Versatile High Resolution Dispersion Measurements in Semiconductor Photonic Nanostructures using Ultrashort Pulses

Thesis submitted to the University of St Andrews
in application for the degree of Doctor of Philosophy

Matthew Richard Bell

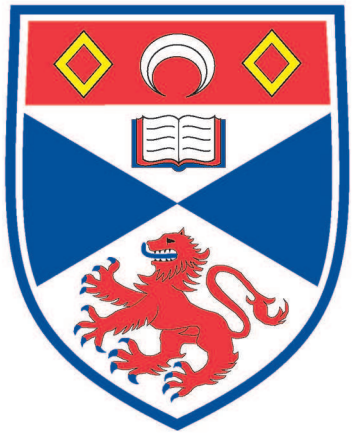

University
of
St Andrews

School of Physics and Astronomy
University of St Andrews
Fife, Scotland

December 2006

Declaration

I, Matthew Richard Bell, hereby certify that this thesis, which is approximately 30000 words in length, has been written by me, that it is the record of work carried out by me and that it has not been submitted in any previous application for a higher degree.

Date Signature of candidate

I was admitted as a research student in October, 2001 and as a candidate for the degree of Doctor of Philosophy in 2002; the higher study for which this is a record was carried out in the University of St. Andrews between 2001 and 2006.

Date Signature of candidate

I hereby certify that the candidate has fulfilled the conditions of the Resolution and Regulations appropriate for the degree of Doctor of Philosophy in the University of St. Andrews and that the candidate is qualified to submit this thesis in application for that degree.

Date Signature of supervisor

Unrestricted

In submitting this thesis to the University of St. Andrews I understand that I am giving permission for it to be made available for use in accordance with the regulations of the University Library for the time being in force, subject to any copyright vested in the work not being affected thereby. I also understand that the title and abstract will be published, and that a copy of the work may be made and supplied to any bona fide library or research worker.

Date Signature of candidate

CONTENTS

1. INTRODUCTION

1.1 Abstract

This thesis describes the process of developing a robust phase measurement technique with which to analyse semiconductor based devices intended for use in optoelectronic/all optical networks. The devices measured are prospective dispersion compensators, based either on planar photonic crystal waveguides or coupled microcavities connected by ridge waveguide. The technique was validated by measuring the phase transfer function of a Fabry Perot etalon. This demonstrated that even when detecting low optical powers (sub μW), accurate measurement of phase could quickly be carried out over a significant spectral range ($\sim$10nm). Comparison of experimental data taken from the prospective dispersion compensators with theory showed excellent agreement, which provided qualitative (cavity spacing and reflectivity) and quantitative (loss) measures of device performance.

The phase measurement technique has been designed to be capable of measuring other classes of device also, including active devices such as semiconductor optical amplifiers. This suggests the phase measurement technique may be valuable in analysing the variation of dispersion as a function of applied bias, peak power or temperature for a variety of devices.

1.2 Introduction

In developed countries it is increasingly rare for any member of the population to not use the miles of optical fibres that provide the backbone of long haul data transmission. While telephony was the initial driving factor in creating a global telecommunication network, increasing demands are being placed upon

existing infrastructure to deliver more data, whether audio, video or computer data. Copper cables became ubiquitous because of their low cost and ease of production. At low transmission rates the disadvantages of high signal attenuation, interference sensitivity and susceptibility to ground potentials were not prohibitive. Attenuation losses can be offset by increasing cable diameter, but this quickly becomes costly and impractical. More significantly, loss increases with modulation frequency, which limits maximum data transmission rates. The electrical field from copper cables (even with shielding) is susceptible to interference with other electromagnetic fields, increasing error rate. The difference in electrical potential between each end of a copper cable can give rise to DC currents, again increasing error rate.

Although optical fibres were uncompetitive originally due to undeveloped methods of production and relative cost, they have the potential to provide orders of magnitude improvement over current data transmission rates [1]. The commercial pressure to improve data transmission rates is unarguable; the U.S. telecommunications industry is worth ∼$1 trillion alone [2], and worldwide it accounts for over 2% of the global GDP. Previously dominant technologies have been unable to match the bandwidth offered by optical fibres and as such have lost their place as the technology of choice for long haul transmission. For example, in 1988 only 2% of the world transoceanic flow of messages and data was carried out by undersea cables (satellites were then the dominant carrier). By 2000 this had increased to 80%, due to the introduction of fibre optic cables [3]. Accepting that the current network of subterranean fibres (which were designed and manufactured for lower transmission rates than currently available) would be overly expensive and time consuming to be replace, further increases in data transmission rates will result from superior signal regeneration. Signal regeneration is necessary to compensate for jitter, loss and dispersion of pulses.

Long-haul communications currently avoid the problem of dispersion by converting optical pulses to electrical signals before significant pulse degradation occurs. The electrical signals are then used to re-transmit the data with newly generated pulses. This is robust and reliable, but limits data transmission rates and offers little room for improvement.

Keeping the signal entirely optical requires retiming, amplification and dis-

persion compensation. Research in to the use of modelocked lasers[4] and semiconductor optical amplifiers or fibre amplifiers (most significantly erbium doped fibre amplifiers) to perform retiming and amplification is ongoing.

All-optical networks may be envisaged as the next step after all optical long haul data transmission. To realise these kinds of networks, additional functions must be performed optically: muliplexing/demultiplexing, switching/routing and wavelength conversion. All optical components, such as fibres and amplifiers will exhibit dispersion. To achieve all optical long haul transmission or all optical networks, dispersion compensators must be employed.

A dispersion compensator would ideally produce exactly the opposite dispersion to that induced by the transmission medium (i.e optical fibre) or system (such as a combination of amplifiers, fibres and retiming elements). The pulse would then be identical (in frequency dependent phase) before and after transmission through the system of dispersion compensator and other optical component(s). Ideally a dispersion compensator would be without associated losses and in a compact, inexpensive and reliable package. One proven method of dispersion compensation that has shown the capability of far higher transmission rates than currently available is the use of fibres with engineered dispersion profiles [5]. Lengths of fibre of approximately half the total transmission distance are needed for this, negating their use with existing subterranean optical fibres. Other dispersion compensators exist, but none are yet suitable for installation in to existing optical fibre networks, because of cost, reliability, size or performance limitations.

Semiconductor nanostructures, including photonic crystals, which were previously hampered by fabrication limitations or extremely time consuming modelling are now beginning to demonstrate many of the qualities sought after in dispersion compensators. Compact and inexpensive, the dispersion of semiconductor nanostructures can be engineered to fully compensate for dispersion imparted to pulses by many optical components or many kilometers of fibre. As theory, modelling and fabrication technology improves, experimental dispersion measurements of these structures are needed to complete the iterative cycle of improvement that should eventually yield commercially useful devices and a greater understanding of fundamental physics.

REFERENCE LIST

[1] L E Nelson Y Qian L Cowsar S Stulz C Doerr L Stulz S Chandrasekhar S Radic D Vengsarkar Z Chen J Park K S Feder H Thiele J Bromage L Gruner-Nielsen B Zhu, L Leng and S Knudsen, "3.08 Tbit