10 dB attenuation to protect the RF spectrum analyser.

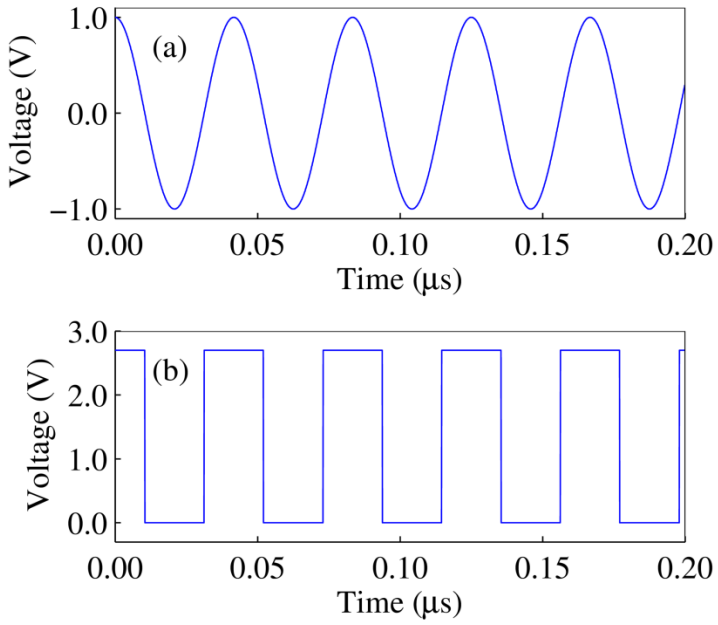

Figure 4.5. (a) Amplified, filtered 24 MHz signal from the APD; (b) TTL output from the comparator.

4.1.4.2 Phase-frequency detector

A digital phase-frequency detector (PFD) compares two TTL input signals and provides the phase/frequency difference between them [24]. The first input is the formatted CEO frequency signal, while the second input is the reference signal to which the CEO frequency is to be locked. There are two standard choices for a reference signal. The first is a fixed reference such as a stabilized 10-MHz Rubidium clock or the output from an RF signal generator. The second choice of reference source is a subharmonic of the pulse repetition rate, such as $f_{REP}/4$, as used in this work.

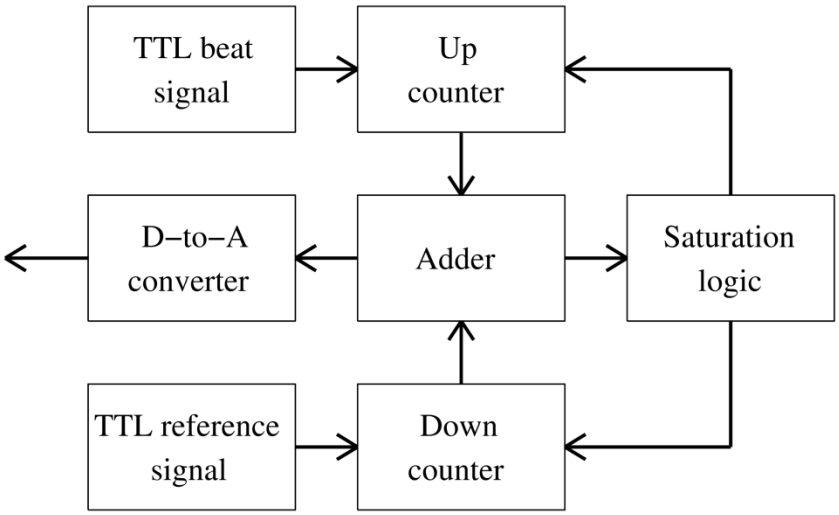

Figure 4.6. Block diagram of the digital phase-frequency detector used in the experiments carried out in this thesis. Adapted from [24].

A block diagram of the PFD used in this thesis is shown in Figure 4.6, and the circuit diagram of the PFD is shown in Figure 4.7. The TTL frequency input from the beat (reference) signal is passed to D-type flip-flop A (B), then divided down in 4-bit synchronous UP (DOWN) counter C (D). The two inputs are added together by 4-bit adder E. If the frequency of the UP input is slightly faster (slower) than the DOWN input then the UP counter will count faster (slower) than the DOWN counter, resulting in an output at the adder that is slowly incrementing (decrementing). When the adder reaches 15 or 0 the output becomes saturated, causing additional logic components (F) to stop the appropriate flip-flop circuit from triggering. When the UP and DOWN inputs have the same frequency but different phases the adder output is constant and reflects the phase difference between the signals. The phase difference is converted to an analogue signal using a simple D-to-A converter (G).

The circuits used in these experiments had a linear range of ±32π, about ±100 radians. The reference frequency input range is determined by the D-to-A converter, and is limited in our case to 5-32 MHz. The analogue output from the PFD is used as the input to a proportional-integral amplifier for CEO control. The DC offset of the output voltage could be adjusted to ensure that two identical frequency inputs would produce an output of 0 V.

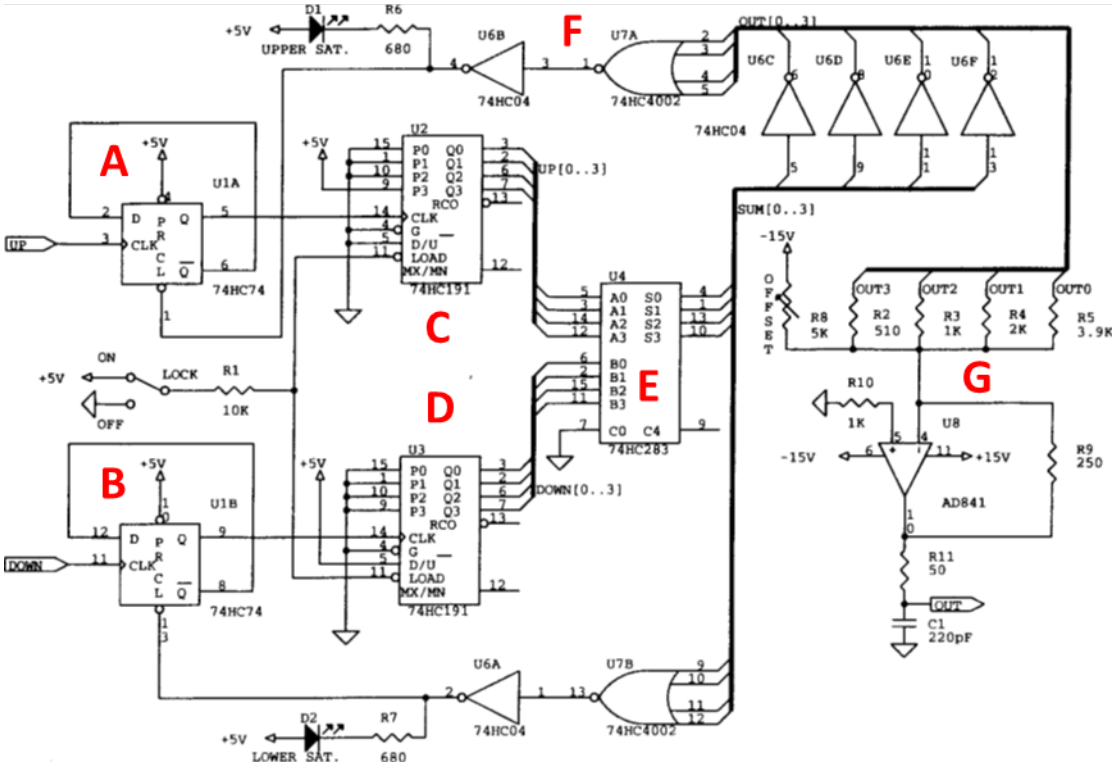

Figure 4.7. Circuit diagram of the digital PFD used in this thesis. See text for a full description. Adapted from [24].

4.1.4.3 Proportional-integral amplifier

The analogue signal from the PFD is sent to a proportional-integral (PI) amplifier to form a stable locking loop. For locking schemes employing a PZT stack in the laser cavity the stack can normally be driven directly by the PI amplifier. For schemes using an AOM to modulate the laser pump power, the PI amplifier output is used to modulate the driving signal of the AOM driver. Correct adjustment of the PI amplifier is one of the most important and difficult sections of the locking system. For the purposes of this discussion the PI amplifier is assumed to be the New Focus LB1005 Servo Controller shown in Figure 4.8, the same model used throughout this work.

Figure 4.8. Front panel of the New Focus LB1005 PI amplifier. 1, input ports; 2, input offset control; 3, PI corner frequency control; 4, proportional gain control; 5, gain regime control; 6, output port; 7, low frequency gain control; 8, output offset control.

The PI amplifier must be calibrated before use, particularly if it will be used to drive a PZT. While monitoring the output voltage (6) with no input (1), the output offset control (8) is turned to its minimum setting and a fine control on the rear of the device adjusted to set the minimum voltage to 0 V. Similarly the output control is adjusted to its maximum setting and the voltage set to 10 V. These voltage rail settings will provide the PZT with sufficient drive to capture and lock the CEO frequency, but will also prevent the PZT from becoming damaged by being driven with a negative voltage.

After calibration, the output from the PFD can be connected to input A of the PI amplifier. A 20-dBm attenuator is placed before the PI amplifier input to reduce the voltage level from the PFD, providing greater sensitivity for the locking loop. The output port is connected to either a PZT or an AOM driver. The input offset (2) of the PI amplifier should be turned to 5.0, a setting corresponding to zero offset. The acquire switch (5) should be set to LFGL (low-frequency gain limit) for initial locking; this limits the DC gain and prevents integral gain domination at lower frequencies, making it easier to acquire a lock. Once locking is achieved the switch can be set to Lock On for longer-term locking stability.

While monitoring the CEO frequency on an RF spectrum analyser, voltage should be applied to a course CEO adjustment method (often a long-travel PZT in the laser cavity) to bring the CEO frequency into the capture range of the locking system. With an initial corner frequency (3) setting of 3 kHz, a proportional gain (4) of 4.0 and a low-frequency gain limit of 50 dB it should be possible to achieve coarse locking. If the

CEO frequency does not stabilize but instead jumps over the target frequency then the input polarity must be changed by switching to the other input port, or by sweeping the CEO beat signal with the course adjustment until it approaches the locking frequency from the other direction.

Increasing the proportional gain should improve the locking stability, which can be measured by monitoring the bandwidth of the locked CEO frequency on a high-resolution RF spectrum analyser. The output from the PFD can also be monitored, moving from a noisy signal to a smooth line as the locking is stabilised. If the proportional gain or the corner frequency is too high then the output from the PFD will oscillate. This signal will be fed back into the locking PZT, causing the locking signal on the RF spectrum analyser to broaden significantly. Careful optimisation of the gain parameters will increase the stability of the locking signal. Over time the CEO frequency will drift out of the range of the locking PZT. It can be brought back into locking range by adjusting the coarse PZT.

4.2 Zero-offset CEO frequency stabilisation of a Ti:sapphire laser and a femtosecond OPO

The optical and electronic techniques described in Section 4.1 were used to detect and stabilize the CEO frequencies of the pump and OPO described in Chapter 3 to 0 Hz, a process known as zero-offset locking. The layout of the locking scheme and relevant results will now be discussed.

4.2.1 Zero-offset locking

As described in Chapter 1, the fundamental prerequisites for coherent pulse synthesis between two or more pulses are that the pulses share (*i*) a common repetition frequency and (*ii*) a common carrier-envelope offset frequency. The synchronous nature of the femtosecond OPO immediately satisfies condition (*i*), providing intrinsic low-timing-jitter synchronisation between the OPO and its pump source. The multiple nonlinear mixing frequencies generated in the OPO, as well as their parent pump, signal and idler pulses, all share the same repetition rate, providing an opportunity for coherent synthesis between the individual pulses.

The second prerequisite for pulse synthesis is simple to satisfy when the combined pulses share a common frequency, however when combining multiple pulses of different frequencies, a common CEO frequency is unlikely to be shared by more than two pulses. In general terms the OPO produces pump (*p*), signal (*s*) and idler (*i*) frequency combs which can be described by

$$f_p = kf_{REP} + f_{CEO}^p \tag{4.11}$$

$$f_s = lf_{REP} + f_{CEO}^s \tag{4.12}$$

$$f_i = mf_{REP} + f_{CEO}^i, \tag{4.13}$$

where *k*, *l* and *m* are integers. Each nonlinear frequency-mixing output is in itself a new comb, which can be generally expressed as

$$f_{NL} = nf_{REP} + qf_{CEO}^p + rf_{CEO}^s \tag{4.14}$$

where *n*, *q* and *r* are integers. The idler CEO frequency does not appear explicitly in this equation as it can always be eliminated using the relation

$$f_{CEO}^{p} = f_{CEO}^{s} + f_{CEO}^{i}. \tag{4.15}$$

Synthesizing a new pulse sequence from two or more nonlinear mixing outputs requires that the participating frequency combs share a common CEO frequency, implying that

$$qf_{CEO}^{p} + rf_{CEO}^{s} = q' f_{CEO}^{p} + r' f_{CEO}^{s} = q'' f_{CEO}^{p} + r'' f_{CEO}^{s} \text{ etc,} \tag{4.16}$$

which is only generally possible when

$$f_{CEO}^{p} = f_{CEO}^{s} = f_{CEO}^{i} = 0. \tag{4.17}$$

For this reason it is necessary to lock the CEO frequencies of the pump and either the signal or idler to 0 Hz, the other automatically locking to 0 Hz as a result of Equation (4.15). In terms of synthesis, zero-offset locking gives identical electric field structure from pulse to pulse.

The electronic referencing techniques discussed in Sub-section 4.1.4 describe two main references to which the CEO frequency can be locked. The first, an external reference, is of no use when locking to 0 Hz as there is obviously no reference signal that can be generated with this frequency. The second standard reference is a fractional harmonic of the source repetition frequency; however it is insufficient for us to lock both the pump and idler CEO frequencies to this value, as Equation (4.14) shows that each SHG and SFM output will have a different CEO frequency.

Locking to 0 Hz without using a feed-forward mechanism can be achieved by introducing a known frequency into one arm of the nonlinear interferometer used for CEO frequency detection. If we take the case of a standard *f*-to-2*f* interferometer, it was shown in Sub-section 4.1.2 that the CEO frequency can be detected from the equation

$$2\nu_n - \nu_{2n} = \left(2nf_{REP} + 2f_{CEO}\right) - \left(2nf_{REP} + f_{CEO}\right) = f_{CEO}. \tag{4.18}$$

If we introduce a known RF frequency, f_{shift}, into one arm of the interferometer, say the one containing the fundamental frequency ν_{2n}, then the detected beat signal becomes

$$2\nu_n - \nu_{2n} = \left(2nf_{REP} + 2f_{CEO}\right) - \left(2nf_{REP} + f_{CEO} + f_{shift}\right) = f_{CEO} + f_{shift}. \tag{4.19}$$

If we then choose $-f_{shift}$ as our locking frequency then the CEO frequency will be locked to zero through

$$2\nu_n - \nu_{2n} = f_{CEO} + f_{shift} = -f_{shift} + f_{shift} = 0. \quad (4.20)$$

In practice frequency shifting is achieved by placing an AOM in one arm of the interferometer and driving it at f_{shift}, such that the first-order diffracted beam from the AOM carries the additional frequency shift. This diffracted beam is then interfered in the nonlinear interferometer in the standard way. By detecting and locking two CEO frequencies from the pump and OPO in this manner it is possible to achieve broadband zero-offset CEO locking.

4.2.2 A nonlinear interferometer for zero-offset CEO locking

Sub-section 4.1.2 described a number of methods of detecting the CEO frequency of a mode-locked laser using an octave-spanning supercontinuum, however in this work it is necessary to detect two independent CEO frequencies that relate to the pump and OPO, and also frequency-shift one arm of the interferometer to achieve zero-offset locking. Two primary nonlinear interferometer configurations for zero-offset locking will now be discussed.

4.2.2.1 Interferometer 1 - an f-to-2f variant

The schematic for the first interferometer design is shown in Figure 4.9, and is based around a traditional *f*-to-2*f* interferometer. 150 mW of pump power was launched into a 30 cm length of PCF (*NKT Photonics NL-2.0-750)* in order to generate an octave spanning supercontinuum. A cold mirror reflected the visible wavelengths of the supercontinuum while transmitting the infrared solitons. Infrared light at 1.060 µm was frequency doubled in a 5-mm-long KDP crystal, with the resultant 0.530 µm light transmitted through a polarising beam cube. The fundamental 0.530 µm light was frequency shifted using an AOM (*IntraAction ASM-803B47*) driven at $3f_{REP}/4$. The reasoning behind this frequency choice will be discussed in Sub-section 4.2.3.

The first-order frequency shifted 0.532 µm beam was sent through a variable-delay arm before being recombined with the second-harmonic 0.530 µm beam after a polarising beam cube. The spatial overlap between the two beams was confirmed over a long distance before the beams were passed through a 0.530 µm interference filter with a 10-nm bandwidth and focused onto an APD. An angled polariser was placed before the lens to balance the power between the two arms of the interferometer and to resolve a common polarisation, as required for observing interference.

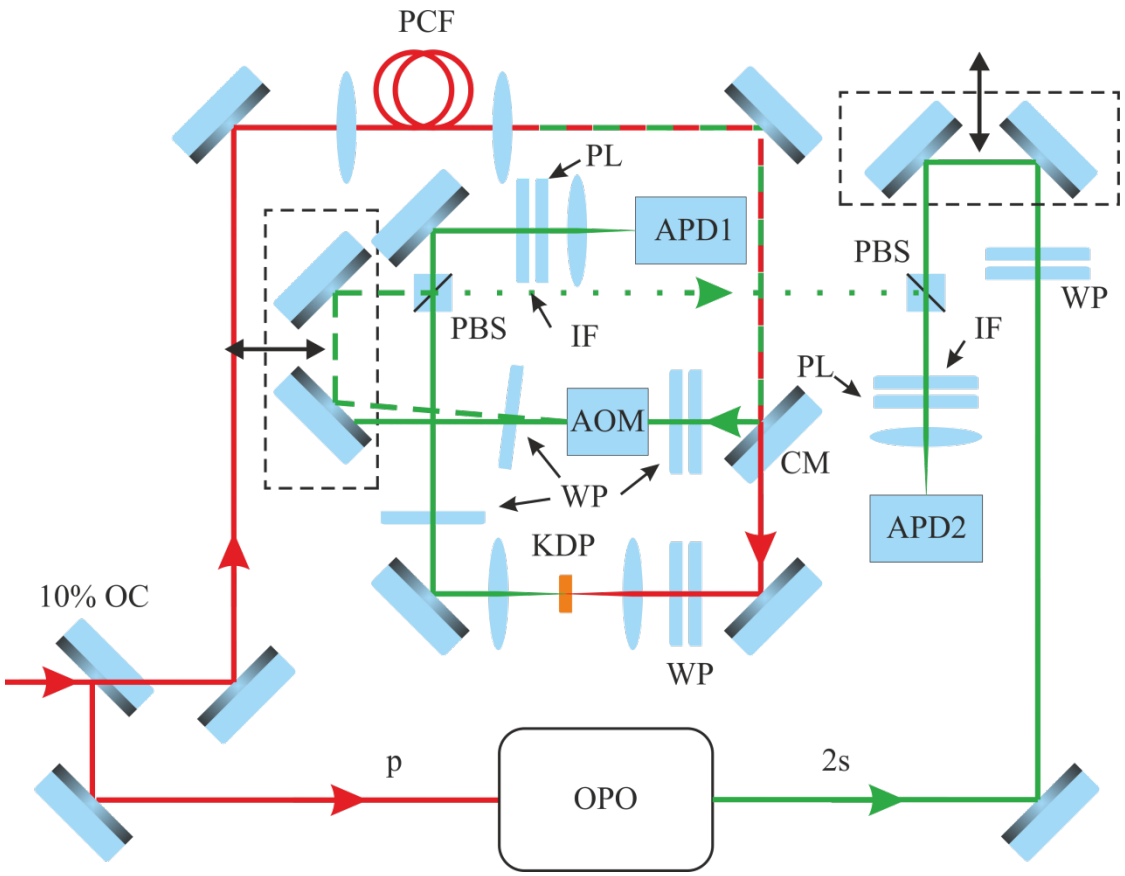

Figure 4.9. Schematic of interferometer 1 – see text for a detailed description. AOM, acousto-optic modulator; APD, avalanche photodiode; CM, cold mirror; IF; interference filter; OC, output coupler; PBS, polarising beam splitter; PCF, photonic crystal fiber; PL, polariser; WP, wave plate.

The output from the APD was monitored on an RF spectrum analyser while the variable-delay arm was scanned through a region of zero path difference. When the path lengths were equal a weak CEO beat of 25 dBm signal-to-noise was observed. While the signal strength was poor, it was believed that this could be improved by altering the beam overlap and by adjusting the alignment into the PCF.

This variation on the *f*-to-2*f* interferometer would allow the pump CEO frequency to be locked to zero, however it was also necessary to detect a second CEO frequency that related to the OPO. This second CEO beat signal must also be frequency-shifted, and so residual light from the first-order 0.530 µm beam that passed through the polarizing beam cube was interfered with 0.530 µm SHG signal light from the OPO. The path difference between the pump and OPO beam were measured from the 10% output coupler in Figure 4.9. Fortuitously the path lengths could be made equal after only a small addition to the OPO beam path. As with the *f*-to-2*f* interferometer described above, the two beams were overlapped spatially, passed through a 0.530 µm interference filter and polariser, and then focused onto an APD. Unfortunately the

signal strength from this second nonlinear interferometer was poorer than the first, primarily a result of using the residual light from the frequency shifted pump beam.

This design would allow the pump and OPO CEO frequencies to be measured independently, however in practice there was insufficient power in the frequency shifted 0.530 μm light to achieve strong beat signals on the APDs. With this in mind a new design was implemented.

4.2.2.2 Interferometer 2 – independent pump supercontinua

The second nonlinear interferometer design is shown in Figure 4.10. 200 mW of pump light was compressed using a GTI and used to pump a pair of independent pump supercontinua, one with strong wavelength components in the green and the other with strong components in the red. In this configuration the AOM was used to frequency shift the visible outputs of the OPO. The pump supercontinuum with the strong green component was interfered with the first-order frequency shifted 0.532 μm light from the OPO using the same interference filter and polariser set up as described above, however there was sufficient power in both arms that a focusing lens before the APD was not required. Similarly, the supercontinuum with the strong red component was interfered with frequency-shifted 0.633 μm light from the OPO, generated by $p+i$ SFM in the OPO crystal.

The signal-to-noise of the CEO frequencies from both the red and green nonlinear interferometers was over 40 dB, a marked improvement to the first design. By employing two separate supercontinua it was possible to adjust the alignment into each PCF independently. A mechanically simpler configuration using only one supercontinuum and a beamsplitter to divide the green and red spectra components is possible, however it lacks flexibility in the ability to tune the output wavelengths of the PCF. A small change in alignment to increase the power at 0.532 μm would often lead to a decrease in power at 0.633 μm, making optimisation of that design a challenge. The use of two PCFs also allows for their output wavelengths to be independently optimised to match almost any combination of SHG and SFM wavelengths, which is advantageous when you consider the tuning range of a broadband femtosecond OPO.

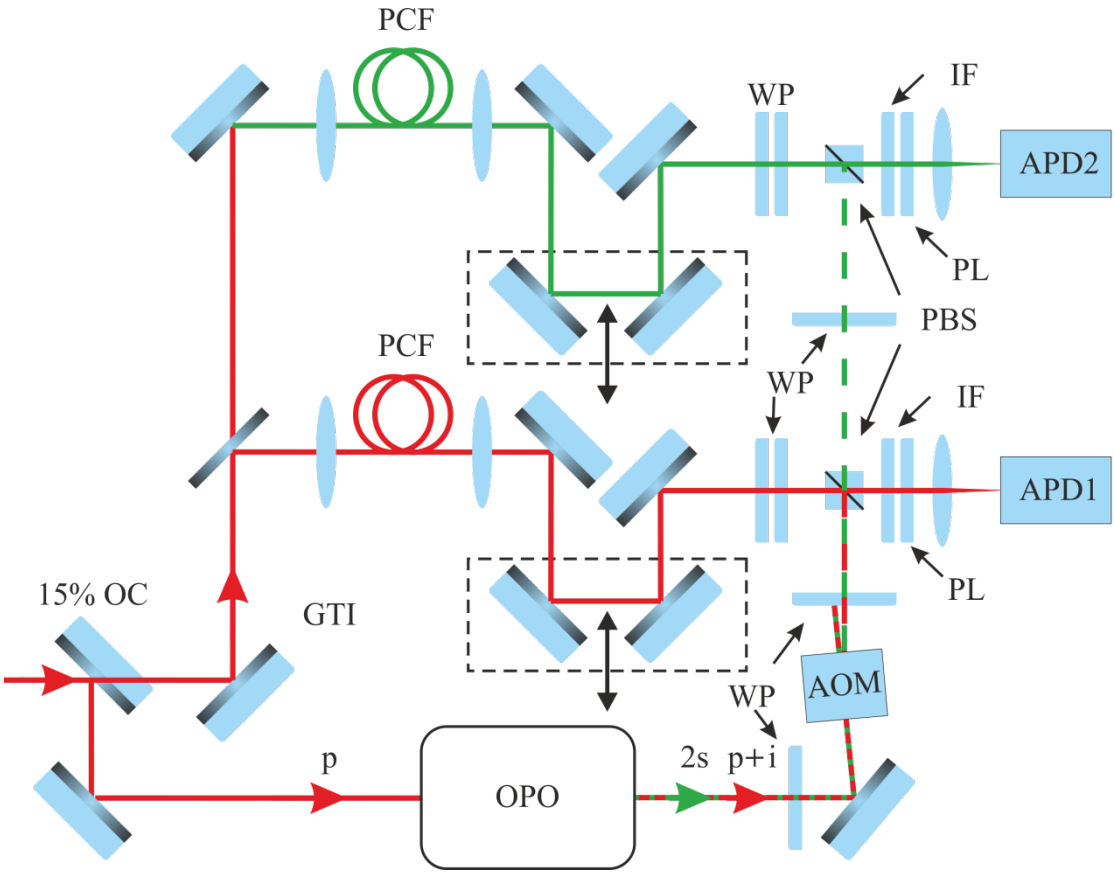

Figure 4.10. Schematic of interferometer 2 – see text for a detailed description. AOM, acousto-optic modulator; APD, avalanche photodiode; IF; interference filter; OC, output coupler; PBS, polarising beam splitter; PCF, photonic crystal fiber; PL, polariser; WP, wave plate.

4.2.3 Zero-offset locking results

As discussed above, zero-offset locking was achieved by frequency-shifting the $p+i$ and 2*s* OPO pulses before they entered the nonlinear interferometers by using an AOM driven at $f_{shift} = 3f_{REP}/4$ (75 MHz). A fast photodiode located in the Ti:sapphire cavity was used to detect the repetition rate; the output from the photodiode was internally processed in the laser control box into a TTL signal. This TTL signal was split into two, and the first signal passed to a divide-by-*n* frequency divider and used to obtain a frequency of $f_{REP}/4$. The second f_{REP} signal was combined with the divided signal in an RF mixer to generate an output frequency of $3f_{REP}/4$. This signal was used as the input to an RF amplifier that drove the AOM. It was not possible to directly drive the AOM at $f_{REP}/4$ (25 MHz) because of its limited radio-frequency acceptance bandwidth. The AOM can be considered to red-shift the $p+i$ and 2*s* modes by

$-f_{REP}/4$, and for this reason the detected CEO beat frequencies were referenced to $f_{REP}/4$.

Detecting a heterodyne beat between a pump supercontinuum and the AOM-shifted SFM and SHG OPO outputs requires that the first-order diffracted beam be used for detection. This beam carries less power than the zero-order beam, and the diffraction efficiency of the AOM can only be optimised across a limited range of wavelengths. For this reason a component of the pump-idler SFM light closest in wavelength to the signal SHG output was chosen for overlap with the pump supercontinuum with the strong red component, as illustrated in Figure 4.11.

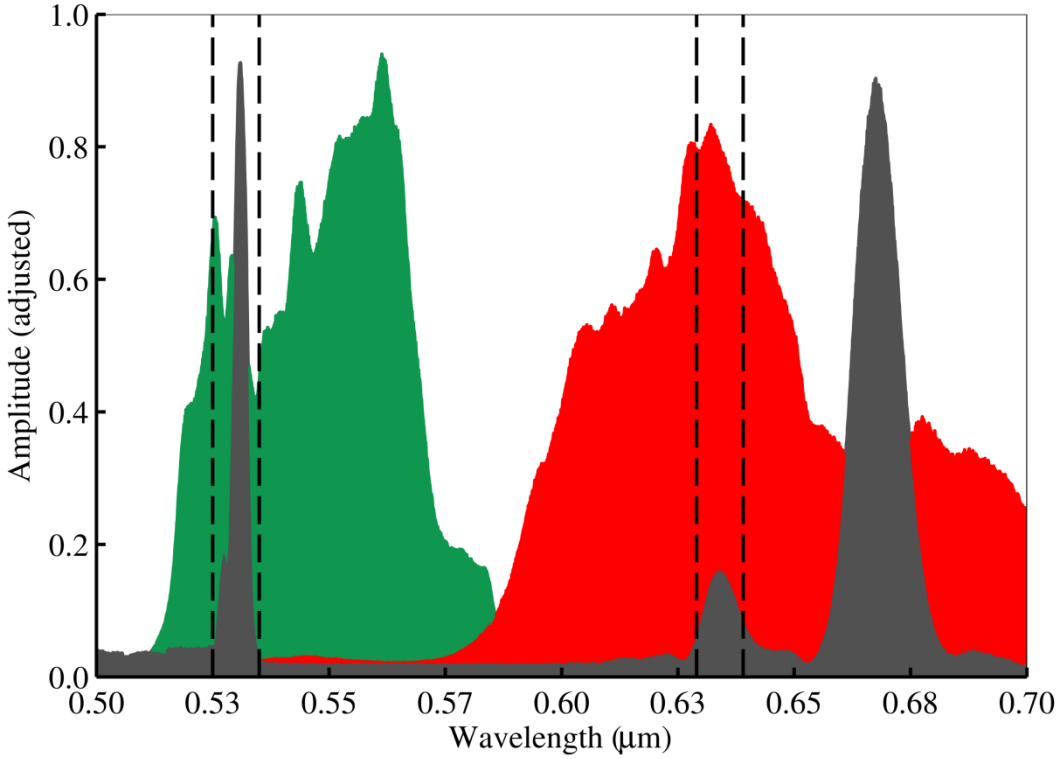

Figure 4.11. Spectral overlap regions between the visible OPO outputs (grey) and the two pump supercontinua (green and red). The 0.532 μm SHG signal beam was overlapped with the supercontinuum with a strong green component, while the 0.633 μm SFM pump + idler beam was overlapped with the supercontinuum with a strong red component. The dashed lines indicate the bandpass interference filter regions used to detect a heterodyne beat.

Heterodyne beating at APD1 between the red pump supercontinuum and the $p+i$ OPO pulses produced an idler CEO beat. Similar beating between the green pump supercontinuum and the SHG signal pulses at APD2 produced a CEO beat containing information about both the pump and OPO CEO frequencies. If the first CEO beat signal has been locked then the second signal will track with the pump CEO frequency.

The locking scheme was implemented as discussed in Sub-section 4.1.4, with the detected CEO frequencies being referenced against the derived $f_{REP}/4$ signal using a pair of PFD circuits. The output from the PFD referenced against APD1 was sent to a PI amplifier that supplied voltage to a PZT attached to a folding mirror in the OPO cavity (*Thorlabs AE0203D04F*; 261 kHz unloaded resonance frequency; 4.6 µm maximum displacement, 150 V maximum drive voltage). The corner frequency was 10 kHz, the LF gain limit was 50 dB and the proportional gain was 5.1. This feedback loop was used to lock the idler CEO frequency. Similarly the output from the PFD referenced against APD2 was sent to a PI amplifier that drove a PZT in the Ti:sapphire pump laser (>500 kHz unloaded resonance frequency) which was mounted on the cavity end-mirror which received spatially dispersed light from the intracavity dispersion-compensating prism pair. The PI amplifier corner frequency was 10 kHz, the LF gain limit was 50 dB and the proportional gain was 7.3. This feedback loop was used to lock the pump CEO frequency.

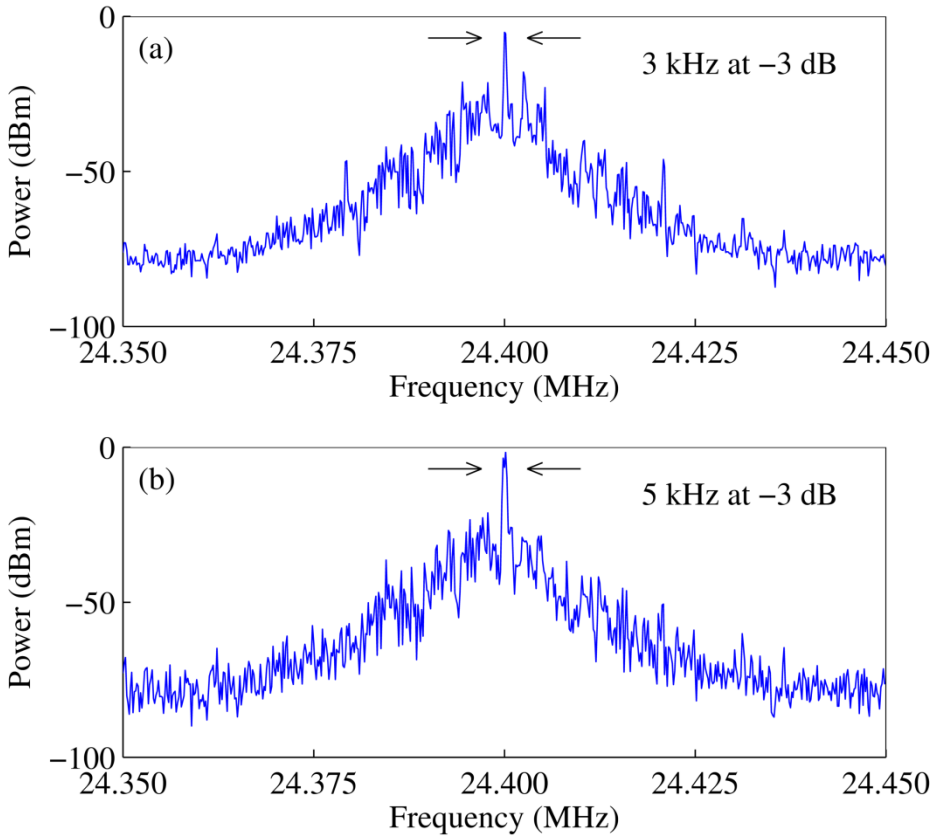

Figure 4.12. RF spectra of the locked CEO frequencies for (a) the idler CEO frequency used to lock the OPO, and (b) the $p + 2s$ CEO frequency used to lock the pump. The span is 100 kHz with a resolution of 300Hz.

As the CEO beat detected at APD2 was a product of CEO frequencies from the pump and OPO it was characteristically unstable, moving at tens of MHz per second. In contrast the idler CEO frequency detected at APD1 moved less than five MHz per second. In order to lock the CEO beat frequency from APD2 it was first necessary to lock the idler CEO frequency. This increased the stability of the beat signal at APD2, allowing the PI amplifier controlling the pump CEO frequency to acquire a lock. The CEO beat signals for both locking loops are shown in Figure 4.12 for the case where both loops are locked. The bandwidth of the idler CEO signal was 3 kHz at the -3 dB point. The $p+2s$ CEO signal was noisier with a 5 kHz bandwidth which is likely due to noise in the idler locking loop feeding into the $p+2s$ locking loop. Characterisation of the locking stability will be discussed in Chapter 5.

4.3 Confirming lock-to-zero

Section 4.2 detailed the mechanism used to implement zero-offset locking of both the pump and OPO, achieving broadband phase coherence across all pulses on the optical bench. However, locking the CEO frequencies alone is not sufficient to truly ensure zero-offset locking, as it is necessary to check for phase coherence between the pulses. The requirement to confirm phase coherence will now be discussed, along with relevant results.

4.3.1 The need for an all-optical confirmation of coherence

When the system is locked, optical heterodyning at the APD in each nonlinear interferometer produces a frequency at f_{REP} with sidebands at $\pm f_{REP}/4$, shown in Figure 4.13. Supplying voltage to a long-travel PZT located in each cavity allows the monitored CEO frequency to be shifted by multiples of 2π radians. Consequently, either beat frequency can be locked to $f_{REP}/4$ or $3f_{REP}/4$, with no electronic means of distinguishing between the two scenarios. This gives a total of four possible locking combinations, only one of which achieves the desired condition of $f_{CEO}^{p} = f_{CEO}^{s} = f_{CEO}^{i} = 0\,\text{Hz}$. Because of the inherent ambiguity of purely electronic detection, confirming lock-to-zero (and therefore broadband phase coherence) requires optical confirmation through either a spectral or temporal interferometric measurement.

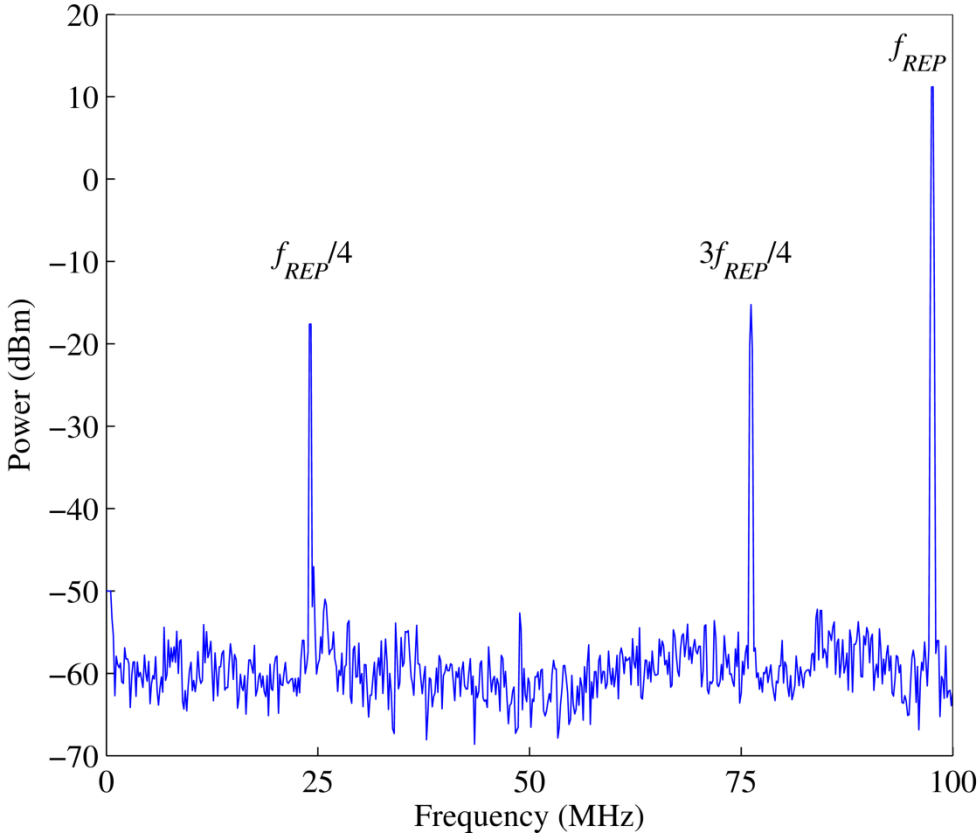

Figure 4.13. Raw RF signal from by APD1. The detected CEO frequency has been locked to $f_{REP}/4$, and so sidebands are detected at both $f_{REP}/4$ and $3f_{REP}/4$.

4.3.2 Confirmation by time-domain interferometry

An interferometer was constructed in which light from the pump supercontinuum containing a strong 0.530 µm component and a weaker 0.633 µm component was interfered with visible SFM and SHG light exiting an OPO folding mirror (Figure 4.14). A temporal interferometry experiment was implemented in which the OPO beam path was modulated using a PZT stage with a frequency of 1.4 Hz and a displacement of 400 µm. The beams were combined and passed through an appropriate interference filter before being detected on a silicon photodiode (*Thorlabs DET10A/M*). A temporal approach was preferred over a spectral approach due to the poor spatial overlap of the interfering beams and low measurement resolution of the spectrometer.

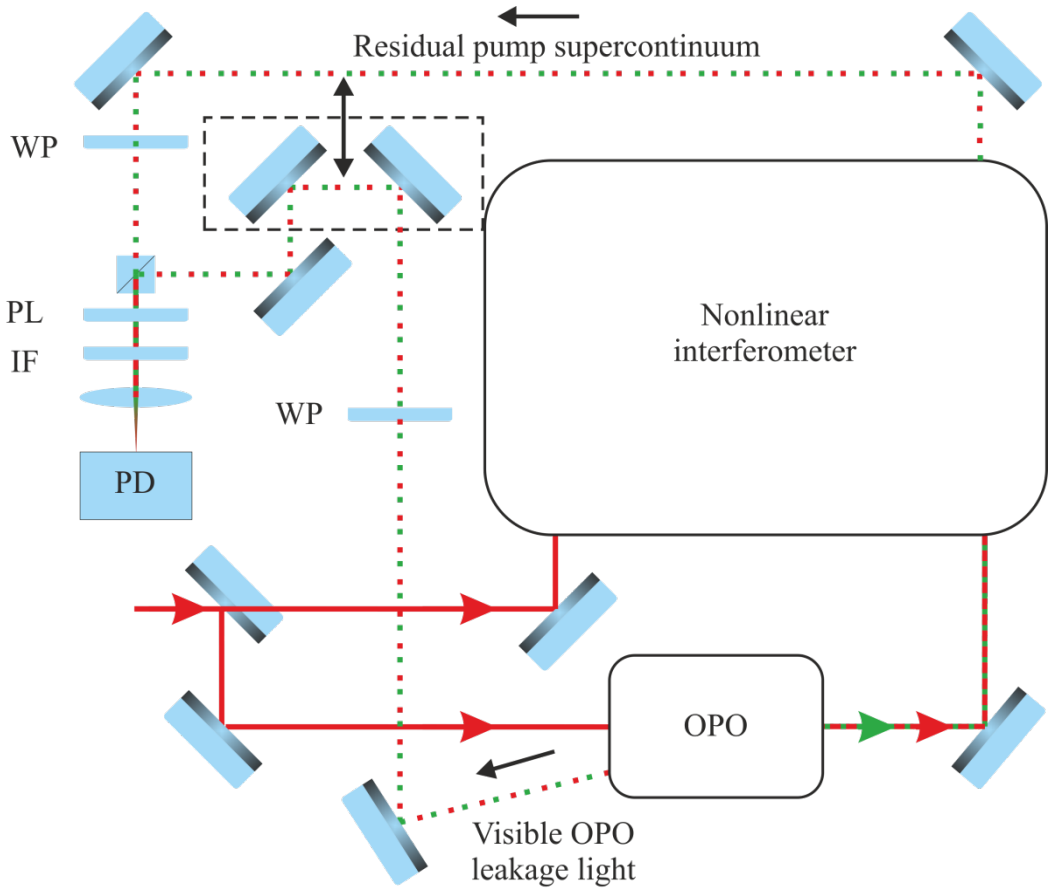

Figure 4.14. Schematic of the measurement interferometer used to confirm zero-offset locking of the pump and OPO. IF, interference filter; PD, silicon photodiode; PL, polariser; WP, wave plate.

4.3.3 Results

With the CEO frequencies of the pump and OPO correctly locked, fringes were observed between the pump supercontinuum pulses and the $p+i$ and 2*s* pulses (Figure 4.15, blue lines), indicating strong coherence over the 100 ms acquisition time of the interferogram. When either CEO frequency was unlocked, or locked to a different beat frequency, no fringes were observed, indicating a lack of coherence between the pulses (Figure 4.15, red and green lines). In practice it was possible to cycle between all four locking combinations through the application of a DC offset PZT located in each laser cavity. Approximately 30 V would shift the CEO frequency of either system by $f_{REP}/2$, allowing each source to be locked to the correct locking frequency without realignment of the interferometers.

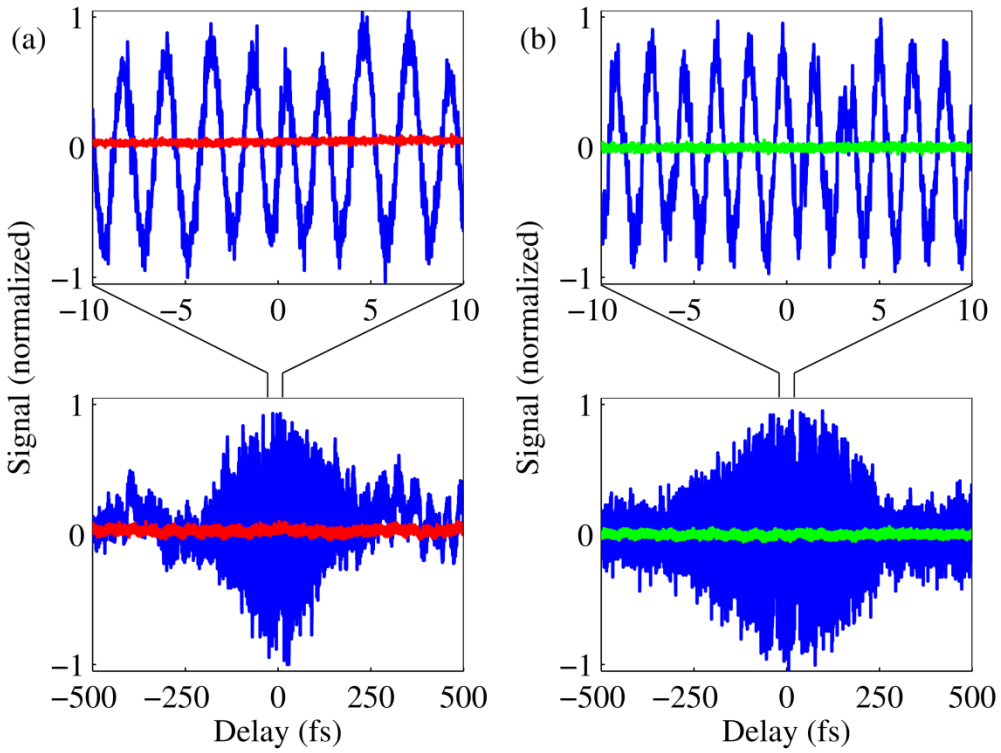

Figure 4.15. Interferograms between a pump supercontinuum and (a) $p + i$ and (b) 2s leakage light from the OPO. Interference fringes (blue) were observed when the CEO frequencies of the pump and OPO were locked in the correct configuration. Any other locking configuration resulted in no interference fringes (red, green).

Observing interference simultaneously at two distinct wavelengths demonstrated that all the CEO frequencies from the pump and OPO were locked to 0 Hz, confirming phase coherence across all the pulses on the optical bench. This coherent bandwidth extends from 0.4 µm to 3.2 µm and comprises of an ensemble of pulses sharing a common zero-offset frequency comb, the broadest zero-offset comb demonstrated to date.

4.4 Conclusions

The concept of the carrier-envelope offset frequency of an ultrafast laser was introduced, and various methods relating to its detection and control discussed. A formal approach for locking the CEO frequencies of a Ti:sapphire laser and a synchronously pumped OPO was presented, along with a discussion of the electronic feedback loops necessary to achieve a lock. By shifting the detected CEO frequencies by a known reference and then locking to that same reference, a zero-offset frequency comb has been established. A full schematic of the locking mechanism employed is shown in Figure 4.16. The presented results show that such a frequency comb can be achieved without the use of f-2f self-referencing, instead using two separate nonlinear interferometers to detect two independent beat frequencies between the pump and OPO.

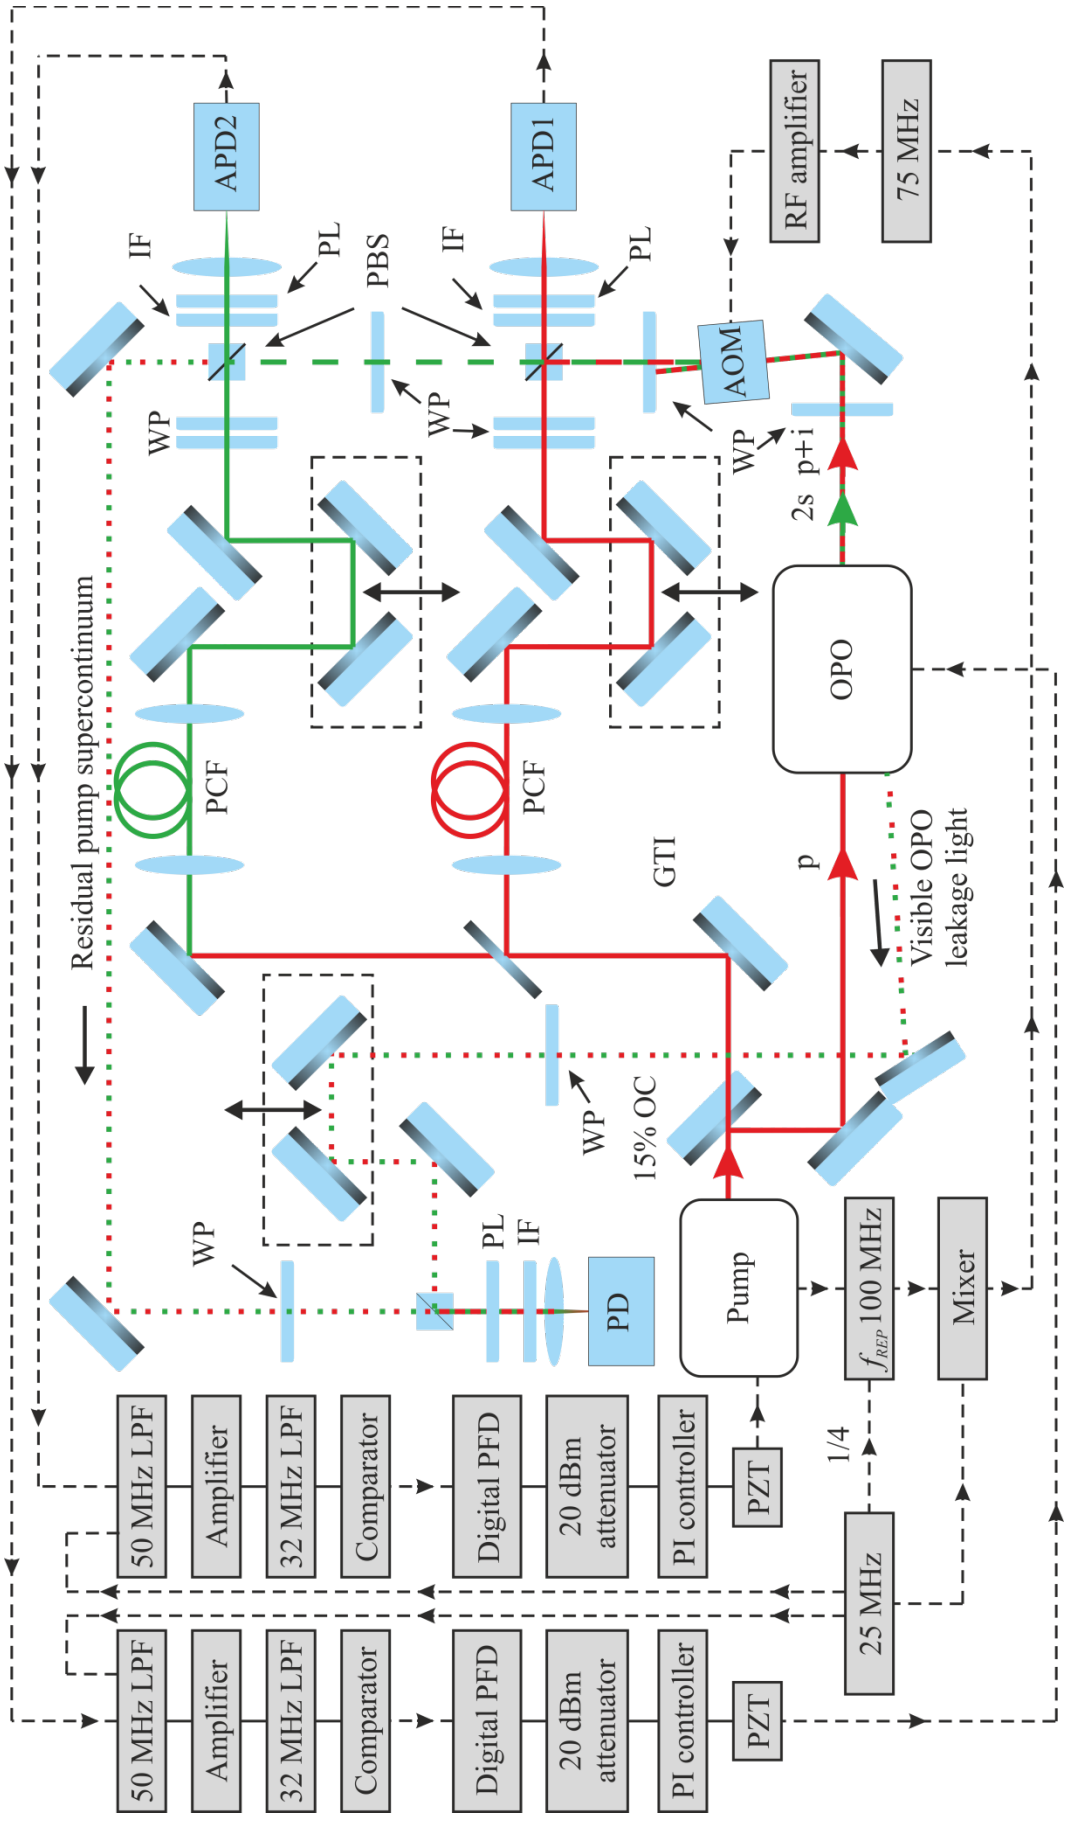

Figure 4.16. Full schematic of the locking mechanism employed to achieve a broadband zero-offset frequency comb. See Figures 4.3, 4.10 and 4.15 for details.

4.5 References

1. R. Ell, U. Morgner, F. X. Kärtner, J. G. Fujimoto, E. P. Ippen, V. Scheuer, G. Angelow, T. Tschudi, M. J. Lederer, A. Boiko, and B. Luther-Davies, "Generation of 5-fs pulses and octave-spanning spectra directly from a Ti:sapphire laser," Opt. Lett. **26**, 373–5 (2001).

2. E. Matsubara, K. Yamane, T. Sekikawa, and M. Yamashita, "Generation of 2.6 fs optical pulses using induced-phase modulation in a gas-filled hollow fiber," J. Opt. Soc. Am. B **24**, 985–9 (2007).

3. D. J. Jones, S. A. Diddams, J. K. Ranka, A. Stentz, R. S. Windeler, J. L. Hall, and S. T. Cundiff, "Carrier-envelope phase control of femtosecond mode-locked lasers and direct optical frequency synthesis," Science (80-.). **288**, 635–639 (2000).

4. J. M. Dudley and S. Coen, "Supercontinuum generation in photonic crystal fiber," Rev. Mod. Phys. **78**, 1135–1184 (2006).

5. U. Morgner, R. Ell, G. Metzler, T. Schibli, F. Kärtner, J. Fujimoto, H. Haus, and E. Ippen, "Nonlinear optics with phase-controlled pulses in the sub-two-cycle regime," Phys. Rev. Lett. **86**, 5462–5 (2001).

6. J. M. Dudley, D. T. Reid, M. Ebrahimzadeh, and W. Sibbett, "Characteristics of a noncritically phasematched Ti: sapphire pumped femtosecond optical parametric oscillator," Opt. Commun. **104**, 419–430 (1994).

7. A. Poppe, R. Holzwarth, A. Apolonski, G. Tempea, C. Spielmann, T. W. Hänsch, and F. Krausz, "Few-cycle optical waveform synthesis," Appl. Phys. B **72**, 373–6 (2001).

8. H. R. Telle, D. Meschede, and T. W. Hänsch, "Realization of a new concept for visible frequency division: phase locking of harmonic and sum frequencies," Opt. Lett. **15**, 532–4 (1990).

9. B. J. S. Gale, J. H. Sun, and D. T. Reid, "Towards versatile coherent pulse synthesis using femtosecond laser and optical parametric oscillators," Opt. Express **16**, 1616–22 (2008).

10. S. A. Meyer, J. A. Squier, and S. A. Diddams, "Diode-pumped Yb:KYW femtosecond laser frequency comb with stabilized carrier-envelope offset frequency," Eur. Phys. J. D **48**, 19–26 (2008).

11. S. Koke, C. Grebing, H. Frei, A. Anderson, and A. Assion, "Direct frequency comb synthesis with arbitrary offset and shot-noise-limited phase noise," Nat. Photonics **4**, 462–5 (2010).

12. F. Lücking, A. Assion, A. Apolonski, F. Krausz, and G. Steinmeyer, "Long-term carrier-envelope-phase-stable few-cycle pulses by use of the feed-forward method," Opt. Lett. **37**, 2076–8 (2012).

13. S. Koke, a. Anderson, H. Frei, a. Assion, and G. Steinmeyer, "Noise performance of a feed-forward scheme for carrier-envelope phase stabilization," Appl. Phys. B **104**, 799–804 (2011).

14. T. M. Fortier, D. J. Jones, J. Ye, S. T. Cundiff, and R. S. Windeler, "Long-term carrier-envelope phase coherence," Opt. Lett. **27**, 1436–8 (2002).

15. J. Reichert, R. Holzwarth, T. Udem, and T. W. Hänsch, "Measuring the frequency of light with mode-locked lasers," Opt. Commun. **172**, 59–68 (1999).

16. F. W. Helbing, G. Steinmeyer, J. Stenger, H. R. Telle, and U. Keller, "Carrier-envelope-offset dynamics and stabilization of femtosecond pulses," Appl. Phys. B **74**, s35–s42 (2002).

17. F. W. Helbing, G. Steinmeyer, U. Keller, R. S. Windeler, J. Stenger, and H. R. Telle, "Carrier-envelope offset dynamics of mode-locked lasers," Opt. Lett. **27**, 194–6 (2002).

18. J. Ye and S. T. Cundiff, *Femtosecond Optical Frequency Comb Technology*, First edition (Springer, 2005).

19. R. L. Fork, O. E. Martinez, and J. P. Gordon, "Negative dispersion using pairs of prisms," Opt. Lett. **9**, 150–2 (1984).

20. Y. Kobayashi, H. Takada, M. Kakehata, and K. Torizuka, "Optical phase locking among femtosecond subharmonic pulses," Opt. Lett. **28**, 1377–9 (2003).

21. J. H. Sun, B. J. S. Gale, and D. T. Reid, "Composite frequency comb spanning 0.4 - 2.4 μm from a phase-controlled femtosecond Ti:sapphire laser and synchronously pumped optical parametric oscillator," Opt. Lett. **32**, 1414–16 (2007).

22. T. I. Ferreiro, J. Sun, and D. T. Reid, "Locking the carrier-envelope-offset frequency of an optical parametric oscillator without f-2f self-referencing," Opt. Lett. **35**, 1668–70 (2010).

23. R. A. McCracken, J. Sun, C. G. Leburn, and D. T. Reid, "Broadband phase coherence between an ultrafast laser and an OPO using lock-to-zero CEO stabilization," Opt. Express **20**, 16269–74 (2012).

24. M. Prevedelli, T. Freegarde, and T. W. Hänsch, "Phase locking of grating-tuned diode lasers," Appl. Phys. B **60**, S241–248 (1995).

Chapter 5 - Practical considerations for coherent synthesis: noise characterisation of phase coherent pulses

5.1 Introduction

In Chapter 4 it was shown that phase coherence could be established across a broadband family of pulses from a synchronously-pumped femtosecond OPO using lock-to-zero CEO control. Establishing phase coherence is one of the key prerequisites for coherent pulse synthesis from multi-coloured pulses, as discussed in Chapter 1. A second and equally important prerequisite is that the phase coherent parent pulses have sufficiently low timing jitter between them so as to be suitable for synthesising daughter pulses which are stable over a practical time period. In this chapter the noise characteristics of the pump and OPO will be examined, with the key concepts of relative intensity noise, phase noise and timing jitter introduced and discussed in detail.

Three parent pulses were chosen from the available pump and OPO outputs as candidates for coherent synthesis. The pump (0.800 μm), doubled-signal (0.530 μm) and pump-signal (0.456 μm) pulses were preferred over other pulses for a number of reasons. The first criterion for pulse selection was average power; the power levels of these pulses were higher than those of the doubled-pump (0.400 μm) and pump-idler (0.642 μm) pulses. The second criterion was the frequency bandwidth supported by the parent pulses; the three chosen pulses spanned a 0.28 PHz bandwidth, capable of supporting a 3.6 fs pulse. The addition of the doubled-pump pulses would increase the bandwidth to 0.37 PHz (2.7 fs), but with an increased cost in the complexity of the optical system. Adding the pump-idler pulses would not increase the total frequency bandwidth. The final criterion was one of practicality. The prism-based pulse compression method discussed in Chapter 6 was sufficiently complex with only three input wavelengths. The sum-frequency mixing crystal employed in the XFROG and cross-correlation apparatus, discussed in Chapter 3 and Sub-section 5.4.1 respectively, did not efficiently phasematch between the pump and doubled-pump wavelengths, and so that pulse could not have been characterised. With these considerations in mind the relative intensity noise of the three parent pulses mentioned above was characterised, along with their relative timing jitter.

5.2 Relative intensity noise

The output of a laser can be described as a time-varying intensity, fluctuations of which are known as the relative intensity noise (RIN). For a mode-locked laser this can be thought of as the pulse-to-pulse intensity variation, as illustrated in Figure 5.1. These intensity variations arise due to a number of factors including beam pointing, optical stability and electrical noise.

5.2.1 Power spectral density

Noise can be represented by a power spectral density (PSD) plot, which describes the noise power per unit frequency interval relative to the mean signal. The general equation for PSD can be obtained using the method described below [1]. The Fourier transform equations for a time dependent function $h(t)$ can be written as

$$H(\omega) = \int_{-\infty}^{\infty} h(t) e^{i\omega t} dt \tag{5.1}$$

$$h(t) = \frac{1}{2\pi} \int_{-\infty}^{\infty} H(\omega) e^{-i\omega t} d\omega . \tag{5.2}$$

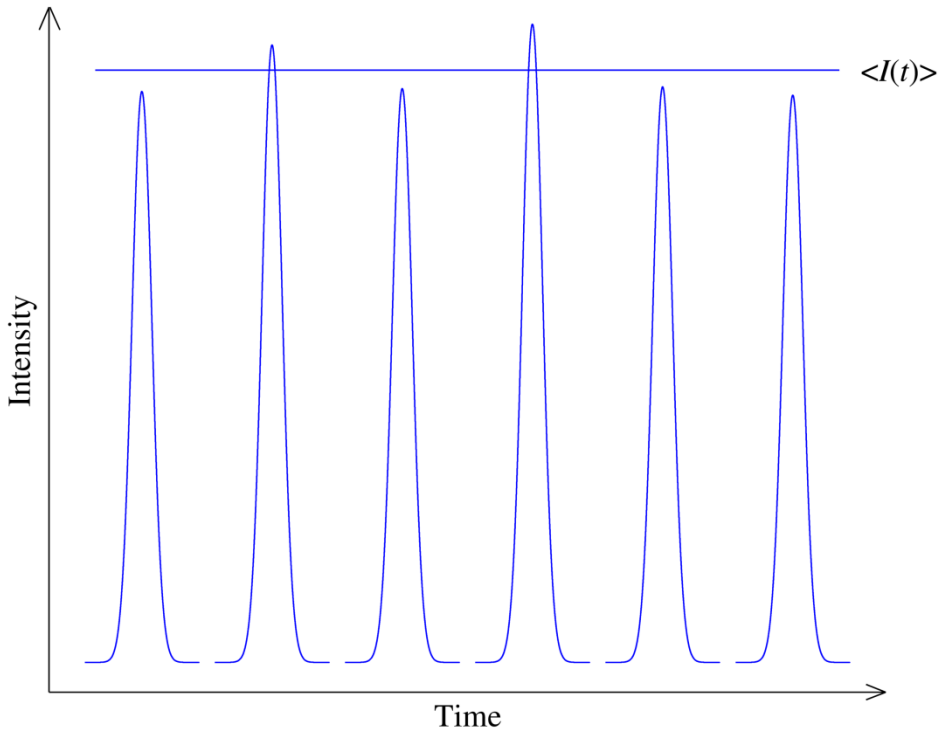

Figure 5.1. Simple illustration of pulse-to-pulse intensity fluctuations from a laser.

We can rewrite the frequency term as $\omega = 2\pi f$ to get

$$H(f) = \int_{-\infty}^{\infty} h(t) e^{2\pi i f t} dt \tag{5.3}$$

$$h(t) = \int_{-\infty}^{\infty} H(f) e^{-2\pi i f t} df \ . \tag{5.4}$$

The expression for total power can be given by Parseval's theorem, which states that the energy contained in a signal is equal in both the time and frequency domains [1], written formally as

$$E_{TOT} = \int_{-\infty}^{\infty} |h(t)|^2 \, dt = \int_{-\infty}^{\infty} |H(f)|^2 \, df \ . \tag{5.5}$$

The power spectral density, or power in a frequency interval, can therefore be described by the combined power of positive and negative frequencies, written as

$$P_h(f) = |H(f)|^2 + |H(-f)|^2 \tag{5.6}$$

for $0 \leq f \leq \infty$.

The analysis of the results presented in this thesis used discrete Fourier transforms as the noise measurements taken contained a finite number of points. The discrete Fourier transform of a quantity h_j containing N points is given by

$$H_k \equiv \sum_{j=0}^{N-1} h_j e^{2\pi i \frac{jn}{N}} \ . \tag{5.7}$$

The PSD can be normalised using the mean-squared-amplitude. The integral of the PSD over the range of 0 to the Nyquist frequency [1] is equal to half the mean-squared-amplitude, and this is used to normalise the PSD for each frequency. The normalised PSD can be integrated over the frequency range to give the cumulative noise, written as

$$\Delta h_{CUM} = \left[2 \int_0^{f_{Nyquist}} P_h(f) df \right]^{1/2} \tag{5.8}$$

where $P_h(f)$ is the normalised PSD and $f_{Nyquist}$ is the Nyquist frequency.

The mathematical methods described above were used to analyse the relative intensity noise, CEO frequency phase noise and timing jitter noise described in this chapter. The latter two measurements will be described in later sections.

5.2.2 Relative intensity noise in the pump and OPO outputs

The relative intensity noise of a signal is its relative fluctuation in a 1 Hz bandwidth, and can be defined as the mean-squared noise power divided by the square of the average power [2]. RIN measurements were carried out for the Verdi pump source, the Ti:sapphire laser and the visible doubled-signal and pump-signal OPO outputs. The outputs from each laser source were monitored on individual silicon photodiodes and amplified before the data were collected using a 12-bit data acquisition card. The amplifier was used to fill the dynamic range of the data acquisition card and also to record and remove the average signal from each data point, ensuring the recorded data were a fluctuation around zero. This is illustrated in Figure 5.2. The data were processed as described above with the PSD calculated relative to the mean signal. The resultant plot has units of dBc^2/Hz, where c is the carrier voltage. The PSD was plotted on a dB scale, a standard method of displaying RIN measurements [1]. The cumulative standard deviation was calculated as a percentage relative to the carrier.

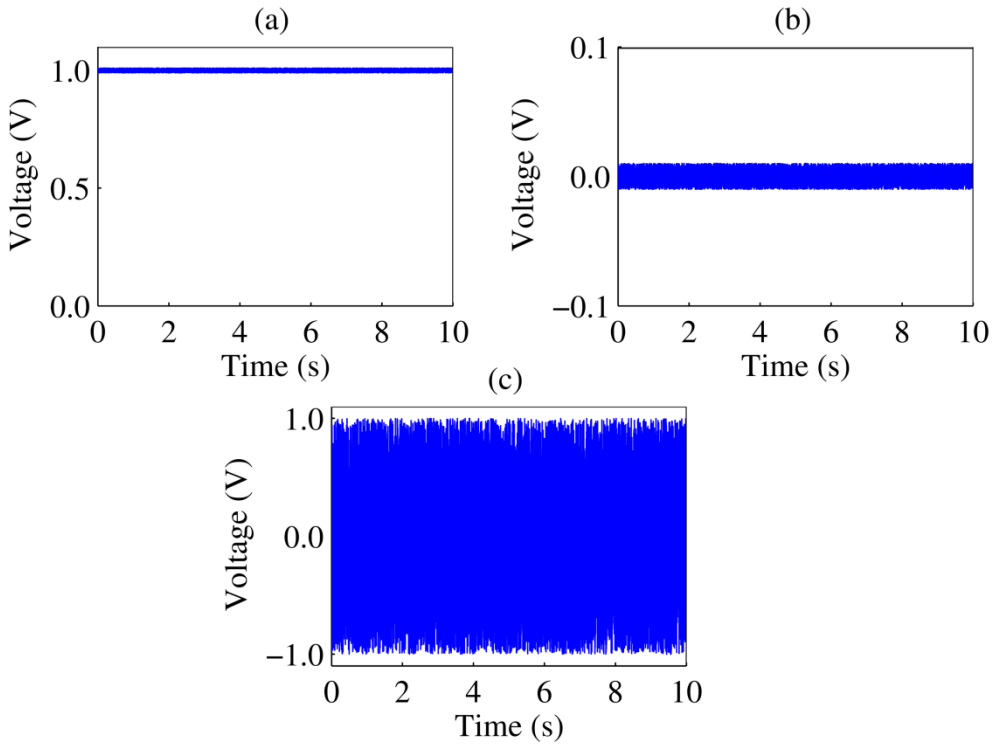

Figure 5.2. Relative intensity noise (RIN) signal at different stages of formatting. (a) The raw signal detected directly by the photodiode, here modelled as a 1 V carrier signal with a 10 mV noise level; (b) the signal after having the DC carrier signal removed; (c) the amplified signal sent to the data acquisition card.

RIN for the Verdi pump source is shown in Figure 5.3. The cumulative standard deviation was 0.05%, the majority of which is contributed frequencies below 100 Hz and most likely arising from acoustic vibrations and thermal effects. The single-longitudinal-mode operation of the Verdi makes it characteristically quiet at higher frequencies, unlike multi-longitudinal-mode pump sources that suffer from mode competition and beating effects.

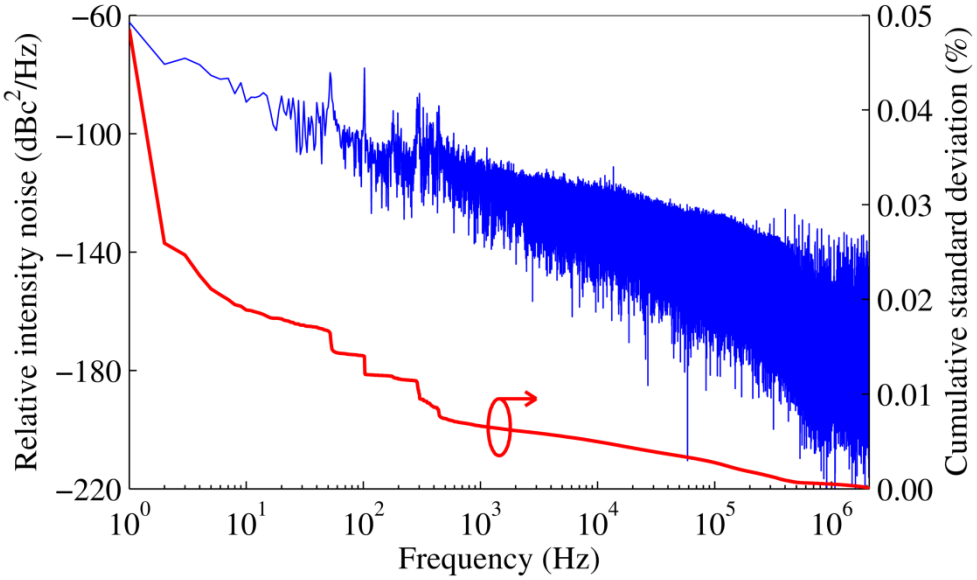

Figure 5.3. Relative intensity noise from the Verdi pump source. The single-mode operation of the laser is characteristically quiet at higher frequencies.

Results of the RIN measurements for the Ti:sapphire pump laser are shown in Figure 5.4 for cases where the CEO frequencies of both the pump and OPO were unlocked and locked. The cumulative standard deviation for both measurements is approximately 0.04%, however the frequency noise contributions from the measurements are quite different. When the CEO frequencies of both lasers were unlocked the RIN from the pump laser showed significant contributions at frequencies from 1 kHz to 1 MHz. Noise contributions from these frequencies were suppressed when the CEO frequencies were locked, with environmental noise dominating the RIN.

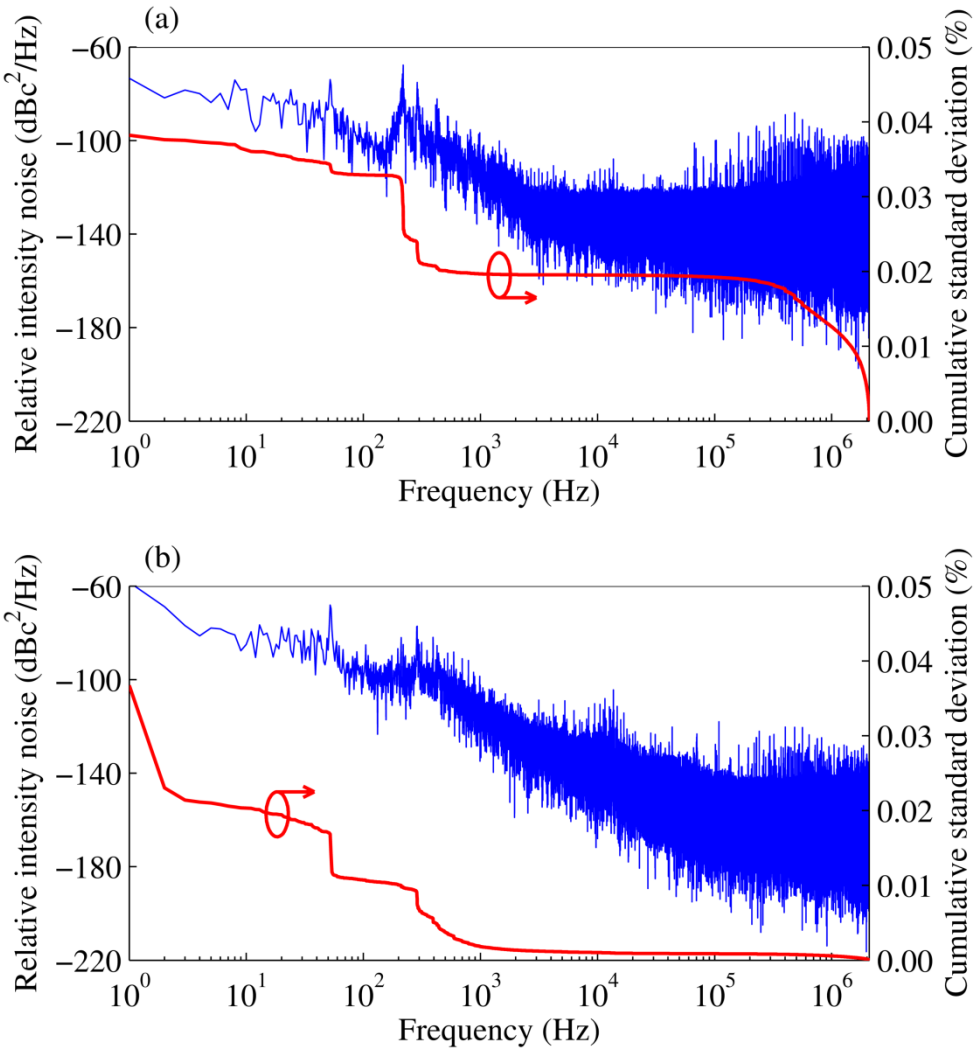

Figure 5.4. Relative intensity noise from the Ti:sapphire pump laser when the CEO frequencies of the pump and OPO were (a) unlocked and (b) locked.

RIN measurements from two OPO outputs, the doubled-signal at 0.530 µm and the pump-signal SFM at 0.456 µm, are shown in Figure 5.5 and Figure 5.6 respectively. Both sets of RIN measurements are very similar in terms of the frequency components that contribute to the overall cumulative standard deviation. When the CEO frequencies of the pump and OPO are unlocked the RIN is dominated by frequencies below 2 kHz. These frequencies arise from acoustic noise, thermal variations across the optical bench and small air currents within the boxed OPO cavity.

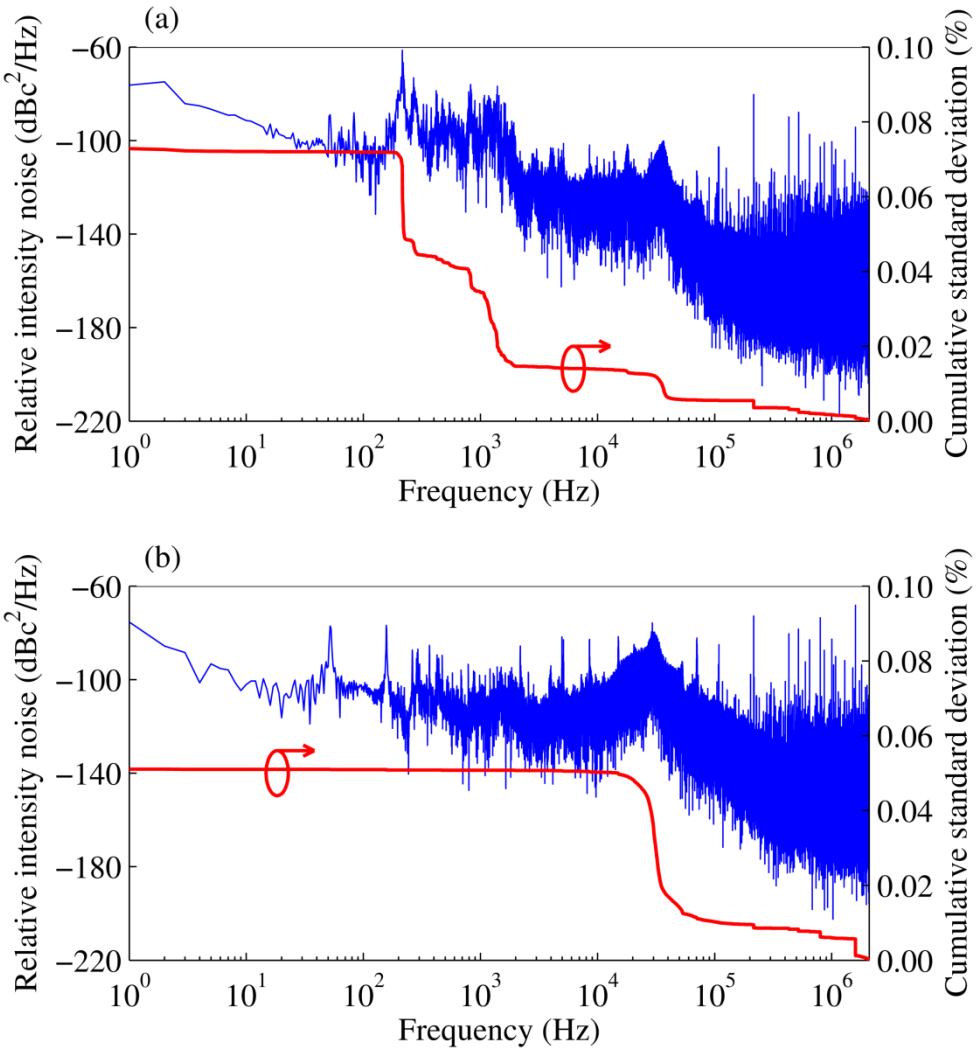

Figure 5.5. Relative intensity noise from the doubled- signal output of the OPO when the CEO frequencies of the pump and OPO were (a) unlocked and (b) locked.

The intracavity dispersion of the OPO is near-zero, with the crystal providing the only significant contribution to the net-positive GDD. Equation 4.9 in Chapter 4 showed that the gross tuning behaviour of a synchronously-pumped OPO is determined by the intracavity dispersion, and that the centre frequency of a cavity with low GDD will tune quickly with cavity length. As a result of this the centre frequency of the signal pulses in this OPO is not stable, but rather fluctuates as environmental effects cause the cavity length to change on the nanometre level. This has important consequences for the OPO visible mixing outputs which are a product of coincidental phasematching within the crystal. As the centre frequency of the signal shifts the efficiency of the phasematching will change, causing fluctuations in the intensity of the visible mixing pulses. These fluctuations were measured as cumulative standard deviations of roughly 0.08% for the case of unlocked CEO frequencies.

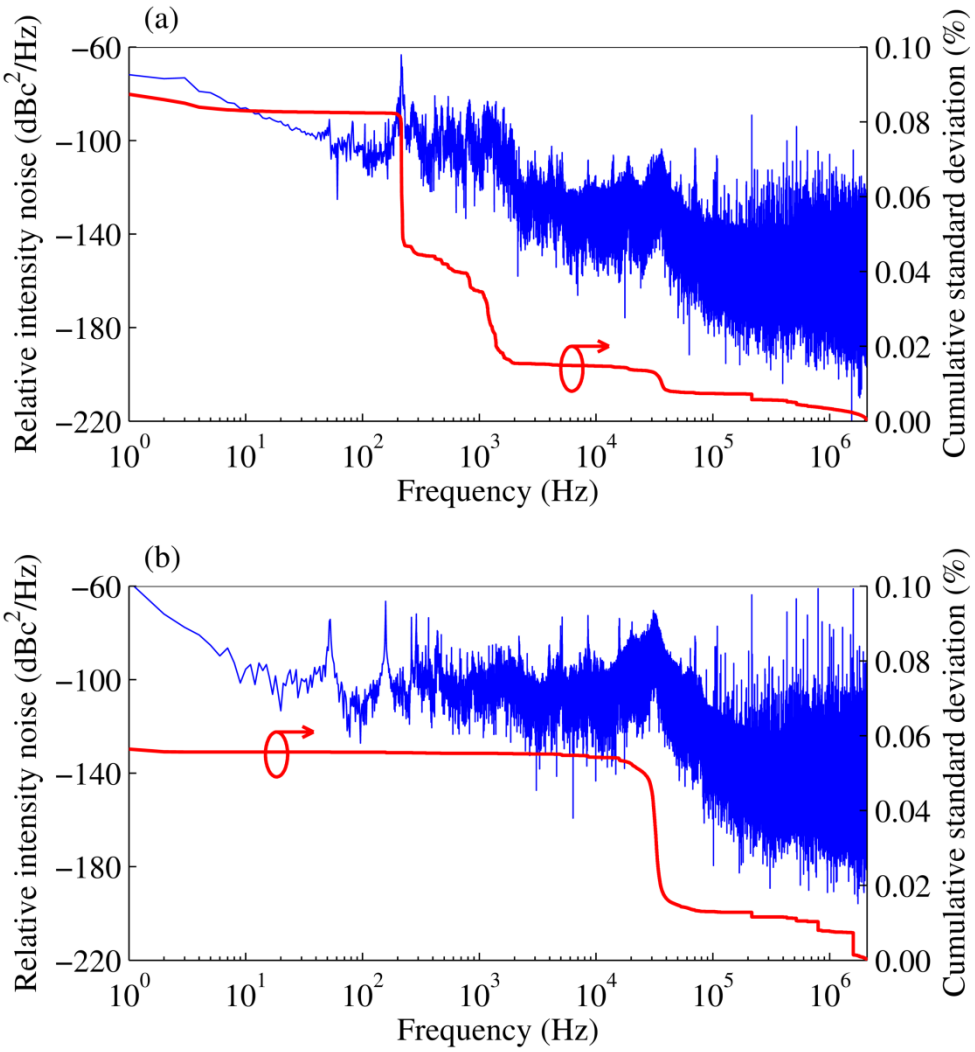

Figure 5.6. Relative intensity noise from the pump-signal SFM output of the OPO when the CEO frequencies of the pump and OPO were (a) unlocked and (b) locked.

When one of the CEO frequencies of the OPO is locked the cavity length is well controlled, effectively locking the centre frequency of the signal pulses, as discussed in Chapter 4, Sub-section 4.1.3. Provided environmental noise is kept to a minimum, for example by floating the optical bench and boxing the OPO cavity, then frequencies below 2 kHz are no longer dominant. Higher-order frequencies above 10 kHz provide the majority of the noise, arising due to amplitude fluctuations in the Ti:sapphire laser coupling into the OPO as both amplitude and phase noise. The latter of these will affect the stability of the OPO CEO frequency lock, again causing small drifts in the centre frequency of the signal pulses and introducing further relative intensity noise. Together, these effects explain why the OPO RIN was lower when the CEO frequencies of the OPO pulses were locked.

5.3 Phase noise

An ideal monochromatic laser has an electric field that can be described by

$$E(t) \propto \mathrm{Re}\left[A_0 e^{i(\omega_0 t+\phi)} \right] \tag{5.9}$$

where the phase ϕ is constant. A true laser output is not ideal as phase and amplitude fluctuations lead to a finite linewidth, with the electric field instead described by

$$E(t) = \mathrm{Re}\left[E_0 e^{i(\omega_0 t+\phi(t))} \right] \tag{5.10}$$

where $\phi(t)$ is a slowly varying phase. Locking the CEO frequency of a laser attempts to minimise fluctuations in the phase of the electric field of the laser pulses. In a method similar to relative intensity noise, phase noise can be described as a fluctuation in radians over 1 Hz bandwidth. The cumulative phase noise is determined by summing the noise across the recorded frequency range, and is equal to the RMS noise for that frequency [3].

5.3.1 Phase noise in the CEO locking loops

In-loop phase noise PSD and cumulative phase noise measurements were made for both CEO frequency locking loops. Part of the output from the PFD circuit was amplified and sent to a 12-bit data acquisition card in a method similar to that for recording relative intensity noise. Calibration of the PFD phase response was necessary to calculate the phase error between the locked CEO frequencies and the reference signal. In order to calibrate the voltage from the PFD as a phase change, two frequency sources with a common reference were set to 10 MHz were used as the input to the PFD. One frequency source was offset by 1 Hz causing the PFD output to move from its minimum to its maximum voltage. In one second the accumulated phase slip is (by definition) 2π, therefore the time and voltage range of this transition yields a slope in V/rad, allowing the voltage recorded in the noise measurements to be converted into a phase in radians.

It is not immediately intuitive what the noise measured at the output of the PFD circuit represents; the signal is clearly a phase change, however it is useful to attach this phase change to a physical mechanism. Recall that (from Sub-section 4.1.1) the CEO frequency is related to the carrier-envelope phase slip (CEPS) by

$$f_{CEO} = \frac{\varphi_{CEPS}}{2\pi} f_{REP}\,. \tag{5.11}$$

If we assume that the CEPS changes by a small amount per pulse then we have

$$\Delta f_{CEO} = \frac{\Delta \varphi_{CEPS}}{2\pi} f_{REP} . \tag{5.12}$$

Small changes in the CEO frequency are observed as a phase change over the observation time, and so we write

$$\frac{\Delta \phi_{CEO}}{t_{OBS}} = \varphi_{CEPS} f_{REP} \tag{5.13}$$

which can be rearranged to find the accumulated optical CEP slip, given by

$$\Delta \phi_{CEO} = \varphi_{CEPS} (f_{REP} t_{OBS}) \tag{5.14}$$

where φ_{CEP} is the phase slip per pulse and $(f_{REP} t_{OBS})$ is the number of pulses in the observation time. Therefore the change in CEO frequency, as measured at the output of the PFD circuit, tells us directly about the total accumulated CEPS in the observation time.

The results of the phase noise measurements are shown in Figure 5.7. The cumulative phase noise was integrated up to 1 MHz, giving an error of 0.18 rad for the $p+2s$ (pump) locking loop and 0.11 rad for the idler (OPO) locking loop. These results compare well with previous results published within the group [4,5], and some conclusions can be drawn from the data. The increased noise in the locking loop used to control the pump CEO frequency could be attributed to a number of factors. The CEO behaviour of laser cavities that employ dispersive prisms has been shown to be significantly noisier than those without [6,7] due to the extra optical elements in the cavity. The mirror attached to the PZT in the pump cavity is larger than that attached to the PZT in the OPO, and the increased mirror inertia may affect the frequency response. The CEO frequency derived in the pump locking loop is also determined by two separate CEO frequencies, one from the pump and the other from the doubled-signal. This is in contrast with the CEO frequency used in the OPO locking loop which is purely from the idler. Small changes in either the pump or OPO CEO frequencies will therefore couple into the pump locking loop, which is not the case for the OPO locking loop as the pump CEO frequency is rejected at the APD.

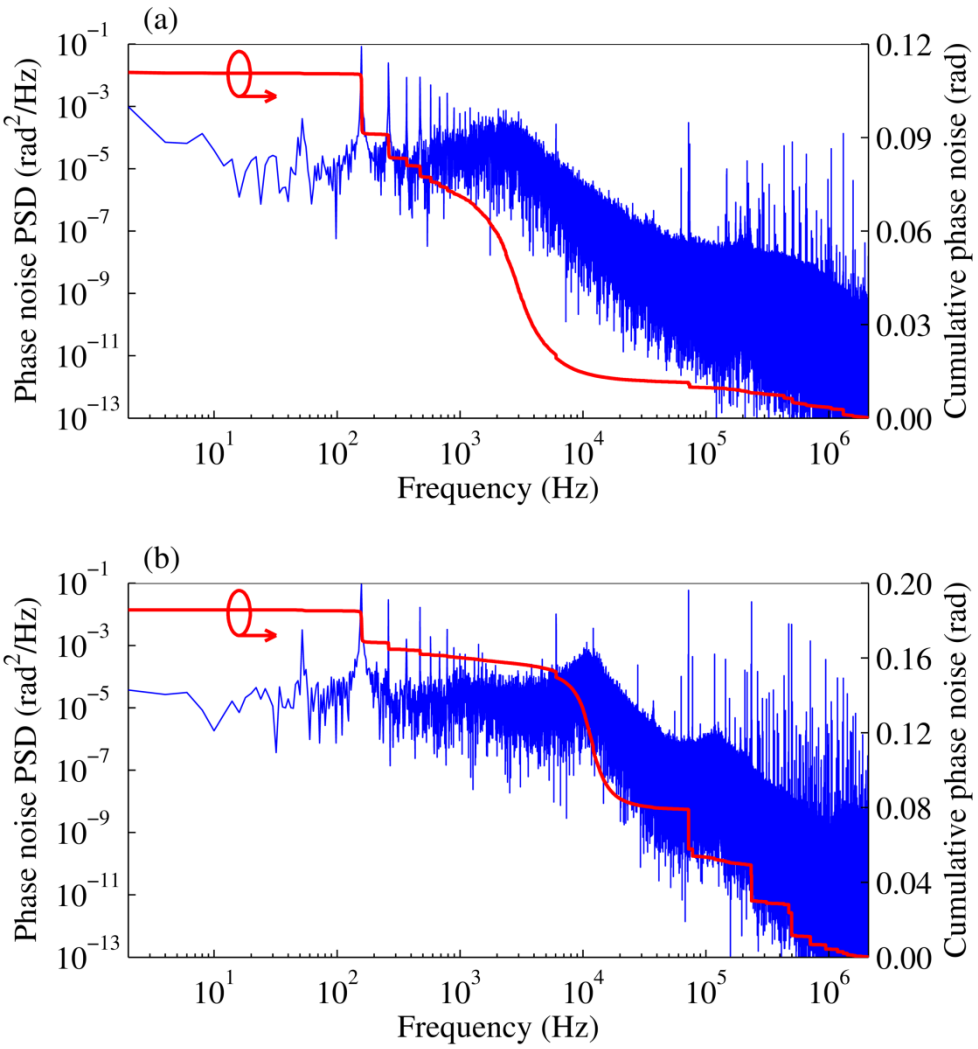

Figure 5.7. In-loop measurements of the phase noise PSD of the pump and OPO CEO frequencies for CEO frequencies locked. Phase noise PSD (blue) and cumulative phase noise (red) for (a) the idler and (b) the *p* + 2s CEO frequency. The cumulative phase noise integrated up to 1 MHz is (a) 0.11 rad and (b) 0.18 rad over a 1 second observation time.

With this reasoning the cumulative phase noise from the OPO locking loop should be significantly lower than that of the pump, however the two are comparable. This can be attributed to intensity fluctuations in the pump coupling into the OPO as both amplitude and phase noise. Finally, neither the Ti:sapphire laser nor the OPO were repetition rate stabilised. Frequency comb breathing due to repetition rate changes would create additional noise that could be suppressed with an additional feedback loop in the Ti:sapphire cavity.

In physical terms the cumulative phase noise describes the RMS variation of the CEO frequency over the observation window. It is worth noting that this is only true for the case of lock-to-zero CEO stabilisation; when the CEO frequency is locked to another

value then the cumulative phase noise describes the RMS variation in the carrier-envelope phase-slip. As the CEO frequencies are locked to zero, the peak of the carrier of the electric field sits under the peak of the envelope. The RMS CEO frequency variation describes the excursion of the CEO frequency from this central position. As the RMS variation is given in radians, it is sensible to restate this variation as a fraction of the carrier period. This is particularly important when discussing timing jitter as it will be a determining factor in the observation window for phase coherence. The RMS for both the pump and CEO locking loops was less than 0.18 radians in a 1 second observation window, less than 3% of an optical cycle, and so the position of the CEO frequency has slipped by 3% relative to the reference frequency.

5.4 Timing jitter

To establish phase coherence among a set of pulses, either from independent lasers or from a synchronously-pumped source, it is essential to achieve sufficient synchronisation such that the mutual remaining timing jitter between the pulses is less than the oscillation period of the optical carrier, given by

$$\tau_c = \frac{\lambda_0}{c} \tag{5.15}$$

where λ_0 is the centre wavelength of the pulse. The oscillation periods for relevant pulses are shown in Table 5.1.

Table 5.1. Oscillation period of the optical carrier for a pulse with a given centre wavelength.

Pulse origin	Centre wavelength (μm)	Oscillation period (fs)
p	0.800	2.67
p+i	0.642	2.14
2*s*	0.530	1.76
p* + *s	0.456	1.50
2p	0.400	1.33

For viable pulse synthesis between the pump and visible OPO pulses it is therefore necessary that the mutual timing jitter between the parent pulses be much less than 1 fs, ideally less than 100 as, over a given observation window, a length of time dependent on the specific experiment that requires a synthesised pulse.

5.4.1 Timing jitter between the pump and visible OPO pulses

Timing jitter arises from two primary phenomena, vibrations and group delay dispersion (Figure 5.8). Consider two laser sources with the same centre wavelength. Each laser will experience its own acoustic vibrations as they are constructed from separate optomechanics. A small vibration on a mirror in the cavity for the first laser will not be present in the second laser. This leads to the pulses from each laser having slightly different timing characteristics, both from minute repetition rate differences and extra-cavity mirror vibrations, and so in a jitter-sensitive experiment great care must be taken to stabilise the pulse trains from both lasers [8–11].

The principal approach for controlling the jitter between two pulse trains that have intrinsic jitter, e.g. two independent mode-locked lasers, is to employ a balanced cross-correlator [12–14]. The cross-correlator yields a photodiode signal that depends on the relative arrival time of the pulse from the independent lasers. This photodiode signal can be used as the input to a feedback loop that controls the repetition rate of one of the lasers, allowing them to be locked together with attosecond timing jitter.

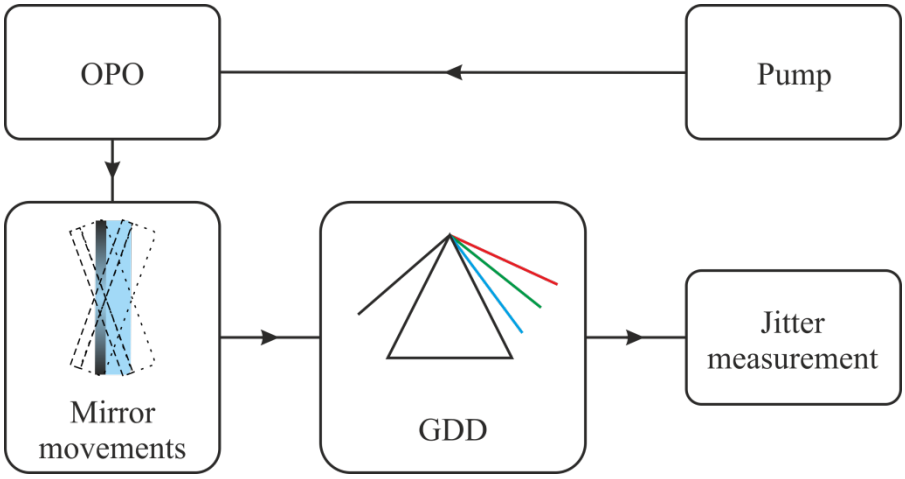

Figure 5.8. Jitter is accrued through mirror vibrations and group delay dispersion.

This work circumvents the problem in two ways. Using a Ti:sapphire pump laser and a synchronously-pumped femtosecond OPO ensures that both sources have the same repetition rate. Keeping the synthesis parent pulses on a common path ensures common-mode rejection of mirror vibrations. In these experiments the only non-common path occurs after the beam splitter in Figure 5.9.

The other primary source of timing jitter is from group delay dispersion, but for ease of explanation it is instructive to use the simpler terminology of chromatic dispersion. As different wavelengths propagate through a medium they travel at different group velocities as they experience wavelength-dependent refractive indices. This will give rise to a group delay difference between pulses of two different wavelengths. Consider the example shown below in Table 5.2 for the pump, doubled-signal and pump-signal wavelengths as the pulses propagate through 6.3 mm of BK7, a standard mirror thickness and substrate. The doubled-signal pulses are now delayed by 105 fs relative to the pump pulses, with the pump-signal pulses further delayed.

Table 5.2. Group delay differences between the pump and two visible OPO outputs after propagating through 6.3 mm of BK7.

Pulse origin	Centre wavelength (μm)	Refractive index	Group delay (ps)	Group delay difference from pump (fs)
p	0.800	1.5108	31.415	-
2*s*	0.530	1.5196	31.310	105
p* + *s	0.456	1.5253	31.187	228

The common-path approach employed in these experiments minimises variations in the group delay dispersion between the pulses. This is only true however if the centre wavelength of each pulse remains constant. If the centre wavelength of the pump or signal (and therefore the double-signal or pump-signal) pulses shifts by a small amount then the group delay dispersion between the pulses will shift also. As an example, take the case of a 1 nm decrease in the centre wavelength of the pump pulses propagating through a BK7 substrate as above. The difference in group delay at the pump wavelength is 5 as, however the difference at the new pump-signal wavelength of 0.4555 μm is 1 fs, almost an entire optical cycle. It is extremely important to try and stabilise the wavelengths of the pulses so that the mutual timing jitter can be kept to a minimum.

As has been discussed in previous sections and chapters, locking a CEO frequency of the OPO is the same as locking its centre wavelength. Additional GDD introduced by changes in the signal wavelength will therefore be extremely small as there will only be small deviations from the mean wavelength. Locking the CEO frequency of the pump laser does not introduce the same wavelength stability however; the PZT used in the CEO feedback loop is placed behind a mirror after dispersing prism pair, and it would require extremely large angular deviations of this mirror to change the centre wavelength in this way. Rather, small changes in the pump wavelength are not observed by the OPO when the OPO CEO frequency is controlled. The phasematching bandwidth of the PPKTP crystal OPO is very broad (see Chapter 3, Figure 3.10), and therefore a pump spectrum with a slightly different centre wavelength will still phasematch efficiently. Moreover the OPO CEO locking loop will accommodate pump wavelength changes by adjusting the cavity length accordingly to produce the same signal output.

5.4.1.1 Method

A balanced cross-correlator was constructed to measure the mutual timing jitter between the pump and two visible pulses from the OPO. A schematic of the cross-correlator is shown in Figure 5.9. The collinear beams were split using a visible-NIR beam splitter, with the two arms travelling separate delay paths before being recombined using a polarising beam splitter cube. As all beams were horizontally polarised a half-wave plate was placed in the visible arm to produce vertically polarised light. The two orthogonal beams were focused into a 30-µm-thick BBO crystal and the resulting UV SFM mixing light focused onto a silicon photodiode after suitable filtering. A more detailed description of this apparatus was given in Chapter 3 while discussing XFROG measurements of the pump and OPO pulses. The balanced cross-correlator is simply the XFROG apparatus with the delay stage kept at a fixed position.

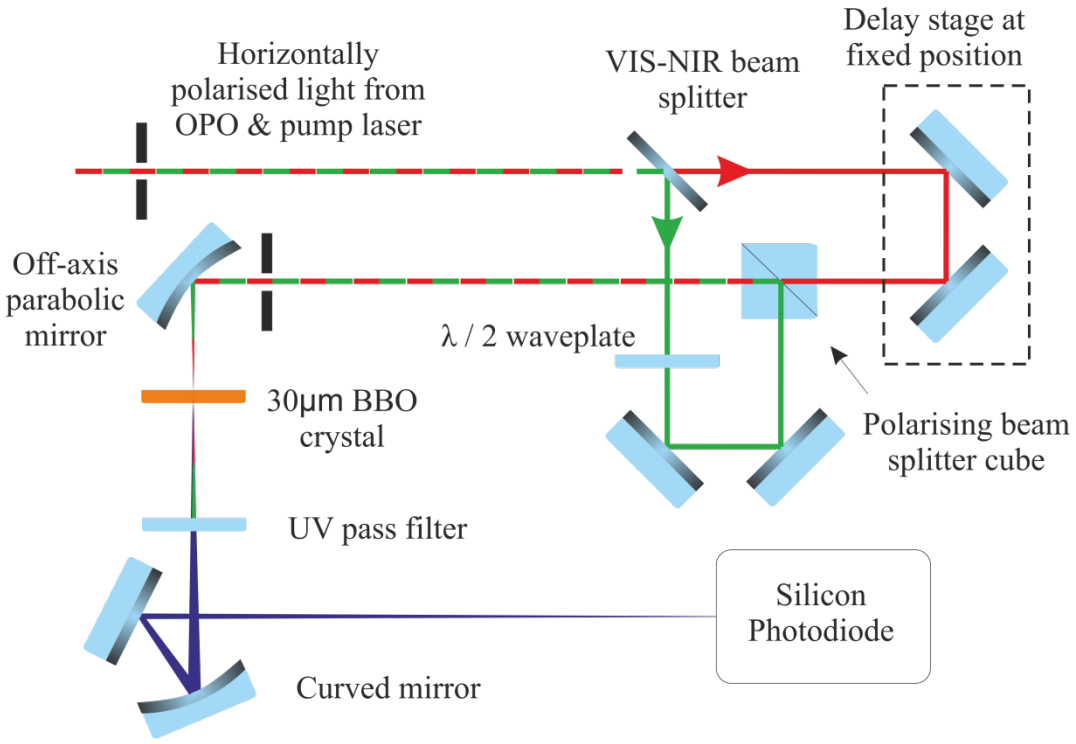

Figure 5.9. Schematic of the balanced cross correlator used to measure the mutual timing jitter between the pump and OPO pulses. This set-up was also used to measure the RIN of the pulses by blocking the unwanted beams, removing the crystal and replacing the UV pass filter with a suitable narrowband interference filter.

As the position of the delay stage was scanned the photodiode voltage increased and decreased as SFM light was generated. This scan was recorded using MATLAB and the signal processed by removing the background, as shown in Figure 5.10 (a) and (b). An area with a linear time-voltage relationship was identified and a mathematical fit performed. The coefficients of this fit were used to determine a linear relationship

between RMS voltage error and timing jitter, as shown in Figure 5.10 (c) and (d). The delay stage was moved to the linear time-voltage position identified and its position fixed. The voltage level of the photodiode was recorded on a 12-bit acquisition card in the same manner as the relative intensity noise measurements described in Sub-section 5.2.2. Excursions from the mean voltage were then converted into time fluctuations to provide a measure of the total timing jitter between the pump and visible OPO pulses. The term total timing jitter will be elaborated upon in Sub-section 5.4.2. Timing jitter results are presented in a similar manner to both RIN and phase noise, however the PSD scale is now calibrated in as^2/Hz, while integrating the PSD provides the cumulative mutual timing jitter between the two pulses, calculated in attoseconds.

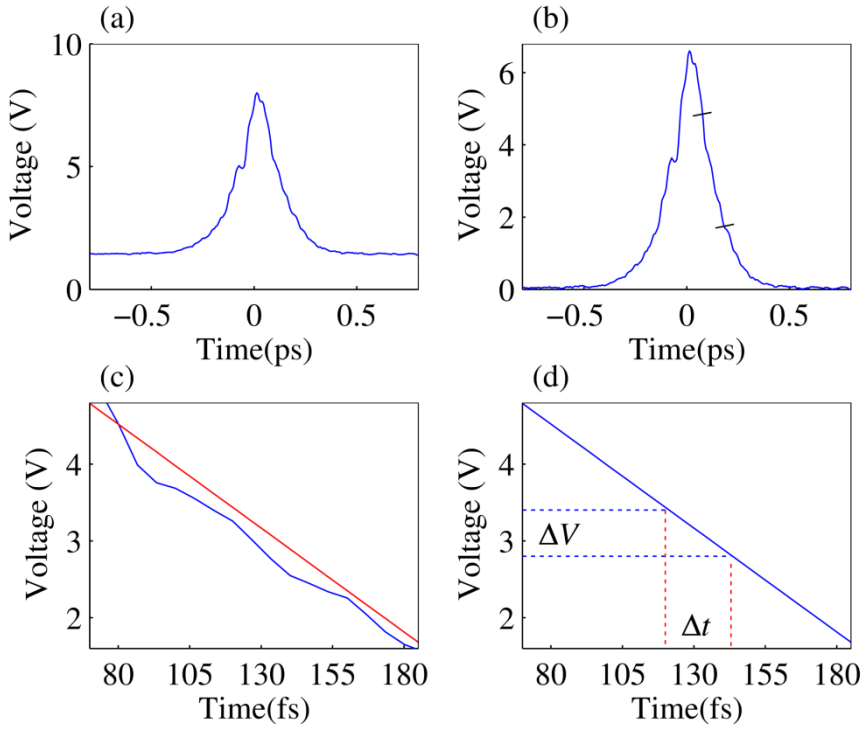

Figure 5.10. Outline of the calibration method used to measure the timing jitter between the pump and OPO pulses. (a) Raw photodiode signal; (b) background-free signal; (c) region of linear time-voltage relationship (blue) with fit (red); (d) fluctuations in voltage correspond to a fluctuation in timing jitter.

5.4.1.2 Results

The results of the mutual timing jitter measurement between the pump and doubled-signal pulses are shown in Figure 5.11 (a) and (b), for both CEO loops unlocked and locked respectively. When the CEO loops were unlocked the relative timing jitter was poor and measured at almost 950 as over a one-second observation period. In this case the centre wavelength of the OPO was not controlled and therefore would fluctuate due

to environmental noise. The delay accrued due to the GDD of various lenses and mirror substrates leads to the large timing jitter.

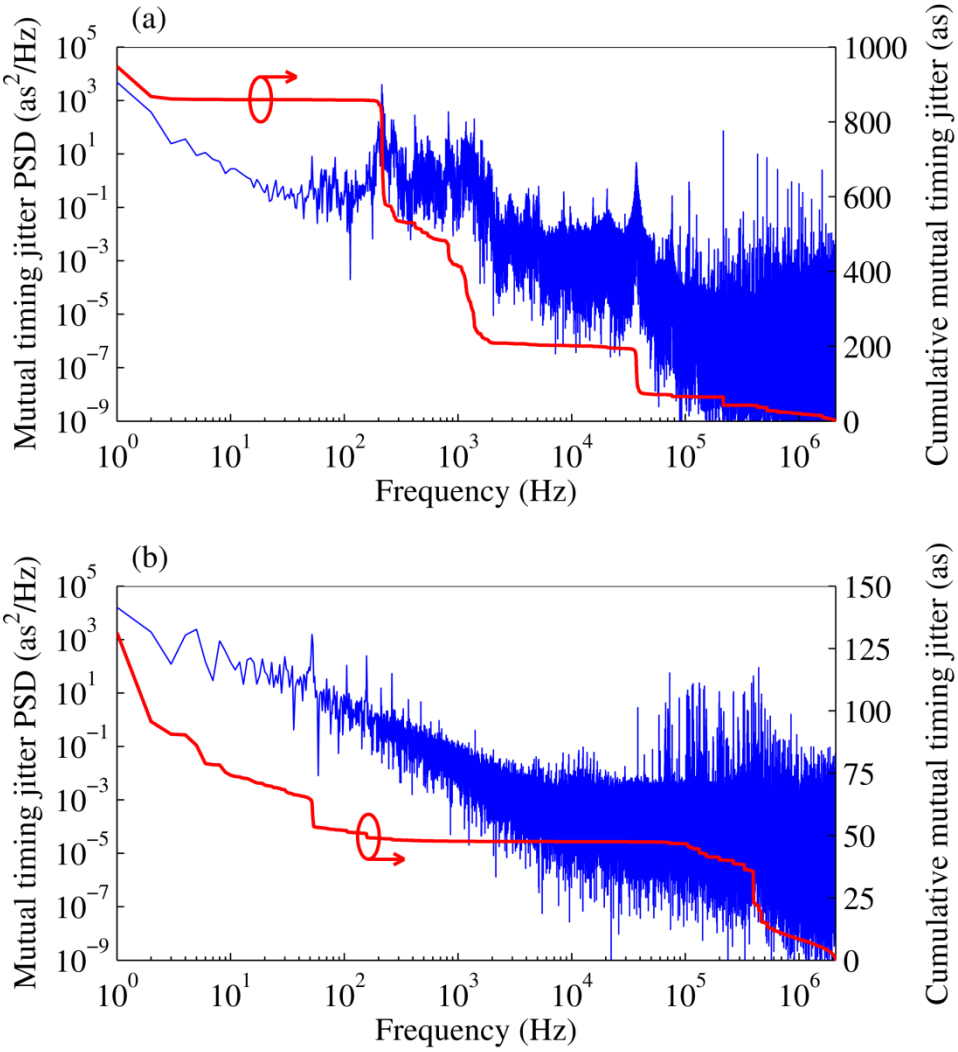

Figure 5.11. Mutual timing jitter between the pump and doubled-signal pulses for both CEO locking loops (a) unlocked and (b) locked. When both CEO loops are locked the total cumulative mutual timing jitter is 132 as.

When the CEO frequencies of both systems were locked the total cumulative timing jitter was 132 as over a one-second observation period. There was a substantial decrease in noise below 10 kHz, with the largest jitter contributions coming at low environmental frequencies. This is most likely due to the common path approach that ensures common-mode rejection of acoustic and vibrational noise up to the beam splitter in the cross correlator. The excellent timing stability between the pump and synchronous OPO, coupled with the wavelength stability of the locked OPO and common optical path approach, ensured that the envelopes of the pump and doubled-signal pulses drifted by only 132 as in the observation time of one second.

Similar timing jitter measurements between the pump and pump-signal SFM pulses are shown in Figure 5.12. When the CEO frequencies were not stabilised the cumulative mutual timing jitter between the pulses was almost 1.1 fs, but when the CEO frequencies were locked the jitter was reduced to 137 as over the 1 second observation period. As with the pump and doubled-signal result the jitter is dominated by lower frequency environmental noise.

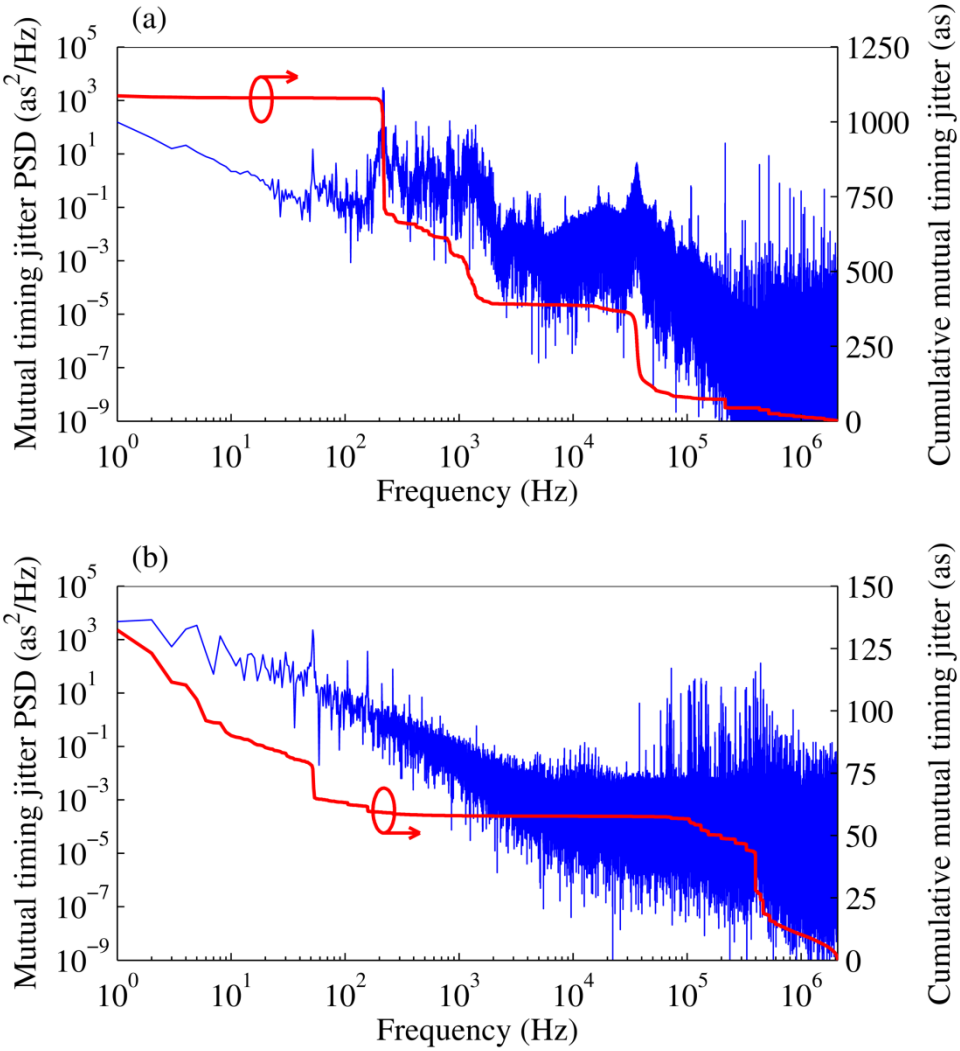

Figure 5.12. Mutual timing jitter between the pump and pump-signal pulses for both CEO locking loops (a) unlocked and (b) locked. When both CEO loops are locked the total cumulative mutual timing jitter is 137 as.

The results presented in this sub-section are critical for the establishing the potential of zero-offset CEO locking for broadband coherent pulse synthesis. The timing jitter between three pulses has been shown to be less than 140 as, indicating a high degree of timing stability. Table 5.1 listed the carrier periods of the candidate parent pulses, with carrier period decreasing linearly with wavelength. The timing jitter is much less than

the carrier period of the parent pulses, with the largest jitter of 137 attoseconds corresponding to just 9% of the carrier period of the pump-signal pulses. If the observation time is reduced to 1 ms then the jitter is reduced to 60 as, or 4% of the carrier period. These results show that even over an observation time of one second the timing jitter does not compromise coherent synthesis.

5.4.2 Contribution of RIN to timing jitter

From the results presented in the previous sub-section it is clear that the timing jitter between the candidate parent pulses for synthesis is sufficiently low to provide a train of phase coherent pulses. The timing jitter results shown are not a measurement of true timing jitter however, as the cross-correlation measurement technique is sensitive to contributions from both timing and intensity instabilities. It is therefore important to try and separate the true timing jitter from the intensity jitter.

The measured timing jitter measured as a cross-correlation τ_{meas} is composed of the true timing jitter τ_{jitter} and the RIN contribution to jitter τ_{RIN}. The two jitter components are a measure of the noise in the arrival time of the pulses, and so are added in quadrature to find the measured jitter in the form

$$\tau_{meas}^2 = \tau_{jitter}^2 + \tau_{RIN}^2 . \tag{5.16}$$

This allows the true timing jitter to be separated from the measured timing jitter.

RIN measurements were carried out for the three interfering pulses by blocking each arm of the interferometer individually and placing an appropriate bandpass filter in front of the photodiode. The RIN could then be recorded using the same method as previously outlined in Sub-section 5.2.2. Once data were recorded the voltage excursions due to relative intensity noise were rescaled so that the mean photodiode voltage matched that of the total timing jitter measurements. The voltage fluctuations could then be rescaled to establish the RIN contribution to timing jitter using the voltage-time conversion factor already established.

5.4.2.1 Results

The results of the RIN contribution to timing jitter for the pump, doubled-signal and pump-signal pulse are shown in Figure 5.13, Figure 5.14 and Figure 5.15 respectively.

A discussion of the frequency sources that contribute to the RIN of these pulses has already been made in Sub-section 5.2.2, and will not be repeated here.

When CEO locking was turned off the RIN contribution to the cumulative timing jitter was of the order of 50 as. When the CEO frequencies of both the pump and OPO were locked the timing jitter contribution dropped to 12 as for the OPO pulses (8 as for the pump).

Two primary conclusions can be drawn from these results. Firstly, the contribution of intensity fluctuations to the timing jitter is small, with the primary contribution arising from temporal drifts in the overlapping envelopes rather than intensity variations. Secondly, locking the CEO frequencies of the pump and OPO to establish both phase and wavelength stability remains the primary factor that determines the mutual timing jitter between the overlapping pulses.

Values for the true timing jitter can be calculated using Equation (5.16), yielding τ_{jitter} values of 131 as and 136 as between the pump and doubled-signal and the pump and pump-signal pulses respectively, showing that almost all the measured jitter is from true timing jitter.

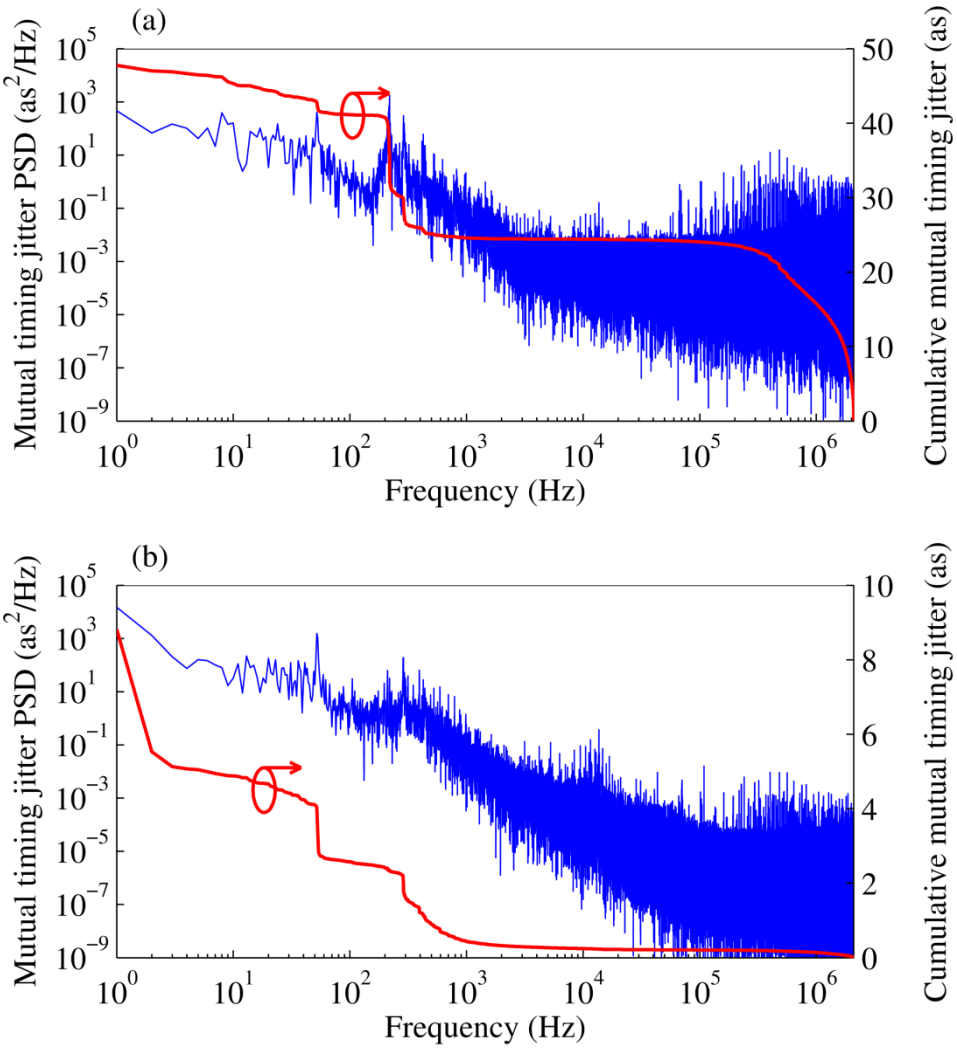

Figure 5.13. Contribution from the pump RIN to the mutual timing jitter for the CEO locking loops (a) unlocked and (b) locked.

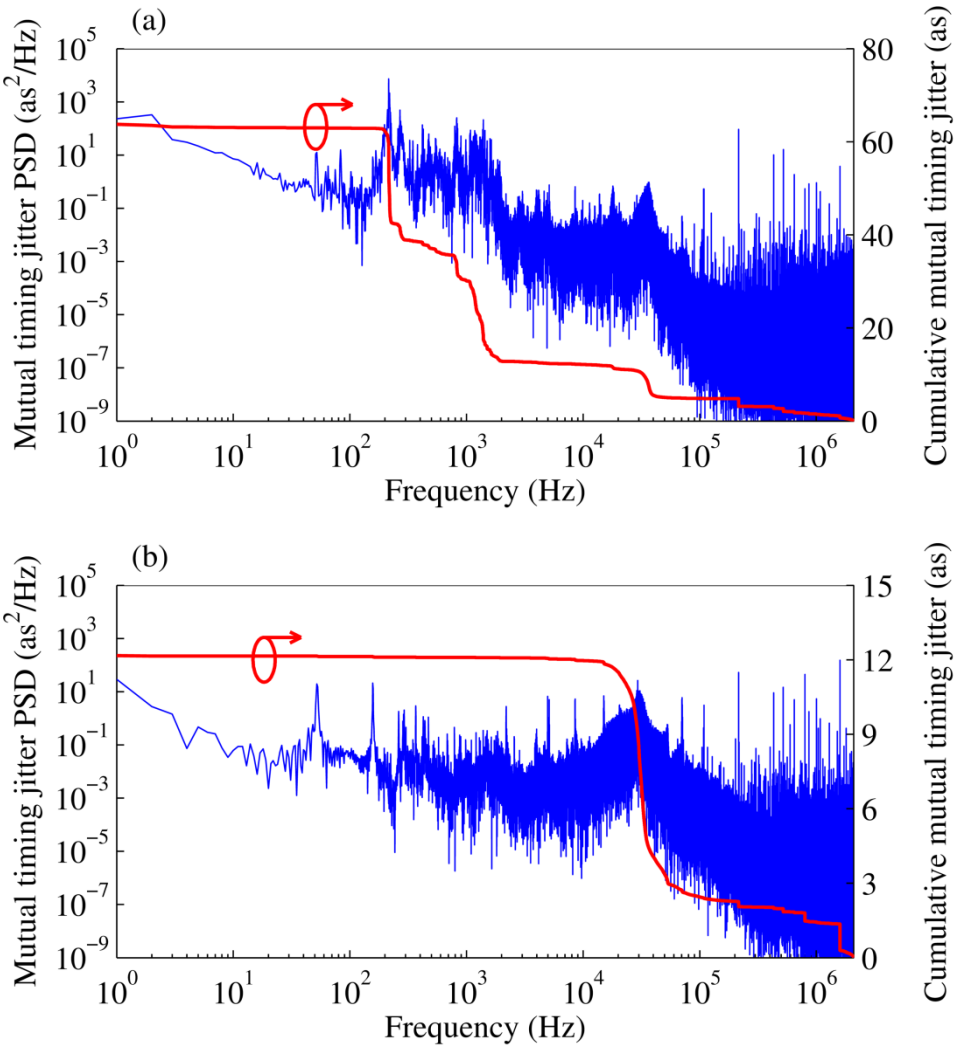

Figure 5.14. Contribution from the doubled-signal RIN to the mutual timing jitter for the CEO locking loops (a) unlocked and (b) locked.

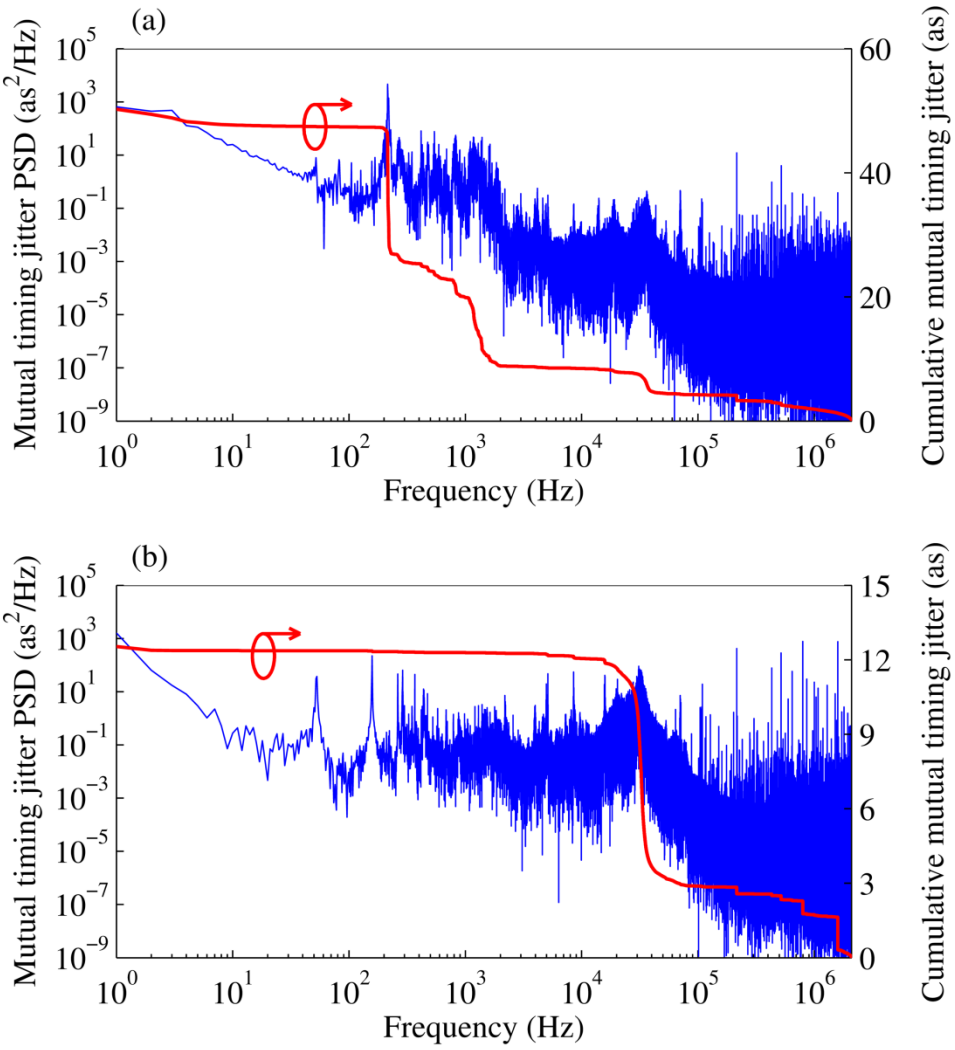

Figure 5.15. Contribution from the pump-signal RIN to the mutual timing jitter for the CEO locking loops (a) unlocked and (b) locked.

5.5 Conclusions

The concepts of relative intensity noise, phase noise and timing jitter were introduced in order to characterise the noise performance of the broadband, zero-offset frequency comb achieved in Chapter 4. Locking the CEO frequencies of both the pump and OPO was shown to greatly improve the overall stability of the system, and reduce the mutual timing jitter between the pump, doubled-signal and pump-signal pulses to less than 140 as. This result satisfies one of the critical prerequisites for coherent pulse synthesis.

5.6 References

1. W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipies: The Art of Scientific Computing*, Third edition (Cambridge University Press, 2007).

2. A. Yariv and P. Yeh, *Photonics: Optical Electronics in Modern Communications*, Sixth edition (Oxford University Press, 2006).

3. B. D. J. Jones and S. T. Cundiff, "Optical frequency electronics," IEEE Circuits Devices **19**, 28–35 (2003).

4. T. I. Ferreiro, J. Sun, and D. T. Reid, "Locking the carrier-envelope-offset frequency of an optical parametric oscillator without f-2f self-referencing," Opt. Lett. **35**, 1668–70 (2010).

5. T. I. Ferreiro, J. Sun, and D. T. Reid, "Frequency stability of a femtosecond optical parametric oscillator frequency comb," Opt. Express **19**, 24159–64 (2011).

6. F. W. Helbing, G. Steinmeyer, J. Stenger, H. R. Telle, and U. Keller, "Carrier-envelope-offset dynamics and stabilization of femtosecond pulses," Appl. Phys. B **74**, s35–s42 (2002).

7. F. W. Helbing, G. Steinmeyer, U. Keller, R. S. Windeler, J. Stenger, and H. R. Telle, "Carrier-envelope offset dynamics of mode-locked lasers," Opt. Lett. **27**, 194–6 (2002).

8. L.-S. Ma, R. Shelton, H. Kapteyn, M. Murnane, and J. Ye, "Sub-10-femtosecond active synchronization of two passively mode-locked Ti:sapphire oscillators," Phys. Rev. A **64**, 021802 (2001).

9. A. Leitenstorfer, C. Fürst, and A. Laubereau, "Widely tunable two-color mode-locked Ti:sapphire laser with pulse jitter of less than 2 fs," Opt. Lett. **20**, 916–8 (1995).

10. Z. Wei, Y. Kobayashi, and K. Torizuka, "Relative carrier-envelope phase dynamics between passively synchronized Ti:sapphire and Cr:forsterite lasers," Opt. Lett. **27**, 2121–3 (2002).

11. A. Bartels, S. A. Diddams, T. M. Ramond, and L. Hollberg, "Mode-locked laser pulse trains with subfemtosecond timing jitter synchronized to an optical reference oscillator," Opt. Lett. **28**, 663–5 (2003).

12. J. Kim, J. Chen, Z. Zhang, F. N. C. Wong, F. X. Kärtner, F. Loehl, and H. Schlarb, "Long-term femtosecond timing link stabilization using a single-crystal balanced cross correlator," Opt. Lett. **32**, 1044–6 (2007).

13. J. Kim, J. Chen, J. Cox, and F. X. Kärtner, "Attosecond-resolution timing jitter characterization of free-running mode-locked lasers," Opt. Lett. **32**, 3519–21 (2007).

14. T. R. Schibli, J. Kim, O. Kuzucu, J. T. Gopinath, S. N. Tandon, G. S. Petrich, L. A. Kolodziejski, J. G. Fujimoto, E. P. Ippen, and F. X. Kaertner, "Attosecond active synchronization of passively mode-locked lasers by balanced cross correlation," Opt. Lett. **28**, 947–9 (2003).

Chapter 6 - Practical considerations for coherent synthesis: pulse compression and group delay compensation

6.1 Introduction

In Chapter 5 the noise characteristics of the locking system were examined, showing that the mutual timing jitter between the selected parent pulses was sufficiently low to allow coherent pulse synthesis. To synthesise the shortest possible pulses it is necessary to compress the parent pulses and to control their temporal overlap so that they arrive at the target synthesis location at the same time. In this chapter, two methods of compressing the pulses and controlling their relative delay are introduced and results obtained from these systems evaluated and discussed, in the context of coherent synthesis.

6.2 Pulse compression using a quasi-common-path prism delay line

As discussed in Chapter 5, the selected parent pulses for synthesis were the pump (0.800 μm), doubled-signal (0.530 μm) and pump-signal (0.456 μm) pulses, all of which were chirped as they exited the OPO through a cavity folding mirror. Such chirp can affect the shape of the synthesised pulses dramatically, as illustrated in Figure 6.1. The upper panel in Figure 6.1 shows the intensity of a synthesised pulse composed of three parent pulses centred at the wavelengths mentioned above, assuming they are transform limited with 15-fs durations. The FWHM duration of the synthesised pulse is 2.14 fs, close to the single-cycle limit for a pulse generated over this bandwidth. The lower panel in Figure 6.1 shows the synthesis result when each of the parent pulses has acquired 150 fs^2 of group delay dispersion. While the frequency bandwidths of both the parent and daughter pulses are the same, the synthesised pulse has a much longer duration and also shows signs of pulse break-up.

It is clear from this example that the parent pulses must be compressed as much as possible before the shortest synthesised pulses can be generated, however a problem presents itself when implementing such compression. The wide bandwidth spanned by the parent pulses precludes the use of chirped mirrors or GTIs, and so a prism or diffraction grating approach is necessary. The power level of pump-signal pulses was lower than 30 mW and so a diffraction grating approach was not favourable due to the associated losses. Finally, the requirement of low mutual timing jitter between the pulses indicates that a common-path approach is needed for increased stability. Using

these criteria, a quasi-common-path prism delay line was constructed to both compress the pulses and compensate for their relative timing delay.

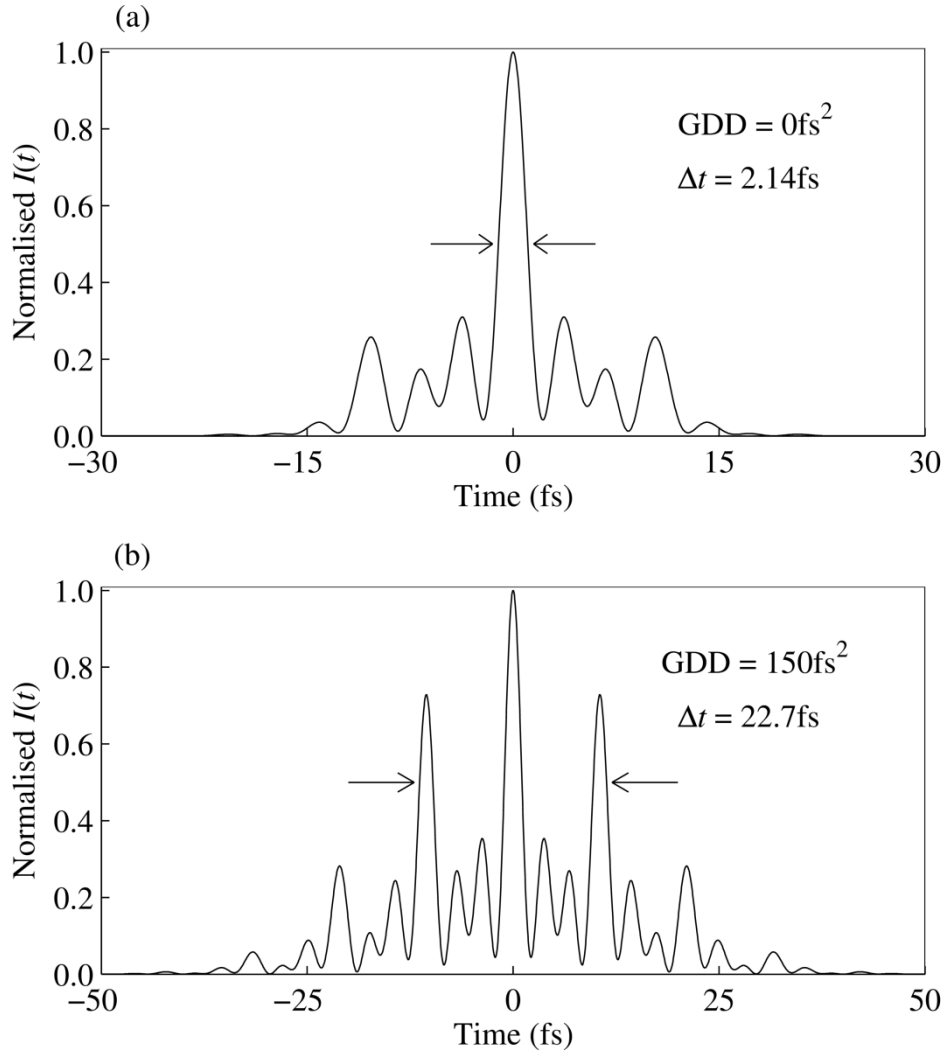

Figure 6.1. Modelled pulse synthesis results from 15-fs parent pulses centred at 0.800, 0.530 and 0.456 µm. Image (a) shows the resulting pulse intensity for chirp-free pulses. Image (b) shows the result when the parent pulses are chirped.

The prism delay line is shown in Figure 6.2. The collinear pump, doubled-signal and pump-signal beams exiting the OPO were steered into a fused silica prism with a 10-mm base. Each of the three beams was separately collected by a second fused silica prism with a 25-mm base. A mirror was placed after each prism pair to form a folded 4*f* system for each pulse. The folding mirrors were mounted on translation stages to compensate for the temporal walk-off between the pulses, which is discussed in Sub-section 6.2.2. The pulses were vertically displaced in order to be picked off and steered into the XFROG apparatus which was detailed in Chapter 3.

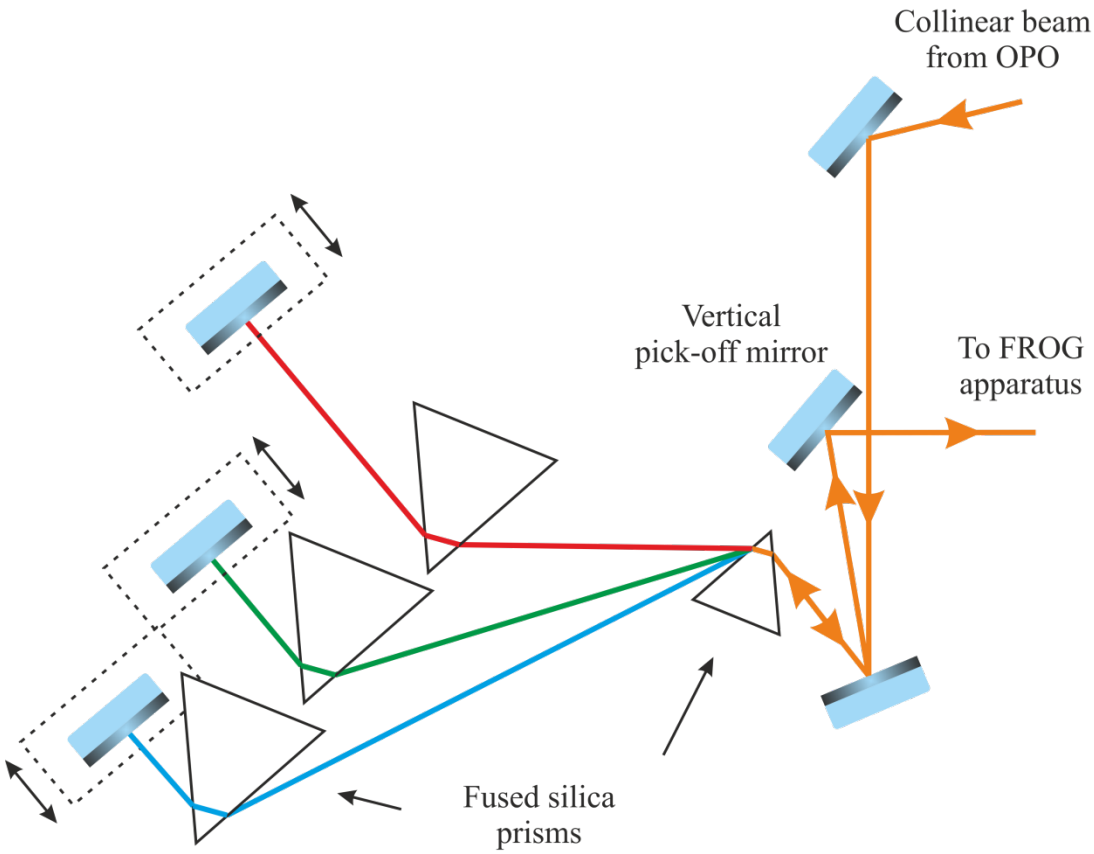

Figure 6.2. Schematic of the prism-based method for compressing the pump and visible OPO pulses. The folding mirror after each prism pair is mounted on a variable delay stage, as discussed in Sub-section 6.2.2.

6.2.1 Pulse compression results

The results from the quasi-common-path prism compressor are discussed in this sub-section. Each pulse was characterised individually in the XFROG apparatus, with the pump pulses employed as the XFROG gating pulse being compressed first. This ensured that the XFROG results for the doubled-signal and pump-signal pulses would not be skewed by the use of a chirped gating pulse.

6.2.1.1 Pump compression

The pump light exiting the OPO was steered directly into the XFROG apparatus in order to determine the initial level of chirp across the pulses. The results are shown in Figure 6.3. The interference fringes present in panel (a) were due to Type I phasematching in the XFROG BBO crystal which generated 0.400 μm second harmonic background light. The results of the XFROG retrieval algorithm, discussed in Chapters 2 and 3, are shown in panels (b) – (d). The FHWM bandwidth of the retrieved pulses was 297 fs, indicating that they are highly chirped.

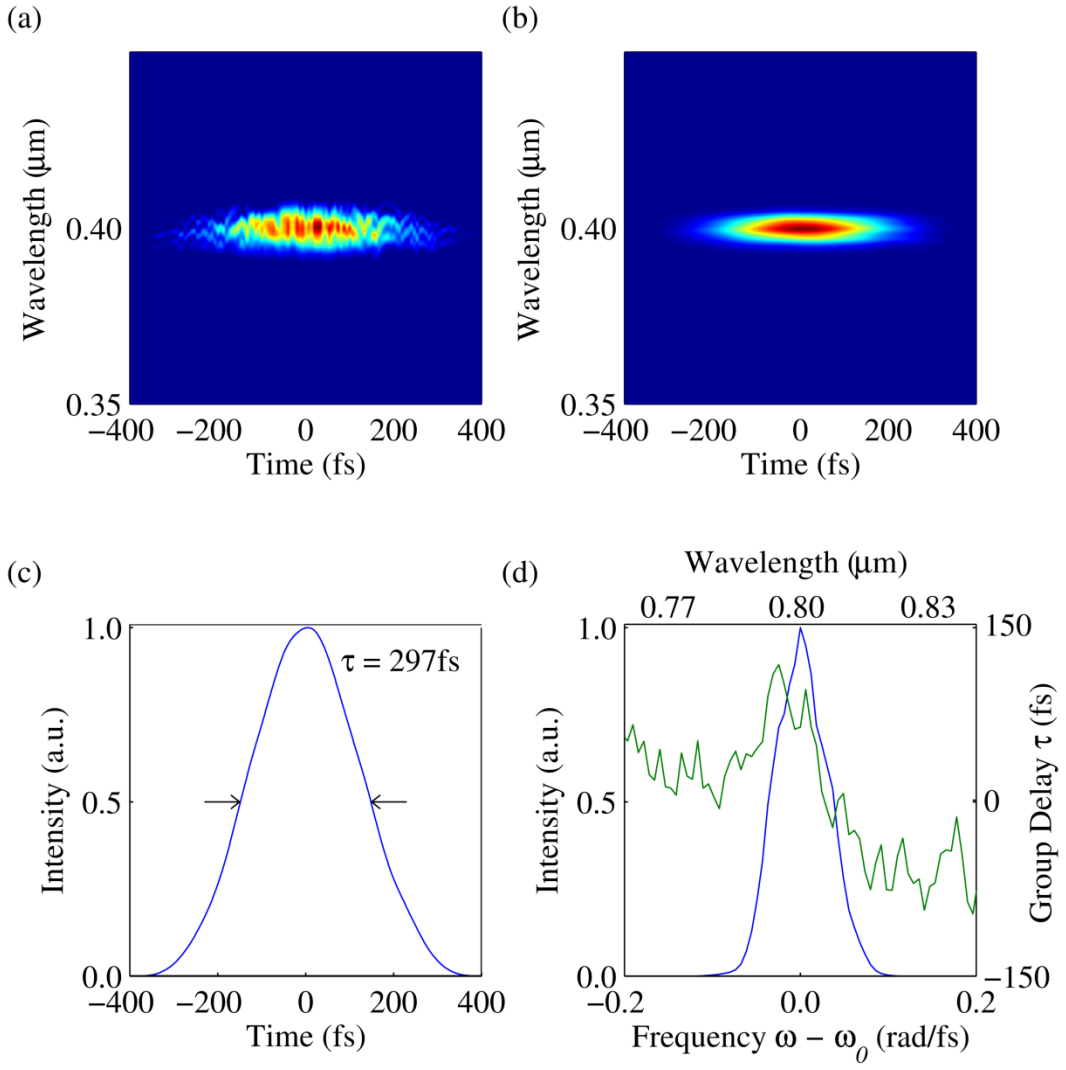

Figure 6.3. XFROG retrieval of the chirped pump pulses. (a) Raw trace; (b) retrieved trace; (c) temporal pulse profile; (d) spectral pulse profile.

The pump beam was then steered into the prism compressor with an initial prism separation of 38 cm, with the returning beam steered back into the XFROG apparatus. The measured XFROG trace was analysed and the prism separation increased until the minimum pulse duration was achieved. A simple metric was used to determine the pulse duration without the need to retrieve the XFROG trace. Summing the FROG trace over the frequency range of the spectrometer produced an intensity profile over time – an intensity autocorrelation. The FHWM span was calculated from intensity autocorrelation and a graph plotted of the resulting span against prism separation, shown in Figure 6.4. Fitting a quadratic to this data produced a curve allowed the optimal prism separation to be determined, however it was also evident by eye as to which separation yielded the shortest pulses. This method was repeated for the doubled-signal and pump-signal pulses, however the results are omitted here.

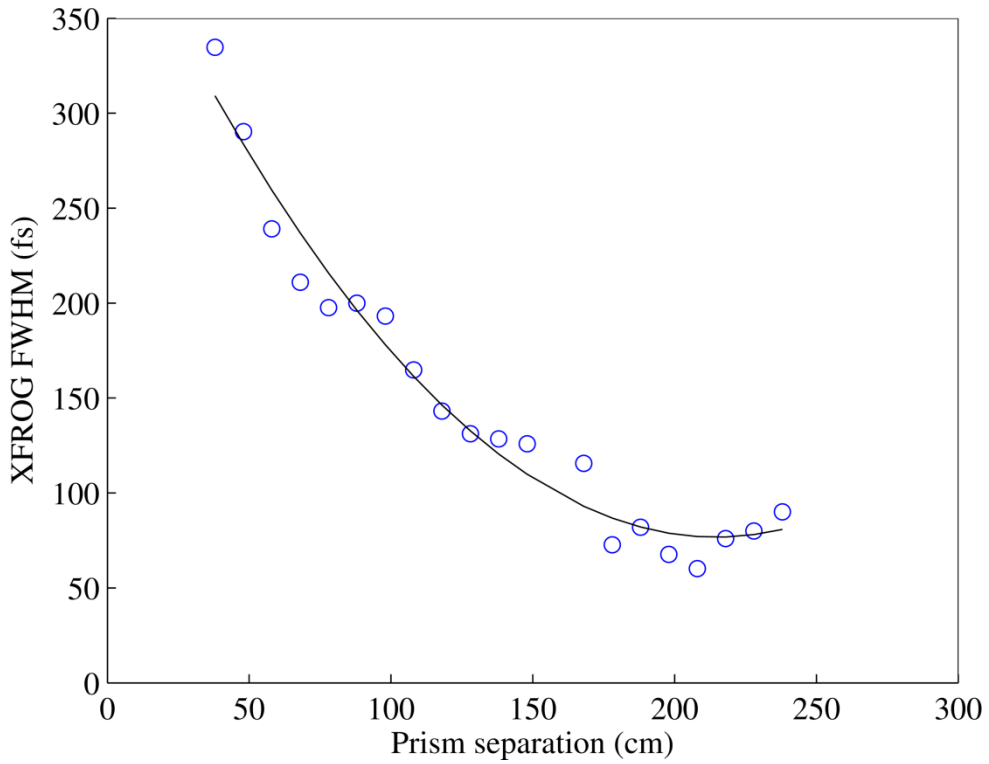

Figure 6.4. FHWM duration of the pump pulses as the prism separation was increased. A fit was used for a rough guide, with the pump pulses being compressed to 38 fs at 200 cm separation.

The results showed that the minimum pump pulse duration was achieved with a prism separation of 208 cm. The results from the XFROG trace with 208 cm prism separation are shown in Figure 6.5. An unconstrained XFROG retrieval algorithm was used to recover the pulse duration rather than a standard FROG retrieval. The standard algorithm assumes that the apparatus is balanced in terms of dispersion and spectral profile; however in these experiments the dichroic beam splitter used to separate the pump and visible wavelengths only reflects as small portion of the pump beam for testing. Different mirror coatings are also used in each arm of the interferometer, and so an XFROG algorithm is more suited for retrieval of the pump pulses, leading to an asymmetry in the trace shown in Figure 6.5.

The pump pulse duration was compressed to 38 fs; while this result is not transform-limited (the spectral bandwidth supports pulse durations of approximately 20 fs) it is the best that could be achieved with the experimental design. A prism separation of 2 metres presented its own problems as the returned beam had diffracted sufficiently as to double in diameter; however the aperture of the common-path prism was wide enough to ensure that the prism compressor was more than 90% efficient.

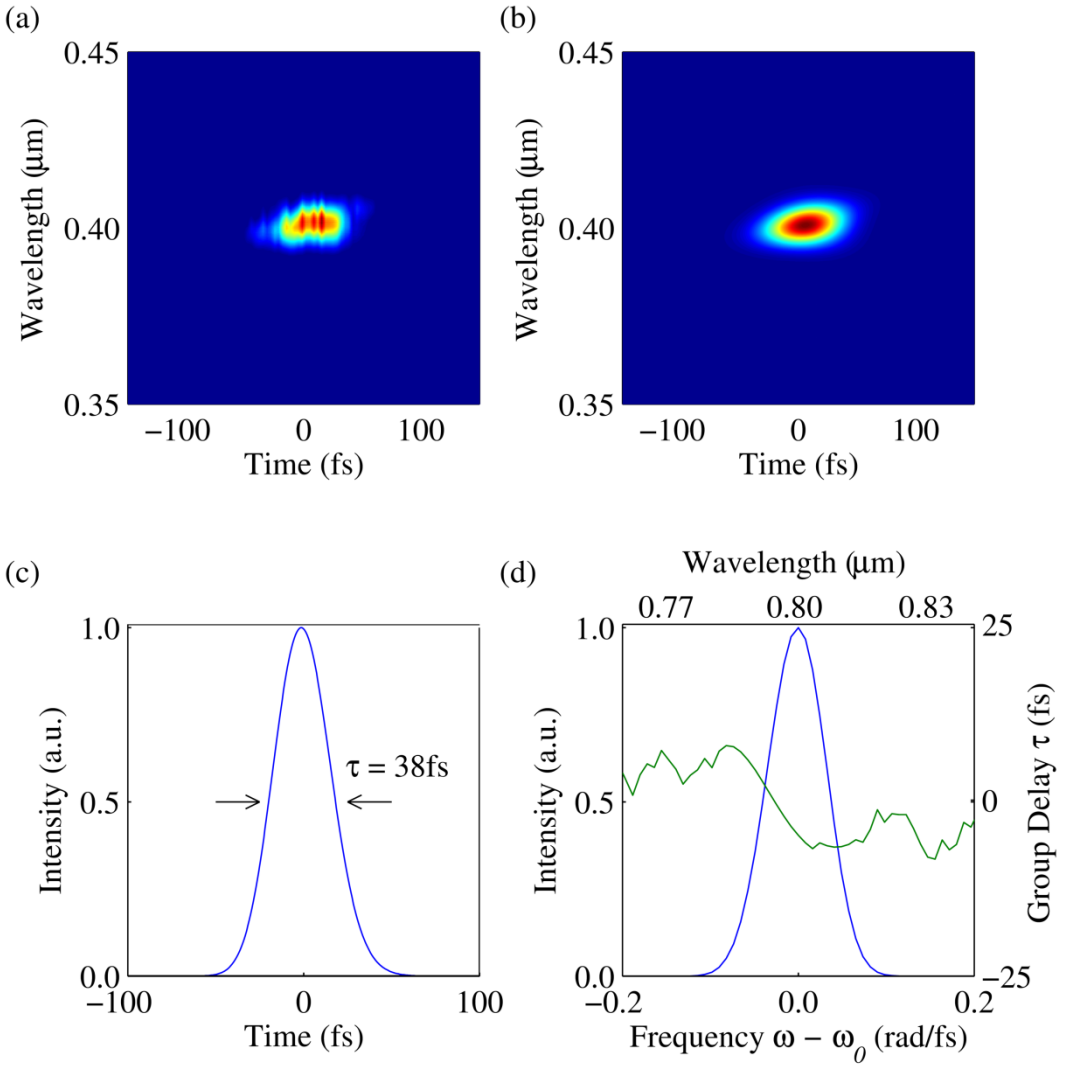

Figure 6.5. Retrieval results from the XFROG measurement on the compressed pump pulses. (a) Raw trace; (b) retrieved trace; (c) temporal pulse profile; (d) spectral pulse profile.

6.2.1.2 Doubled-signal compression

The doubled-signal pulses were compressed in an identical manner to the pump pulses. The prism separation was increased until the minimum pulse duration was obtained. The optimal separation was found to be 150 cm, and the XFROG results are shown in Figure 6.6. The pulses were compressed to 100 fs, slightly longer than the 80 fs transform limited duration.

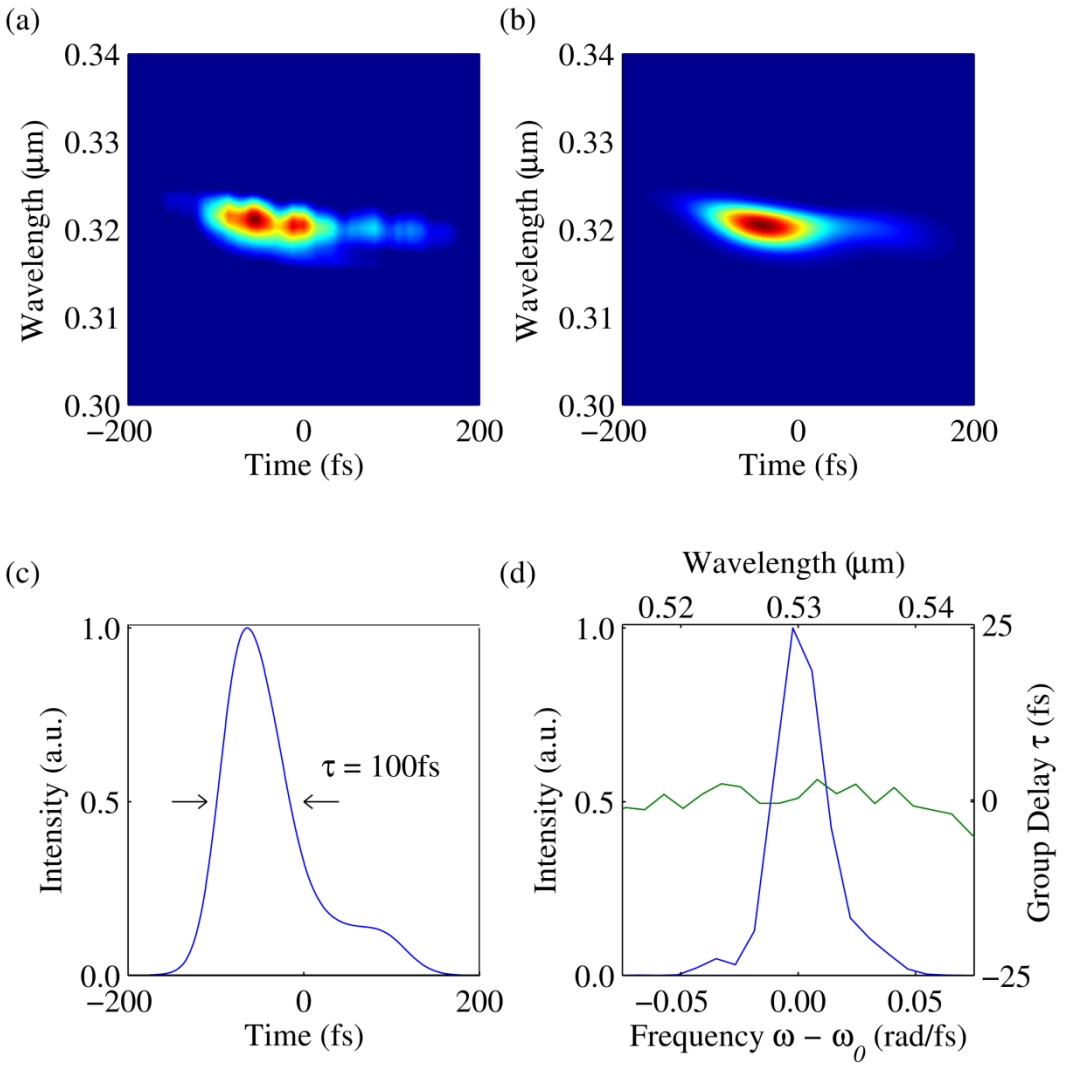

Figure 6.6. Retrieval results from the XFROG measurement on the compressed SHG signal pulses. (a) Raw trace; (b) retrieved trace; (c) temporal pulse profile; (d) spectral pulse profile.

6.2.1.3 Pump-signal compression

The pump-signal pulses were compressed in the same manner detailed above. The prism separation was increased until the minimum pulse duration was obtained. The optimal separation was found to be 120 cm, and the XFROG results are shown in Figure 6.6. The pulses were compressed to 115 fs, indicating that they were still slightly chirped.

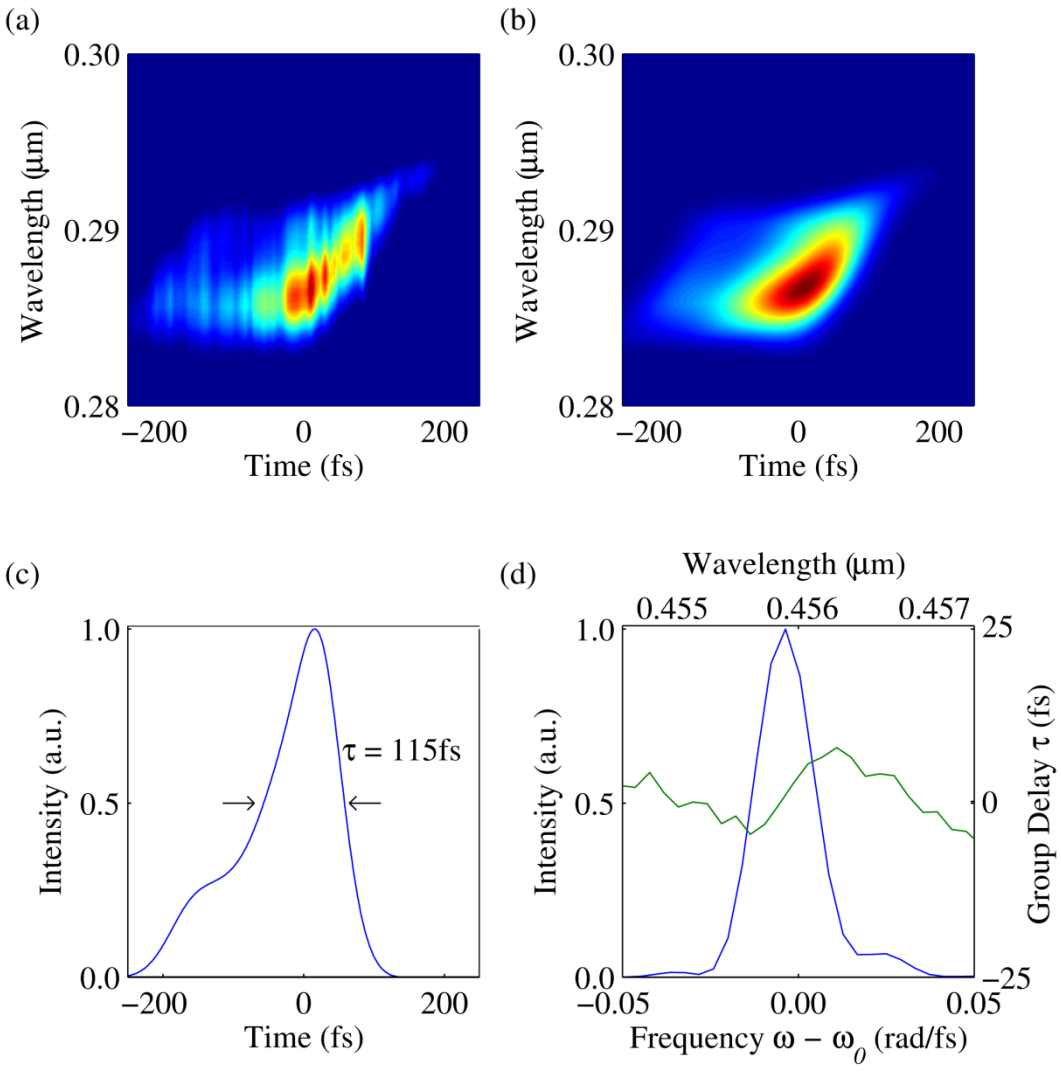

Figure 6.7. Retrieval results from the XFROG measurement on the compressed pump-signal pulses. (a) Raw trace; (b) retrieved trace; (c) temporal pulse profile; (d) spectral pulse profile.

6.2.2 Relative delay manipulation

Once the pulses had been compressed it was then necessary to compensate for walk off due to the group delay dispersion of each pulse. As the pulses propagate through different optical materials they experience a frequency-dependent group delay. Over short distances this group delay dispersion gives rise to timing jitter, but over longer distances the pulses walk away from each other and ultimately lose their temporal overlap. The temporal walk off between the pulses was compensated for by adjusting the distance between each second prism and the folding mirror in each prism compressor, as illustrated in Figure 6.8. This adjustment was used to bring the pulses back into overlap in the crystal of the XFROG apparatus; this is an ideal candidate position for pulse synthesis as temporal overlap can be confirmed before removing the BBO crystal and using the same focal plane for future experiments.

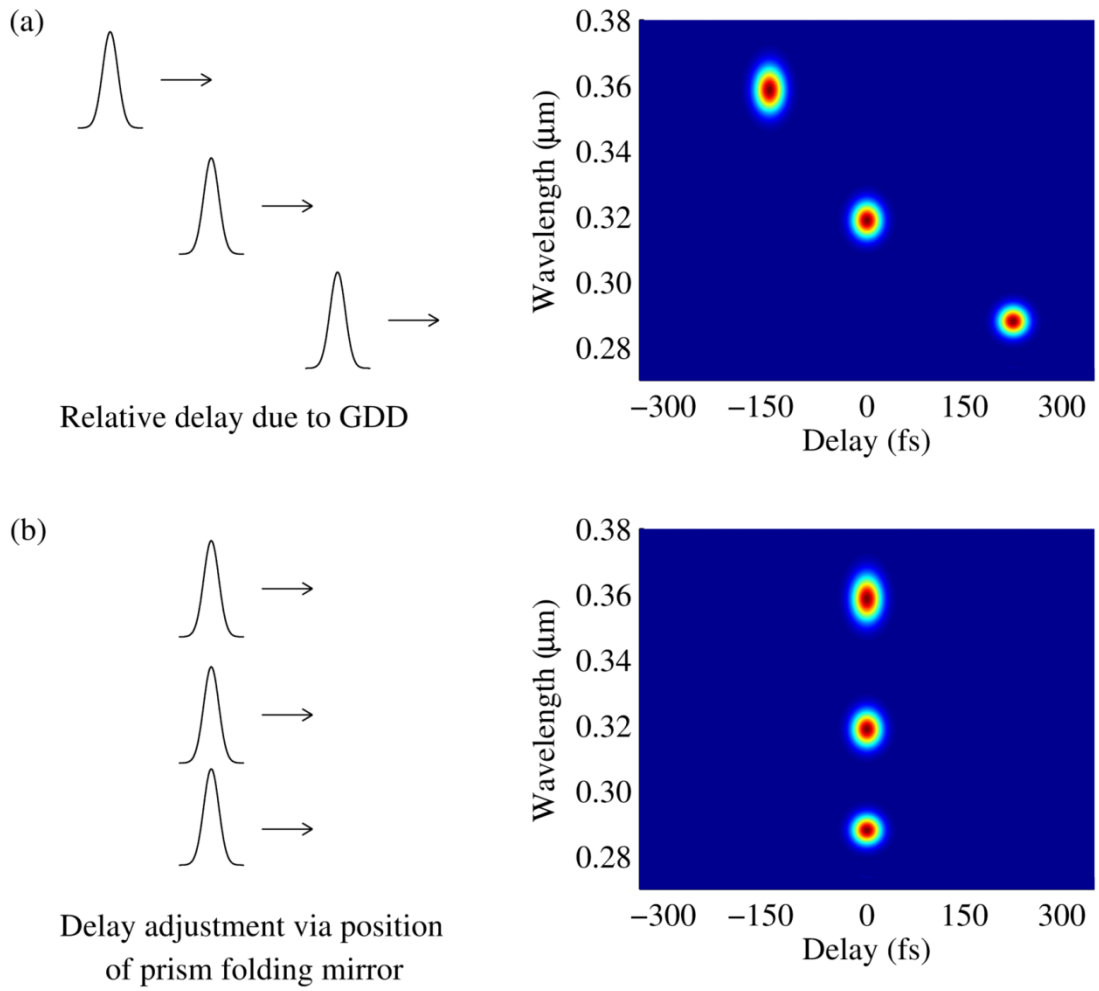

Figure 6.8. Illustration of how XFROG can be used to obtain temporal overlap. The position of the prism folding mirror can be adjusted to bring the pulses into temporal alignment.

6.2.2.1 Relative delay adjustment results

The group delay dispersion between the pulses was roughly compensated for by correctly positioning the folding mirrors after each second prism. The path lengths between the first prism and the folding mirror were made equal for each parent pulse. As the prism separation in the pump prism compressor was over 2 metres, the distance between the second prism and the folding mirror was quite large for both the doubled-signal and pump-signal delay lines. After a rough delay adjustment had been made all three pulses were steered into the XFROG apparatus for measurement, with the results shown in Figure 6.9.

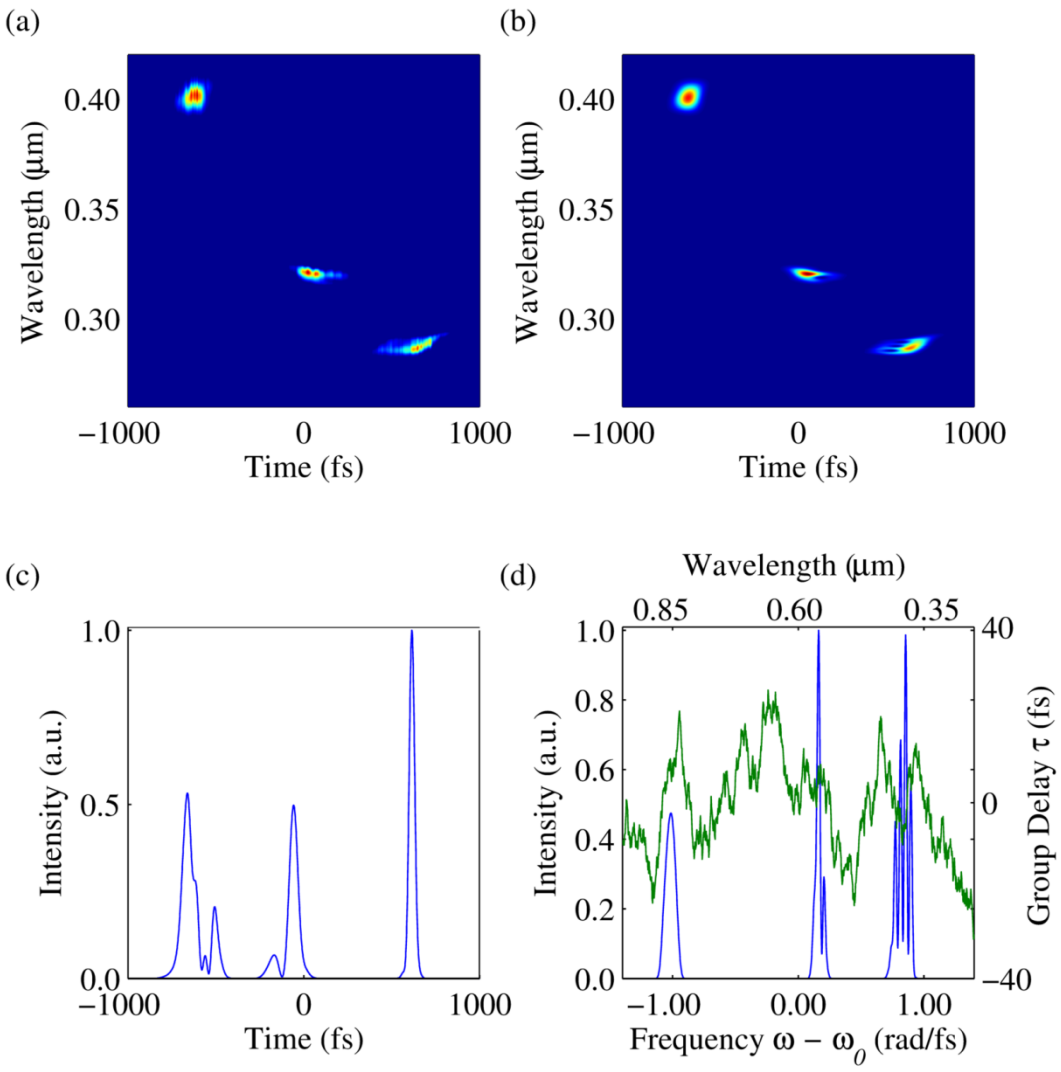

Figure 6.9. XFROG retrieval results from the pulses when they are temporally dispersed. (a) Raw trace; (b) retrieved trace; (c) temporal pulse profile; (d) spectral pulse profile.

As was demonstrated in Chapter 3, the XFROG retrieval algorithm is robust enough to recover the relative delays between the pulses. The leading and trailing pulses are

delayed by approximately 1.5 ps, roughly 0.5 mm propagation distance in air, which shows that even the most accurate of path length measurements will still contain some error due to refractive index differences.

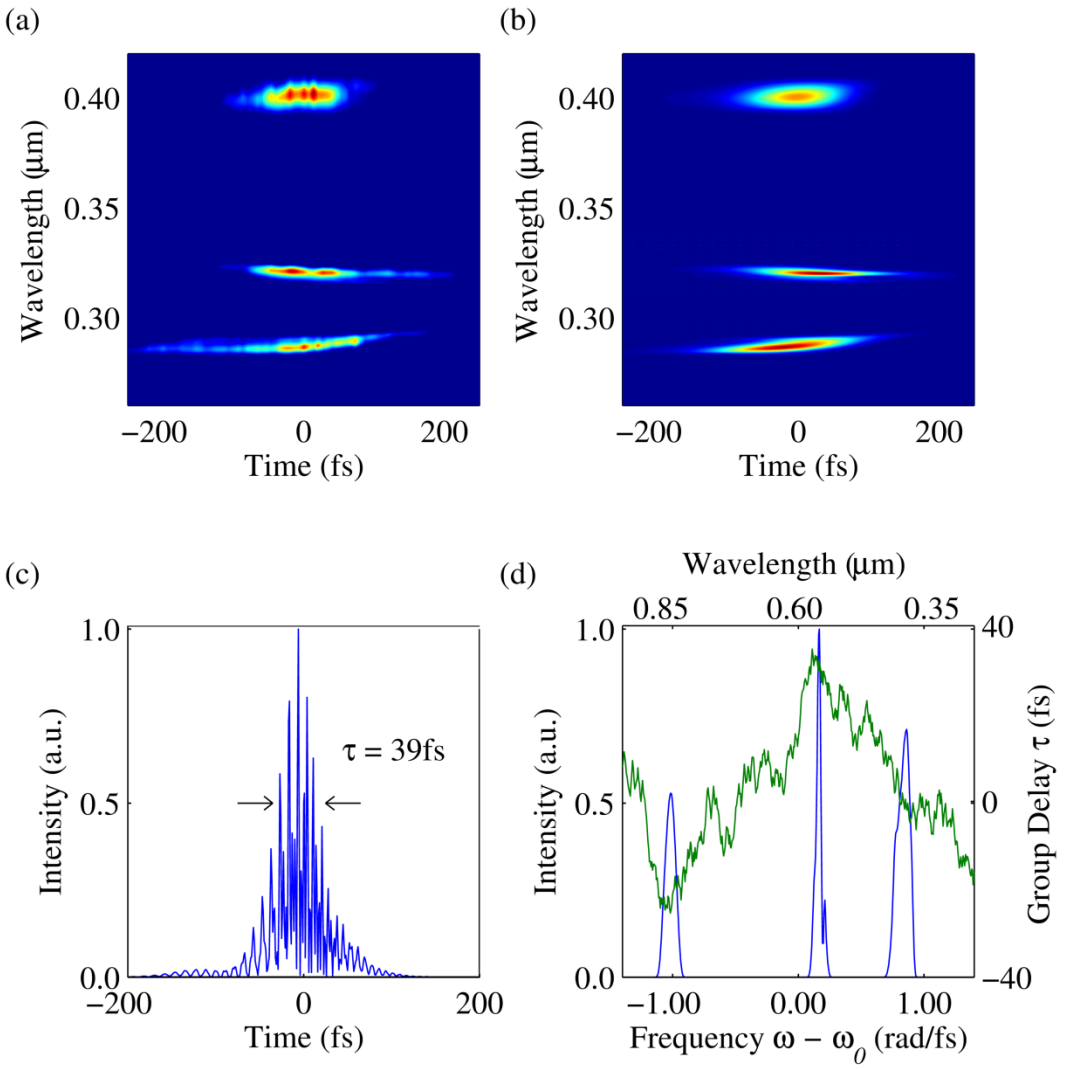

Figure 6.10. XFROG retrieval results from the pulses after delay compensation. (a) Raw trace; (b) retrieved trace; (c) temporal pulse profile; (d) spectral pulse profile.

The position of the doubled-signal and pump-signal folding mirrors in the prism compressor were adjusted in order to temporally align the pulses. This was achieved practically by moving one mirror and acquiring an XFROG trace, a laborious process that could possibly be improved by monitoring the UV spectrometer in real time. An XFROG measurement showing the pulses temporally aligned is shown in Figure 6.10. The XFROG retrieval algorithm has successfully retrieved multiple, spectrally distinct pulses. The retrieved temporal pulse profile shown in panel (c) shows temporal interference fringes, indicating that, as far as the algorithm is concerned, the pulses are indeed temporally overlapped. The FHWM duration of the temporal pulse is 39 fs.

6.2.2.2 Comparison of results with model

The group delay between the pulses, shown in Figure 6.10 (d) in green, is derived from the absolute phase differences between the pulses retrieved by the XFROG algorithm. The XFROG trace, however, is invariant to the absolute phase, or the phase between the individual pulses, and so a true understanding of how the pulses will interfere in time cannot be determined from the XFROG trace alone. It is possible however to isolate the retrieved information for the individual pulses and predict the optimal synthesis result using a minimisation function.

Using MATLAB, the relative delays between the pulses were adjusted by small amounts and the amplitudes of their electric fields weighted relative to the pump pulse. Both of these factors can be adjusted experimentally using a pulse shaper capable of shaping both amplitude and phase, such as the liquid-crystal spatial light modulator discussed in the next section. It is important to note that both the delay and amplitude of a pulse can be altered without changing the pulse shape or duration. The relative delays and amplitudes of the doubled-signal and pump-signal pulses were used as the input variables of a minimisation function in order to improve the contrast of the pulses produced by coherent synthesis. The result of the minimisation function is shown in Figure 6.11 (c).

Improving the temporal overlap and appropriately weighting the amplitudes of the pulses has resulted a train of five high-contrast synthesised pulses with durations of 2.15 fs, close to the single-cycle limit of the highest-frequency parent pulse (pump-signal, 1.5 fs). This result shows that, with appropriate delay adjustment and amplitude shaping, coherent synthesis is possible using the parent pulses measured in the previous sub-section.

It is instructive to determine the effect of timing jitter and CEO phase noise on the synthesised pulse; if the timing jitter is significant then the synthesised pulse will be washed out after only a short observation time. From Chapter 5, the largest timing jitter measured between the parent pulses over a 1 second observation window was 137 as; this includes contributions from relative intensity noise. Similarly, the largest CEO phase noise was 0.18 radians. Both timing jitter and phase can be added or subtracted to each pulse as a delay term. Calculating the contribution to noise as a fraction of an optical cycle, it is clear that the timing jitter is the largest contributor to the noise at 0.1 cycles, while the phase noise contribution is 0.02 cycles. Adding 150 as of delay to one

visible pulse and subtracting the same from the other is a worse-case scenario, and will determine how synthesis is affected by the noise measured previously.

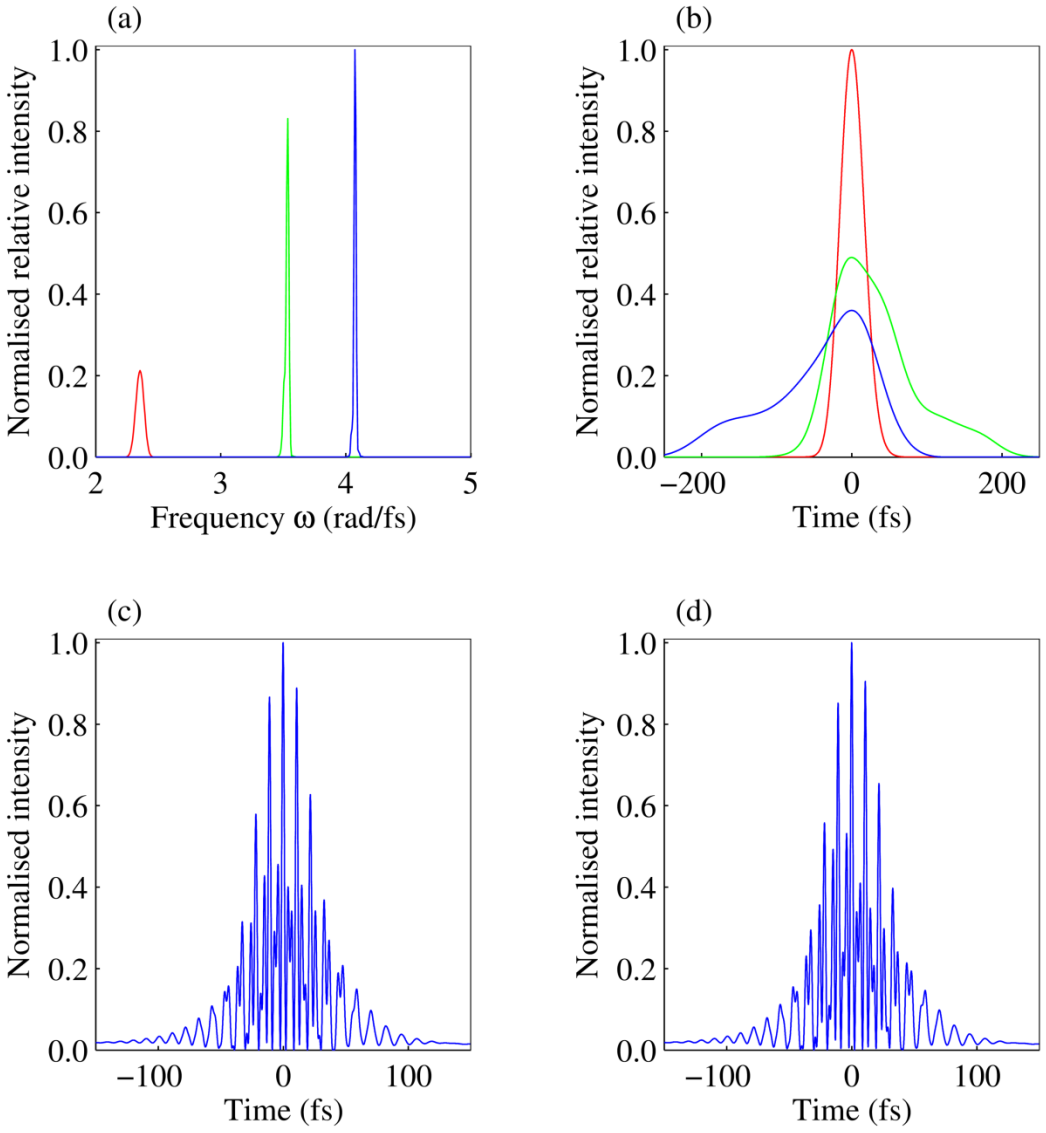

Figure 6.11. Results of the minimisation function used to determine the optimal coherently synthesised pulse that can be produced from the three parent pulses. (a) Spectra of the three pulses; (b) their overlapping individual fields; (c) optimal synthesis result, showing a train of 5 synthesised pulse of 2.15 fs duration; (d) synthesis result with added timing jitter of 150 as.

The result is shown in Figure 6.11 (d). The change in delay has reduced the fringe contrast, however the durations of the synthesised few-cycle pulses remains unaffected. This shows that the synthesised pulses remain coherent and experimentally accessible over a 1 second observation time, an important result in the practical application of synthesised pulses.

6.3 A spatial light modulator for a common-path approach to pulse compression

The quasi-common-path prism delay line was successful in both compressing the chosen parent pulses and compensating for their temporal walk-off; however there are still areas in which improvements could be made. The quasi-common-path approach introduces additional timing jitter that would be rejected in a truly common-path compressor. Precise positional control of the folding mirrors would be required to improve the temporal overlap, most likely through the use of long-travel piezoelectric transducers attached to the positioning stages. More crucially however the pulses cannot be compressed to their transform-limited durations without additional dispersion compensating optics.

An approach that would address all these problems would be to use an ultrafast pulse shaper to control the spectral phase of the pulses. Ultrafast pulse shapers have been used for many years to enable the generation of spectral and temporal profiles tailored to the need of specific experiments [1]. Examples include liquid-crystal spatial light modulators (LC-SLMs) [2,3], acousto-optic programmable dispersive filters (AOPDFs) [4], and deformable mirrors [5], which have been used to shape pulses for coherent control [6], ultraviolet to mid-infrared spectroscopy [7,8] and Raman frequency-comb compression [9].

A suitably characterised broadband, high-resolution LC-SLM can be used to manage the group delay dispersion of the parent pulses in a way that maximises their mutual temporal overlap, while simultaneously compressing each pulse to its transform-limited duration.

In this section the development of a new technique for characterising the broadband phase and voltage characteristics of an LC-SLM is described, in preparation for its deployment in a pulse-shaping experiment. Due to a failure in the LC-SLM firmware, it was not possible ultimately to test the LC-SLM in a true phase-shaping mode, however the data obtained can be used to estimate the range of shaping outcomes achievable.

6.3.1 Liquid-crystal spatial light modulators

An LC-SLM is composed of a thin layer of liquid-crystal sandwiched between two layers of glass. The inside surface of each piece of glass is coated with a thin, transparent, electrically conducting film. Once piece of film is patterned into a number of separate electrodes which act as pixels; the film on the other glass surface is

unpatterned, serving as a ground plane. This is illustrated in Figure 6.12 (a). The pixel size and number varies between LC-SLM designs, however in all cases the distance between pixels is kept to the few-μm level to prevent unwanted diffraction effects. LC-SLMs can operate in reflective or transmissive geometries with appropriate anti- or high-reflectivity coatings applied to the glass layers.

The liquid-crystals within a spatial light modulator are typically uniaxial-nematic; they are rod shaped and possess long-range directional order with the preferred direction of alignment being with their long axes roughly parallel. The liquid-crystals can be easily aligned by an external electric field, with the long axes aligning parallel to the applied field, as shown in Figure 6.12 (b).

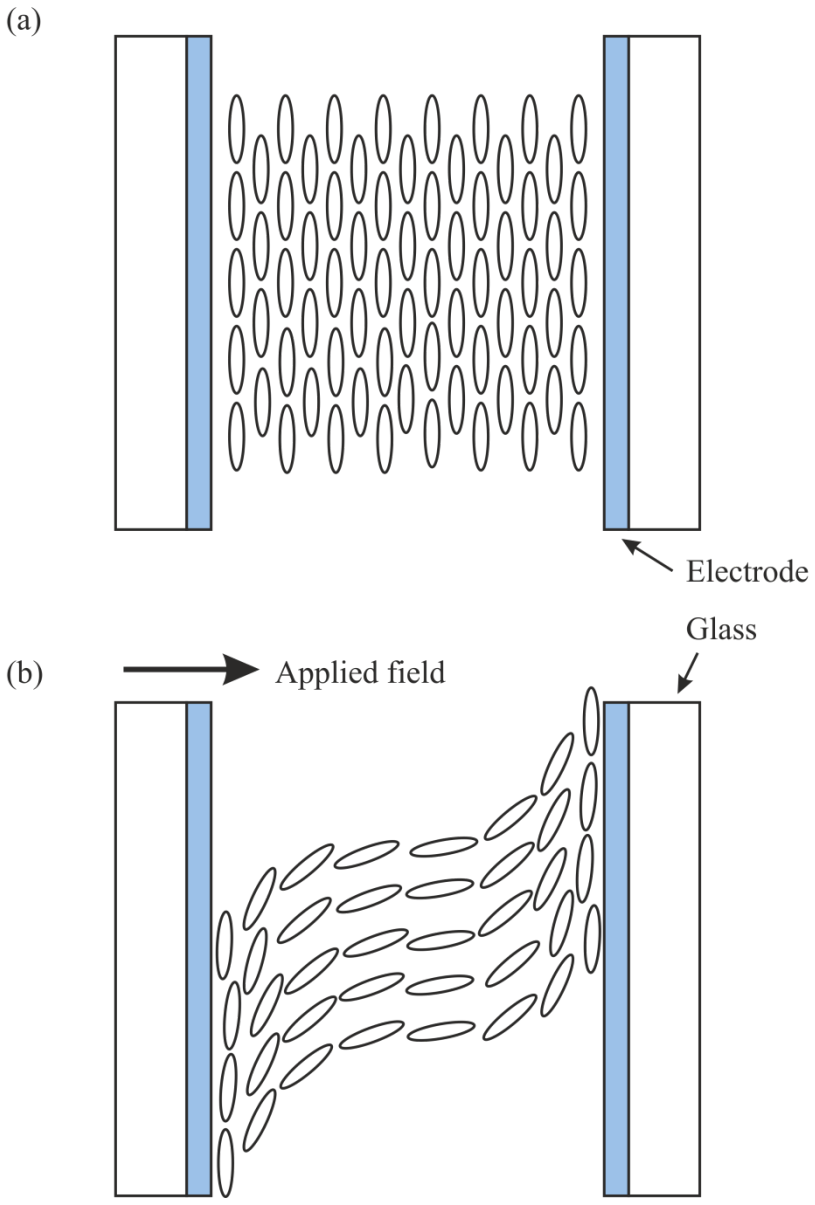

Figure 6.12. Side view of a liquid-crystal pixel with (a) no applied electric field, and (b) an electric field applied in the longitudinal direction. Adapted from [10].

LC-SLMs are of particular interest to ultrafast pulse shaping as the liquid-crystals are birefringent. As an electric field is applied across the pixel the orientation of the crystals will change; incoming light of the correct polarisation will observe a changing refractive index as the crystals rotates and will undergo a corresponding phase shift. In the case of ultrafast pulse shaping this voltage-dependent phase change can be exploited to compensate for phase differences across the bandwidth of the pulse, using the LC-SLM as a compressor. An attainable phase change of at least 2π across the bandwidth of the incoming pulse is required for complete phase control, which determines the minimum thickness of the liquid-crystal layer.

6.3.2 Characterisation of a broadband linear LC-SLM

While the chosen set of parent pulses for coherent synthesis is composed of only three pulses, future opportunities may present themselves that will allow the pump and all four visible OPO outputs to be parent pulses. Implementing programmable phase control over such a wide frequency bandwidth presents specific challenges associated with calibrating the phase response of the LC-SLM, which is characteristically nonlinear with both and phase and applied voltage. The pulse shaper used in the work carried out in this thesis was a folded 4*f* system (Figure 6.13) in which the incident and returning beams were coupled in and out of the system by a prism, with a small relative vertical offset to allow the shaped pulses to be isolated and analysed. The collinear pump and visible beams from the OPO were angularly dispersed with a fused silica prism placed at the front focal plane of an achromatic doublet [11]. A prism-based shaper was preferred over a design using a diffraction grating because it offers greater double-pass efficiency, considerably lower wavelength-dependent loss, and avoids problems with overlapping diffraction orders when the bandwidth approaches an octave. The material dispersion introduced by the prism can be compensated for either before or after the shaper, or by choosing a prism with low intrinsic material dispersion such as the fused silica prism used here.

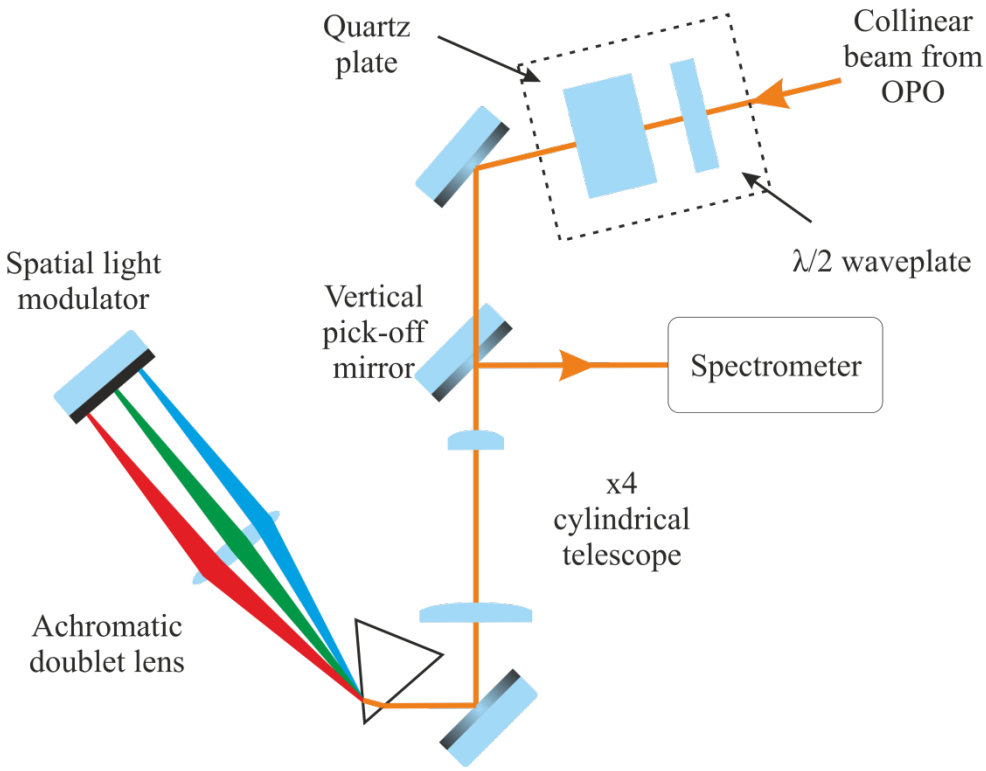

Figure 6.13. Schematic of the LC-SLM characterisation apparatus. The optics in the dashed box were used for the voltage-to-phase calibration only.

A 12,288-pixel reflective liquid-crystal-on-silicon LC-SLM (*12.288 Linear Series, Boulder Nonlinear Systems*) was situated at the back focal plane of the lens. The LC-SLM had an array size of 19.66 mm with a 1.6-µm pixel pitch, comprising a 1.0-µm electrode and a 06-µm gap. A custom dielectric high-reflectivity coating was applied to the backplane of the LC-SLM, with the reflectivity optimised to cover the spectral regions occupied by the doubled-pump, pump-signal, doubled-signal and pump-idler wavelengths. The dielectric coating prevented diffraction from the inter-pixel regions of the backplane and increased the optical efficiency of the device.

Notably, the LC-SLM used in this work had a much higher pixel density than many of the devices commonly used for ultrafast pulse shaping, which typically possess 640 pixels or less. Normally LC-SLMs require the spectral focus to match the pixel period, in order to maximise the spectral resolution. By contrast the high pixel density of the LC-SLM used here meant that each spectral focus included a large number of pixels, providing the capability for high-resolution phase and amplitude control. The spectral foci at 0.400 µm and 0.800 µm were 170 µm and 340 µm respectively, and therefore covered at least 100 pixels.

The purpose of the LC-SLM calibration is to produce a function that can provide the optical phase written by the shaper as a function of both wavelength and applied voltage. This procedure requires separate functional maps of frequency-to-pixel number and voltage-to-phase.

6.3.2.1 Spatial calibration – frequency-to-pixel number

The spatial distribution of the frequencies across the surface of the LC-SLM is not linear, but is determined by the material dispersion of the fused silica prism. The prism was oriented for minimum deviation in the focal plane of the LC-SLM. The broad bandwidth of the dispersed spectrum provides a slightly different minimum-deviation angle for every wavelength; however each angle produces a similar focal pattern on the LC-SLM surface and so the exact prism angle is not critical. The wavelength map across the device was revealed by writing a phase step about a given pixel. As each wavelength is focused onto more than one pixel (as determined by the resolution of the prism "spectrometer") this phase step produces a discontinuity at the corresponding wavelength in the reflected spectrum. Using MATLAB control of the LC-SLM the phase step was scanned across the array of pixels, with the reflected visible spectra measured on each occasion with a spectrometer (*USB4000, Ocean Optics*). Fitting a quadratic curve to these data points provided the pixel-to-frequency relationship, shown in Figure 6.14 as a black dashed line. The relationship agreed well with the theoretical behaviour calculated from the Sellmeier equation for fused silica [12], shown in Figure 6.14 as a red line.

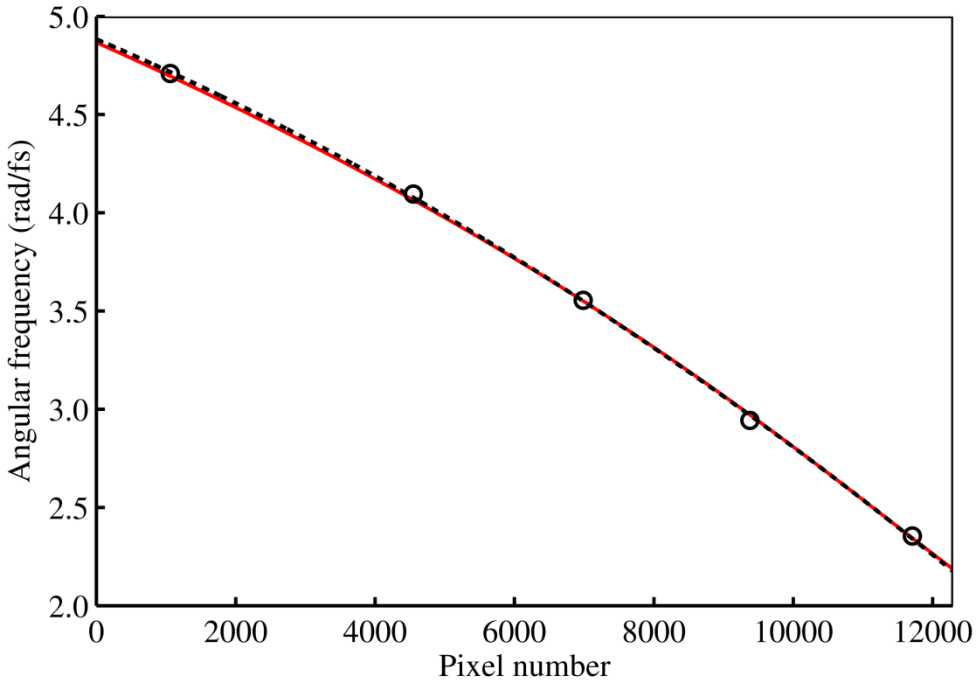

Figure 6.14. Spatial dispersion of the pump and visible OPO frequencies over the LC-SLM. The solid red curve indicates the calculated pixel-to-frequency mapping obtained from the fused silica Sellmeier equation. The measured values are given by the black circles, while the dashed black line is the fitted curve.

6.3.2.2 Phase calibration – voltage-to-phase

A phase calibration as a function of applied voltage was supplied by Boulder Nonlinear Systems for one optical frequency, providing a smooth phase change over 2π radians as the applied voltage was varied from its minimum to maximum value. If this calibration function were applied to every pixel it would provide an incorrect phase response at every other frequency. As a result of this, operation of the pulse shaper across a broad bandwidth requires a separate calibration of the voltage-to-phase relationship for every incident frequency. Combining this with the frequency-to-pixel number relationship determined in Sub-section 6.3.2.1 would enable deterministic phase shaping of pulses across the entire ensemble of pump and visible OPO pulses.

The phase change was calculated using an inline common-path spectral interferometry technique [13]. The horizontally polarised light exiting the OPO was rotated to 45° and double-passed through a quartz crystal whose optic axis was horizontal (parallel to the optical bench). The birefringence of the quartz crystal splits an input pulse into two time delayed replicas with orthogonal polarisations. Only the phase of the horizontally polarised light is affected by the orientation of the liquid-crystals in the LC-SLM, leaving the vertically polarised light as a common-mode reference. After the shaper the two polarisations were resolved with a polariser oriented at 45° and the resulting fringes

observed on a spectrometer. The voltage applied to the LC-SLM pixels was increased from zero to their maximum value in 256 linear steps and the changing interference pattern recorded.

Previous characterisations of pulse shapers which were operated in an open-loop configuration [14] used a Fourier-sideband filtering algorithm [15] to calculate the imparted phase. Due to the broadband nature of the individual pulses used in the LC-SLM characterisation, the number of fringes required to achieve distinct Fourier sidebands would require the pulse to pass through several cm of quartz, introducing significant dispersion. Instead, an alternative approach was adopted, in which the interference pattern was simulated in MATLAB, and a minimisation function used to determine the change in phase. An exact pulse characterisation was not necessary in order to determine the interference spectrum as the important information was encoded only in the change of the interference fringes with voltage.

The simulation proceeded as follows. A linearly sampled Gaussian pulse was constructed with a FWHM bandwidth and centre wavelength matching those of the experimental pulses. The simulated pulse was then Fourier transformed into the frequency domain where a replica was created. Next, a known phase was added to the replica, and the interference spectrum of the combined pulses was calculated. The phase added to the replica pulse consisted of two distinct components. The first corresponded to the group delay and group delay dispersion added by the quartz crystal, which produced the fringe pattern observed on the spectrometer. The second component was the voltage-dependent phase written onto the test pulse by the LC-SLM, which was modelled as a 4-point cubic spline. The simulated spectral intensity variation with wavelength and voltage formed a 2-D dataset, and the RMS error between this and the experimental data was used as the basis for a multi-dimensional minimisation. The minimisation function determined the values of group delay, group delay dispersion, bandwidth and the spline knots that produced the minimum RMS error.

The results of the phase retrieval model for pulses centred at 0.470 μm are shown in Figure 6.15. Arbitrary starting values were selected for the group delay, GDD and pulse bandwidth, resulting in a ~3π phase change as the drive voltage is increased, as shown in panel (a). The minimisation function was carried out using the experimental data as a reference, shown in panel (b). The results are shown in Figure 6.15 (c); the total

modelled phase change is approximately 5π, which matches well with the experimental data.

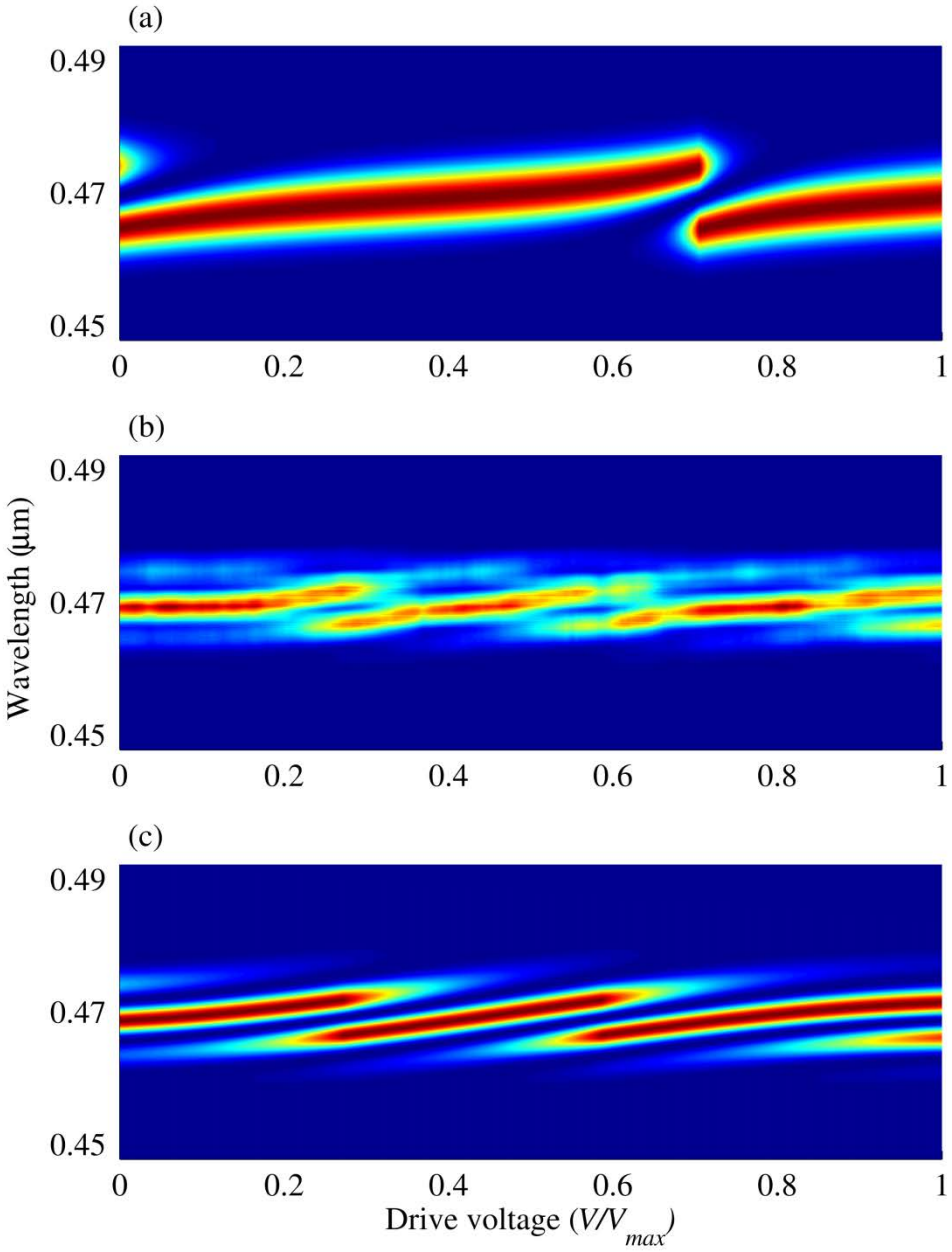

Figure 6.15. Phase change model results for pump-signal pulses centred at 0.47 µm. The minimisation function was given an arbitrary starting condition (a) and the experimental data (b), resulting in the modelled phase change shown in (c).

Comparing the spectral profiles shown in Figure 6.15 (b) and (c), it is clear that the experimental data contains large intensity variations which are not reflected in the model, potentially affecting the results of the minimisation function. By replacing the spectral intensity of the modelled pulse with that of the experimental data, the measured intensity profiles can be accommodated by the model. The results of the spectral

intensity replacement are shown in Figure 6.16. The results of the new model are shown in panel (b), and produce results that are extremely close to that of the original model. The extracted curve of voltage-to-phase is almost identical for both models, and so for simplicity the original model was used for all phase retrieval calculations.

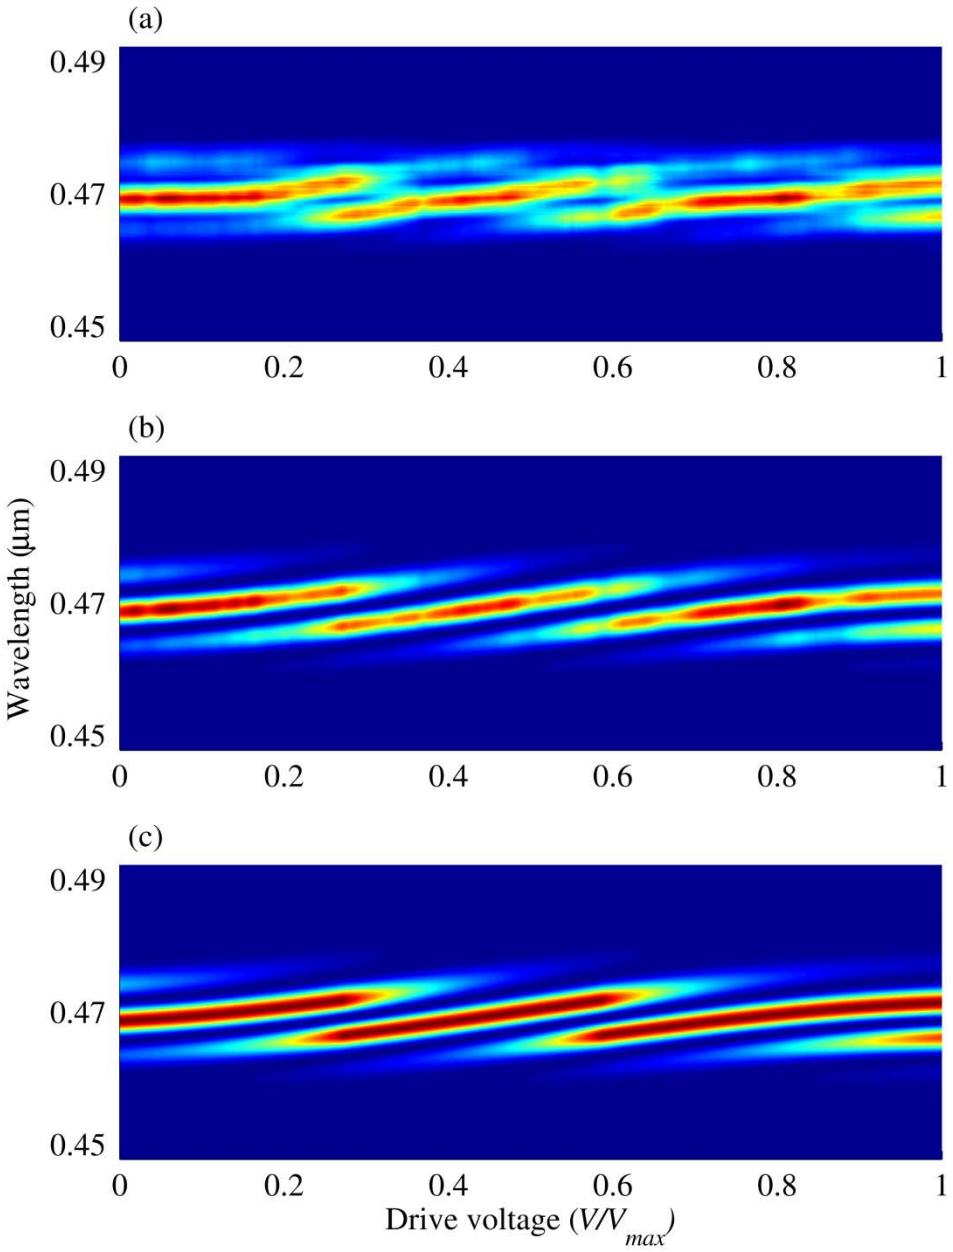

Figure 6.16. Result of the phase change model for pump-signal pulses centred at 0.47 µm. (a) Experimental spectrum; (b) model result with intensity replacement; (c) model result without intensity replacement.

Six wavelength regions were selected for voltage-to-phase calibration of the LC-SLM, spanning the doubled-pump at 0.400 µm to the lower wavelength edge of the pump at 0.770 µm. The HR coating applied to the backplane of the LC-SLM was not effective

at 0.800 µm, an unfortunate consequence of being highly reflective over such a wide frequency range. To protect the LC-SLM from damage it was necessary to block wavelengths higher than 0.780 µm with a knife edge placed after the achromatic doublet.

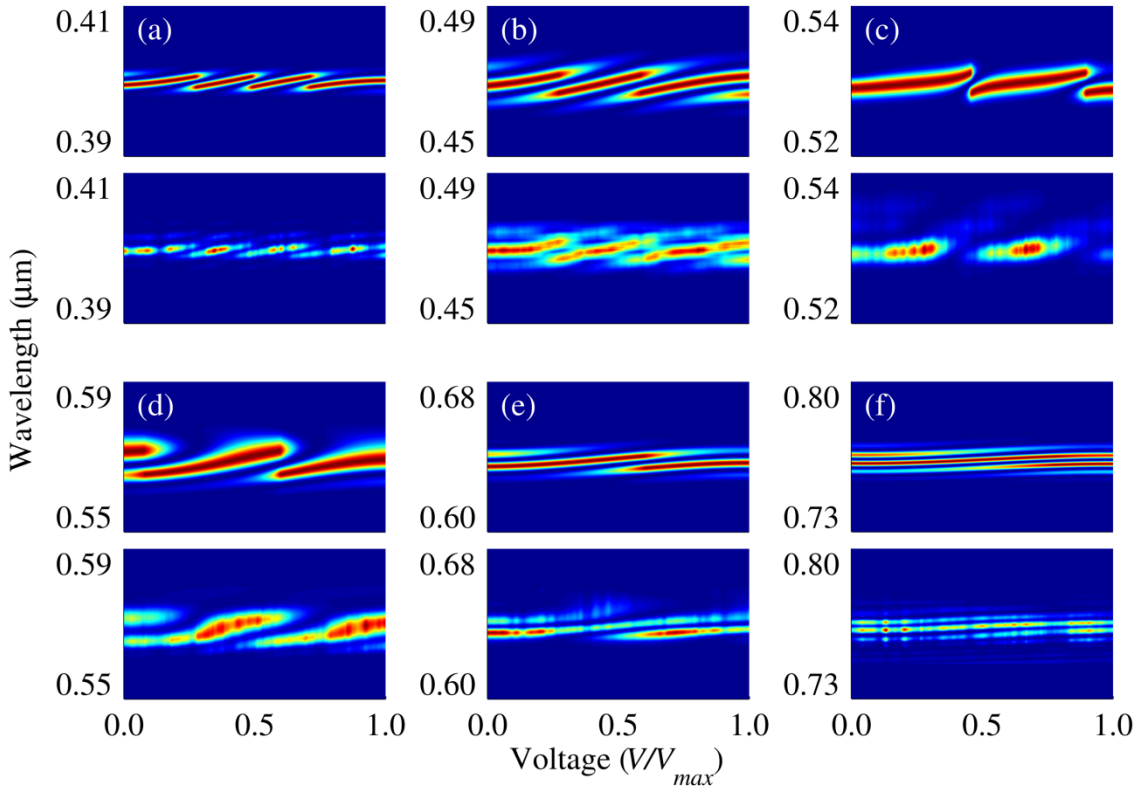

Figure 6.17. Phase change model results for six pulses spanning the 0.400 – 0.770 µm range. (a) Doubled-pump; (b) pump-signal; (c) doubled-signal; (d) doubled-signal after OPO cavity length change; (e) pump-idler; (f) lower wavelength range of pump.

The results of the phase change model for each of the test wavelengths is shown in Figure 6.17. For each wavelength the simulated interference spectrum after the minimisation function had converged (upper panels) was compared with the experimental data (lower panels). In each case the model produced accurate spectral fringe patterns which matched those recorded by the spectrometer. A curve of phase change as a function of drive voltage was extracted from each trace, with the results shown in red in Figure 6.18. The maximum possible phase change for each wavelength correctly fits the design specifications provided by Boulder Nonlinear Systems.

The phase change $\Delta\varphi$ observed experimentally arises from the difference in the optical paths between the test and reference pulses as a voltage is applied to the liquid crystal array. This path difference is a result of the birefringence of the liquid-crystal cells,

which is reduced as the drive voltage is increased. The individual liquid-crystals are positively uniaxial, and at $V = 0$ the birefringence is at a maximum $(n_e > n_o)$, therefore the horizontally polarised test pulses travel a longer optical path than those in the reference arm. The drive voltage is increased until $V = V_{max}$, at which point $n_e \approx n_o$. This change in refractive index can be approximated by

$$\Delta\varphi = \frac{2\pi L n_e(\lambda)}{\lambda} - \frac{2\pi L}{\lambda}\left[\frac{\cos^2\left(\frac{\pi}{2}\frac{V}{V_{max}}\right)}{n_e^2(\lambda)} + \frac{\sin^2\left(\frac{\pi}{2}\frac{V}{V_{max}}\right)}{n_o^2(\lambda)}\right]^{-\frac{1}{2}}, \tag{6.1}$$

where L is twice the thickness of the liquid crystal cell. At lower voltages the $\cos^2$ term dominates and the phase change is small. As the voltage is increased the equilibrium of the equation shifts until the $\sin^2$ term dominates. Equation (6.1) was used to determine the refractive index and liquid-crystal cell length for each individual wavelength. Each fit provided a liquid-crystal cell length of 5.1±0.05 µm. Each phase curve was then fitted simultaneously in order to determine a dispersion equation [16] for the two liquid-crystal optical axes of the form

$$n_{e,o}(\lambda) = A_{e,o} + B_{e,o}\lambda^{-2} + C_{e,o}\lambda^{-4}. \tag{6.2}$$

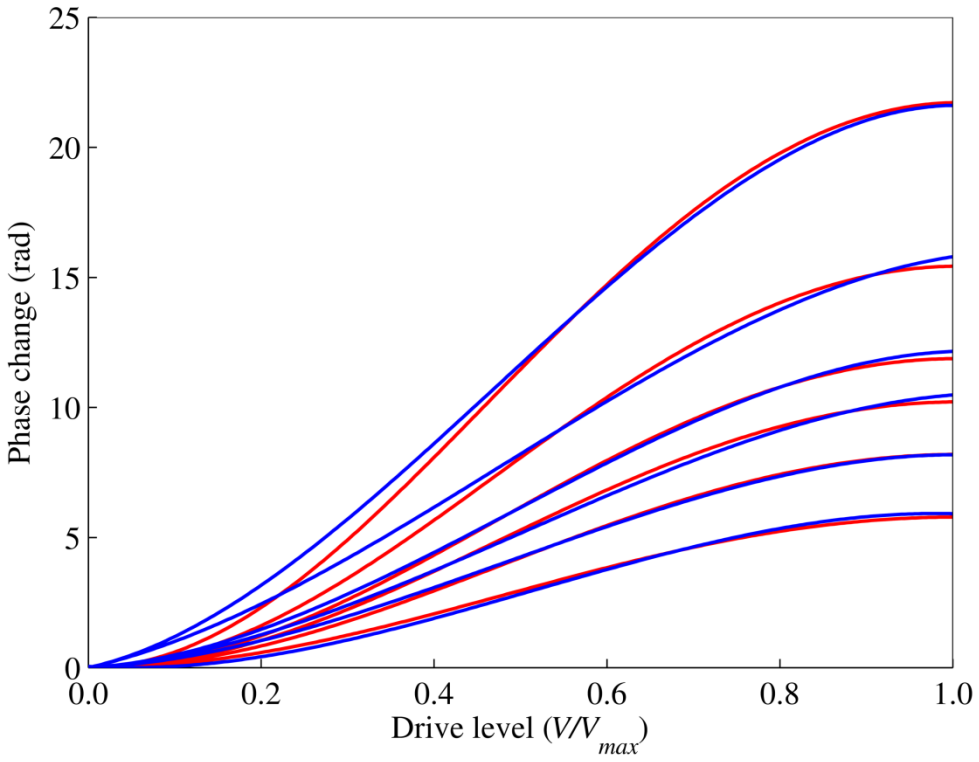

Figure 6.18. Extracted phase changes as a function of voltage applied to the LC-SLM (red) and the phase changes calculated from the fitted dispersion equation (blue).

6.3.2.3 Discussion

The phase change modelled by the best-fit dispersion equation is shown in Figure 6.18 (blue lines) and was compared against the phase response obtained experimentally (red lines). Good agreement was obtained at higher voltages, however at lower voltages the error was larger as Equation (6.1) does not account for forces not associated with the applied field (e.g. Van der Waals) which also act on the liquid-crystals. Experimentally the decrease in birefringence observed is slightly less sinusoidal than assumed, and at maximum drive voltage $n_e \neq n_o$. A plot of the dispersion equations calculated from Equation (6.2) is shown in Figure 6.19. Although details of the exact liquid-crystal used in the LC-SLM were not available from the manufacturer, the birefringence and dispersion of the refractive indices are consistent with documented liquid-crystal data [16].

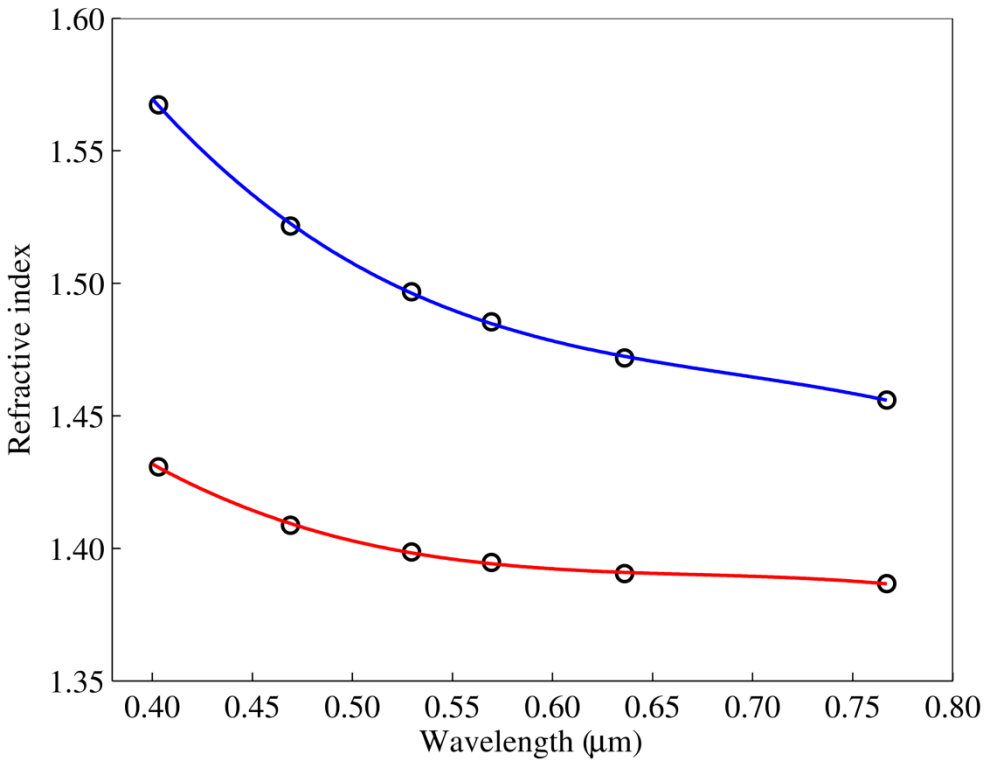

Figure 6.19. Refractive indices of the unknown liquid-crystal material from the modelled phase changes (circles) and best fit to Equation (6.2) for n_o (red) and n_e (blue).

This approach for characterising an LC-SLM over a broad bandwidth has produced a deterministic map of the LC-SLM phase response as a function of wavelength and voltage, providing the necessary information to apply the device to the compression and shaping of femtosecond pulses in the visible region. Unfortunately at this point a

software issue rendered the LC-SLM unusable for several months, and so compression and delay adjustment of the pump and OPO pulses could not be attempted.

6.4 Conclusions

Two methods for compressing the parent pulse for coherent synthesis were introduced and considered. The prism-based approach showed that individual control of the GDD and group delay of the parent pulses is possible but without fine control. The LC-SLM approach is the most promising candidate for a broadband common-path pulse compressor for coherent synthesis. The optimal result for synthesis between the three parent pulses was modelled, and showed that the synthesised pulses would remain usable over a 1 second observation window.

6.5 References

1. J. P. Heritage, A. M. Weiner, and R. N. Thurston, "Picosecond pulse shaping by spectral phase and amplitude manipulation," Opt. Lett. **10**, 609–11 (1985).

2. A. M. Weiner, D. E. Leaird, J. S. Patel, and J. R. Wullert, "Programmable femtosecond pulse shaping by use of a multielement liquid-crystal phase modulator," Opt. Lett. **15**, 326–28 (1990).

3. A. M. Weiner, D. E. Leaird, J. S. Patel, and J. R. Wullert, "Programmable shaping of femtosecond optical pulses by use of 128-element liquid crystal phase modulator," IEEE J. Quantum Electron. **28**, 908–920 (1992).

4. F. Verluise, V. Laude, Z. Cheng, C. Spielmann, and P. Tournois, "Amplitude and phase control of ultrashort pulses by use of an acousto-optic programmable dispersive filter: pulse compression and shaping," Opt. Lett. **25**, 575–7 (2000).

5. J. Garduño-Mejía, A. Greenaway, and D. Reid, "Designer femtosecond pulses using adaptive optics," Opt. Express **11**, 2030–40 (2003).

6. A. Assion, T. Baumert, M. Bergt, T. Brixner, B. Kiefer, V. Seyfried, M. Strehle, and G. Gerber, "Control of chemical reactions by feedback-optimized phase-shaped femtosecond laser pulses," Science (80-.). **282**, 919–922 (1998).

7. B. J. Pearson and T. C. Weinacht, "Shaped ultrafast laser pulses in the deep ultraviolet," Opt. Express **15**, 4385–8 (2007).

8. S.-H. Shim, D. B. Strasfeld, E. C. Fulmer, and M. T. Zanni, "Femtosecond pulse shaping directly in the mid-IR using acousto-optic modulation," Opt. Lett. **31**, 838–40 (2006).

9. M. Shverdin, D. Walker, D. Yavuz, G. Yin, and S. Harris, "Generation of a single-cycle optical pulse," Phys. Rev. Lett. **94**, 033904 (2005).

10. A. M. Weiner, "Ultrafast optical pulse shaping: A tutorial review," Opt. Commun. **284**, 3669–3692 (2011).

11. T. Binhammer, E. Rittweger, R. Ell, F. X. Kärtner, S. Member, and U. Morgner, "Prism-based pulse shaper for octave spanning spectra," IEEE J. Quantum Electron. **41**, 1552–57 (2005).

12. I. H. Malitson, "Interspecimen comparison of the refractive index of fused silica," J. Opt. Soc. Am. **55**, 1205–09 (1965).

13. P. Schlup, J. Wilson, K. Hartinger, and R. A. Bartels, "Dispersion balancing of variable-delay monolithic pulse splitters," Appl. Opt. **46**, 5967–73 (2007).

14. J. W. Wilson, P. Schlup, and R. A. Bartels, "Ultrafast phase and amplitude pulse shaping with a single, one-dimensional, high-resolution phase mask," Opt. Express **15**, 8979–87 (2007).

15. M. Takeda, H. Ina, and S. Kobayashi, "Fourier-transform method of fringe-pattern analysis for computer-based topography and interferometry," J. Opt. Soc. Am. **72**, 156–160 (1982).

16. J. Li, C.-H. Wen, S. Gauza, R. Lu, and S.-T. Wu, "Refractive indices of liquid crystals for display applications," J. Disp. Technol. **1**, 51–61 (2005).

Chapter 7 - Conclusions and future developments

This chapter summarises the experimental work carried out in this thesis and the conclusions that can be drawn from these results. Discussion is then given to the viability of broadband coherent pulse synthesis from an OPO, confirming that pulse synthesis has been successfully achieved, and future opportunities for improving the results that have been presented in earlier chapters.

7.1 Technical summary and conclusions-

In Chapter 3 the design and construction of a synchronously-pumped femtosecond optical parametric oscillator was presented. A mode-locked Ti:sapphire pump laser producing sub-20-fs pulses at an average power of 1.4 W, repetition frequency of 100 MHz and a centre wavelength of 0.800 μm was used to pump a PPKTP-based OPO with a wide range of tunable visible outputs. Both linear and ring cavity designs were built and evaluated for the OPO, with the ring OPO able to produce higher power visible sum-frequency mixing and second-harmonic generation pulses.

In Chapter 4 the carrier-envelope-offset frequencies of the pump, signal, idler and multiple visible outputs from the OPO were locked to 0 Hz. Light from the pump was launched into two photonic crystal fibers to generate independent supercontinua, which were heterodyned against separate visible wavelength components from the OPO to produce two CEO beat frequencies. These CEO frequencies were controlled through a pair of electronic feedback loops. Locking all CEO frequencies to 0 Hz ensured broadband phase coherence from 0.4 – 3.2 μm, the broadest zero-offset frequency comb produced to date. Establishing broadband phase coherence was the first requirement for coherent pulse synthesis.

Chapter 5 discussed the noise characteristics of the pump and OPO. RIN measurements were carried out, and locking the CEO frequencies of the pump and OPO was found to greatly improve the cumulative standard deviation of the visible OPO outputs due to the increased frequency stability the locking loops introduce. Phase noise PSD data were recorded for both CEO locking loops, with the largest cumulative standard deviation being 0.18 radians over a 1 second acquisition time. A balanced cross-correlator was used to measure the relative timing jitter between the pump, doubled-signal and pump-signal pulses. With both locking loops enabled the timing jitter was less than 140 as over a 1 second observation window, indicating that the pulse trains had excellent

timing stability, sufficient for coherent synthesis, and fulfilling the second requirement for coherent pulse synthesis.

In Chapter 6 two complementary methods for compressing and controlling the relative delay of the pump and OPO pulses were discussed. A quasi-common-path prism delay line was built to compress the pump, doubled-signal and pump-signal pulses, with the pulse durations calculated using XFROG. Adequate adjustment of the relative delays between the pulses was shown to be possible. An alternative approach was introduced in which a liquid-crystal spatial light modulator formed the basis of a broadband pulse shaper. The LC-SLM was characterised in order to provide a functional map of frequency-to-pixel number and voltage-to-phase, allowing the device to be used to both compress the pulses and alter their relative delays in a truly common-path approach. Unfortunately, firmware issues prevented further experiments with the device. The successful demonstration of pulse compression and delay compensation with the prism delay line fulfilled the final requirement for coherent pulse synthesis.

7.2 Future developments for broadband coherent synthesis from an OPO

7.2.1 The viability of multi-colour pulse synthesis from an OPO

An important question arises from the summaries given above – *what is the viability of multi-colour coherent pulse synthesis from a femtosecond OPO?* Chapter 1 laid out the experimental milestones that needed to be achieved before synthesis could be attempted. The requirements were

- Matching repetition frequencies for all the parent pulses. **This was achieved by constructing using a synchronously-pumped femtosecond OPO as the source of the parent pulses, as detailed in Chapter 3.**
- Broadband phase coherence across the entire bandwidth of pulses generated by the pump and OPO. **This was achieved through zero-offset CEO locking and is discussed in Chapter 4.**
- Control of the timing jitter between the pulses to attosecond precision. **This was achieved and is discussed in Chapter 5.**
- To correct for group delay dispersion for each pulse and between different pulses. **This was achieved using a prism-delay line and the technique can be refined using a spatial light modulator, as detailed in Chapter 6.**

Using these criteria and the results presented in previous chapters it is clear that multi-colour pulse synthesis from an OPO is not only viable but can be extended beyond the three-parent-pulse scenario presented in this thesis to include all pulses from the OPO.

A question of more specific interest to this thesis also arises – *were daughter pulses synthesised from the considered parent pulses?* This question is harder to answer than the first. Clearly the criteria for pulse synthesis have all been met, but knowing the exact synthesis outcome is challenging. In Chapter 6 the parent pulses were compressed and their relative delays adjusted so that the pulses arrived at the measurement plane (the XFROG crystal) at the same time. The retrieved XFROG trace indicated that the pulses were overlapped in time, while a measurement interferometer confirmed that the pulses were phase coherent, as discussed in Chapter 4. An ambiguity exists in the XFROG results however; the XFROG algorithm is unable to recover the relative phases between the pulses over spectral gaps. Confirming phase coherence at the measurement plane would definitively establish whether pulse synthesis had taken place. This measurement must take place at a chosen measurement plane where the pulses have no delay offset; downstream of the measurement plane the different spectral components of the synthetic pulse will propagate at different group velocities, separating into the original parent pulses.

The concluding analysis in Chapter 6 addresses the synthesis outcomes which could be achieved in the case of optimally chosen relative phases between the parent pulses. Confirming phase coherence between the pulses at the measurement plane is not possible using established pulse retrieval techniques and the current experimental apparatus. A novel approach has been suggested by our collaborators from Oxford, Ian Walmsley, Adam Wyatt and Ilaria Gianani. An ancilla pulse is required to span the spectral gap between the parent pulses. If this broadband ancilla is used as the gating pulse in an XFROG measurement then there will be sufficient information in the resulting trace to establish phase coherence. The simplest approach would be use a small piece of PCF to generate self-phase modulation in the pump beam such that its bandwidth spanned the largest frequency gap between the parent pulses. This spectrally chirped pulse would be temporally compressed and used in an XFROG measurement, as has already been demonstrated in Chapter 6. The results of such an XFROG measurement are simulated in Figure 7.1. The phase retrieval algorithm would have

sufficient information to span the spectral gaps between the test pulses, providing the phase relationship between them and confirming phase coherence.

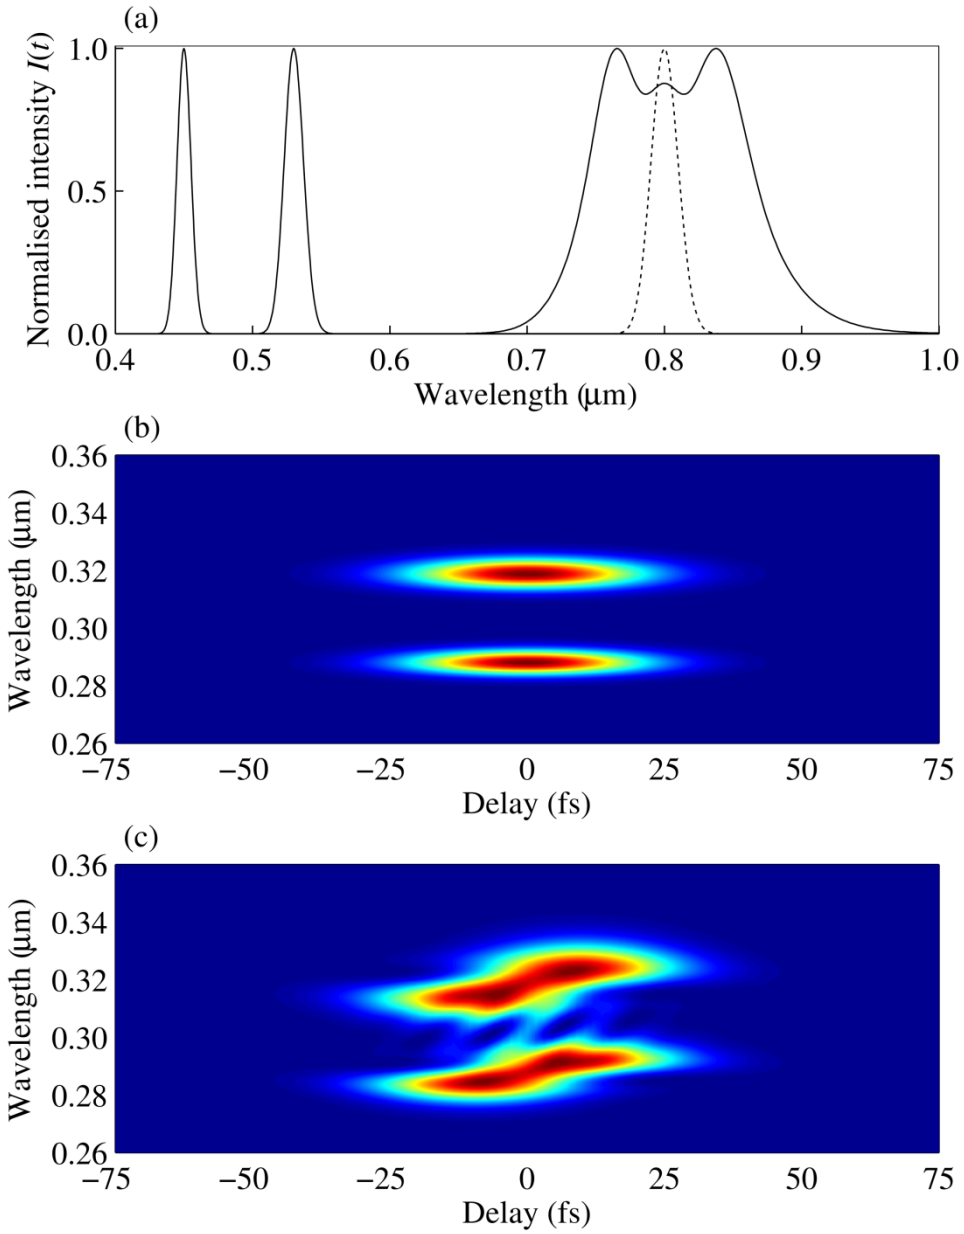

Figure 7.1. Simulation of a suggested method for confirming phase coherence between the pulses. The addition of a spectrally-broadened pump beam would bridge the spectral gap observed in the FROG trace, allowing retrieval of the relative phase information between the pulses.

The result of the spectral confirmation of phase coherence described above would fully determine whether the requirements for pulse synthesis had been met, providing information on phase coherence, pulse compression and temporal overlap at the measurement plane. Several attempts to generate a correct amount of SPM in the pump

pulses were made, however none were wholly successful. The high average power and high peak powers of the compressed pump pulses led to supercontinuum generation when launched into even millimetre lengths of the PCF available; launching chirped pump pulses into the PCF was also considered, however it is much more challenging to remove this additional chirp from the pulses after the PCF. Sourcing PCF of the correct length and dispersion characteristics would make this measurement much more viable.

7.2.2 Technical improvements and future work proposals

The project goal of coherently combining multiple visible outputs from an OPO has been achieved; however there are a number of alterations that can be made to the system that would greatly improve the synthesis result.

7.2.2.1 Producing sub-30-fs visible pulses from the OPO

The shortest pulse durations that can be coherently synthesised are produced by combining pulses that are themselves short in duration. Combining transform-limited pulses centred at 0.800, 0.530 and 0.456 μm with durations of 100 fs will produce an ideal synthesised pulse of 86 fs duration; if the parent pulses are 15 fs in duration then the synthesised pulse will be 2.2 fs duration.

The visible pulse durations available from the OPOs described in this thesis were limited by the phasematching conditions of the PPKTP crystal. SFM and SHG processes were not optimally phasematched and so the generated pulses were typically narrowband with low average powers; the higher powers available from the doubled-signal and pump-signal beams was due to the high intracavity signal power.

A novel approach to improving SFM and SHG efficiency would be to use a grating-engineered quasi-phasematched nonlinear crystal in the OPO. In previous work demonstrated by the Reid group [1–7] the poling period of the nonlinear crystal was chirped to compensate for the temporal walk-off between the pump and signal pulses in the crystal. This technique can be further extended to produce a cascaded crystal (Figure 7.2) with different grating lengths that will produce well-phasematched regions of overlap for the pump and signal, allowing the generation of broadband visible pulses. A crystal of this design was grown for testing in the OPO; however the crystal was cut with a large wedge, making alignment difficult and as a result oscillation was not achieved.

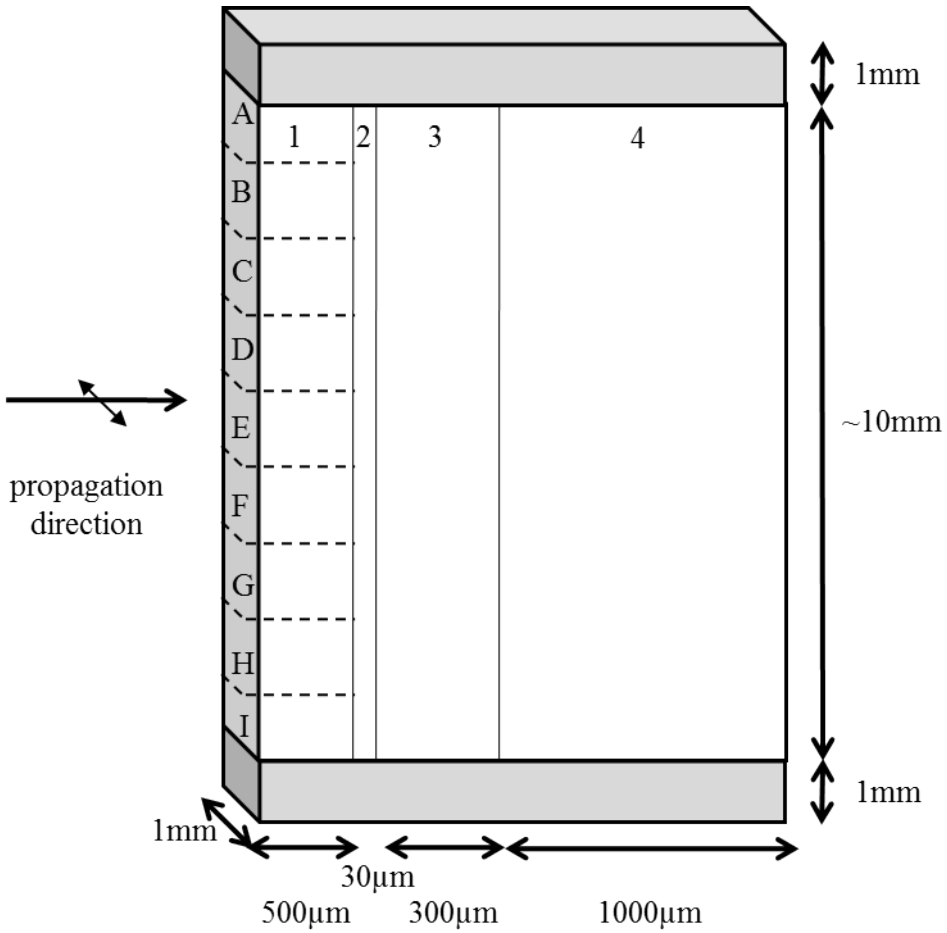

Figure 7.2. Schematic of the cascaded crystal design for efficient SFM and SHG phasematching. Section 1, parametric generation with grating periods increasing from A to I; section 2, signal SHG; section 3, pump-signal SFM; section 4, doubled-pump SHG. The top and bottom 1 mm of the crystal are unpoled.

7.2.2.2 Wavelength-independent CEO frequency locking of an OPO

The CEO frequency locking approach described in Chapter 4 was successful in generating two CEO frequencies without the need for an *f*-2*f* interferometer. A drawback of the chosen design is that CEO frequencies can only be detected at selected wavelengths determined by both the visible wavelengths from the OPO and by the narrowband interference filters placed before the APDs. It would be beneficial to enable the OPO to be tuned to produce different wavelengths and still have the capability to lock the pump and OPO CEO frequencies.

The first possible solution would be to use a diffraction grating in place of the interference filter to spatially disperse the overlapped beams from a pump supercontinuum and the OPO. The position of the overlapped wavelengths of interest are determined by the angle of the diffraction grating; to acquire a CEO frequency with a different OPO wavelength would simply require the repositioning of the APD or

rotation of the diffraction grating. These processes could in theory be automated, so that a CEO frequency could constantly be detected as the OPO wavelength is tuned.

A second solution would be to use an external cavity diode laser (ECDL) referenced to an atomic transition in a gas cell. A well-referenced and stabilised ECDL will produce an extremely stable and well-defined frequency output; this frequency can be heterodyned against a frequency comb to produce a CEO beat frequency. An ECDL referenced to a hyperfine transition in a Rubidium gas cell at ~0.78024 μm would allow such a CEO frequency to be detected directly from the broadband Ti:sapphire laser. The signal pulses in the OPO can be externally broadened in a length of PCF to produce light at 1.56 μm, which can then be frequency doubled and heterodyned against the ECDL beam to produce a second CEO frequency. With sufficient broadening in the PCF the OPO can be tuned with the CEO locking frequency remaining detected.

7.2.2.3 Simultaneous pulse compression and synthesis

The pulse compression methods discussed in Chapter 6 utilised the XFROG apparatus to determine the pulse duration at the measurement plane. Compressing each pulse requires an iterative process that is time-dependent on the acquisition and retrieval of the XFROG trace. A more efficient approach, and one that is suited to the use of a well-characterised LC-SLM, is to use MIIPS to both compress and characterise the pulse in the reference plane. With a sufficiently broadband SHG crystal (or several thin crystal stacked together) it would be possible to compress each visible pulse simultaneously using a modified MIIPS algorithm, providing transform-limited and delay-compensated pulses in a single plane for pulse synthesis. More details on MIIPS can be found in Chapter 2.

7.2.2.4 Future proposals for multi-colour pulse synthesis from an OPO

Building on the successful results presented in this thesis, it is clear that multi-colour pulses synthesis from an OPO is both achievable and can be scaled up to include more than three parent pulses. With an appropriate choice of LC-SLM it is feasible to both compress and temporally combine the visible outputs listed in Table 7.1, a bandwidth capable of supporting sech^2 pulses with duration of 1.12 fs. Coherent synthesis between pulses of these wavelengths with technically-challenging durations of 20 fs would produce a 2.4 fs pulse.

Table 7.1. Mixing wavelengths generated in the OPO crystal for a given pump and signal wavelength.

Primary wavelengths (μm)			Mixing wavelengths (μm)				Visible frequency bandwidth (PHz)
p	**s**	**i**	**2*p***	***p* + *s***	**2*s***	***p* + *i***	**Δ*f***
0.80	1.06	3.26	0.40	0.45	0.53	0.64	0.27

Introduction of the fourth-harmonic of the pump at 0.2 μm would increase the frequency bandwidth to over 1 PHz, capable of supporting pulse durations of ~300 as. Generating 0.2 μm is in itself a challenge, as is being able to compress and control a pulse at that wavelength, and so achieving visible attosecond pulse synthesis will require a great deal of further investigation.

The techniques developed in previous chapters can be utilised in a range of scientific fields, but there are many immediate opportunities in the field of spectroscopy. The wavelength stability demonstrated from the CEO-locked OPO can be used for measurements that require a long exposure time with high pump wavelength stability, while the broadband phase coherence of the frequency comb can be exploited for phase coherent pump-probe spectroscopy.

7.3 References

1. K. A. Tillman, D. T. Reid, D. Artigas, J. Hellström, V. Pasiskevicius, and F. Laurell, "Low-threshold femtosecond optical parametric oscillator based on chirped-pulse frequency conversion," Opt. Lett. **28**, 543–5 (2003).

2. K. A. Tillman, D. T. Reid, D. Artigas, J. Hellström, V. Pasiskevicius, and F. Laurell, "Low-threshold femtosecond optical parametric oscillator based on chirped-pulse frequency conversion," J. Opt. Soc. Am. B **20**, 1309–1316 (2003).

3. D. Artigas and D. T. Reid, "Efficient femtosecond optical parametric oscillators based on aperiodically poled nonlinear crystals," Opt. Lett. **27**, 851–3 (2002).

4. K. A. Tillman and D. T. Reid, "Monolithic optical parametric oscillator using chirped quasi-phase matching," Opt. Lett. **32**, 1548–50 (2007).

5. D. T. Reid, "Engineered quasi-phase-matching for second-harmonic generation," J. Opt. A Pure Appl. Opt. **5**, S97–S102 (2003).

6. U. K. Sapaev and D. T. Reid, "General second-harmonic pulse shaping in nonlinear crystals," Opt. Express **13**, 3264–3276 (2005).

7. Ł. Kornaszewski, M. Kohler, U. K. Sapaev, and D. T. Reid, "Designer femtosecond pulse shaping using grating-engineered quasi-phase-matching in lithium niobate," Opt. Lett. **33**, 378–380 (2008).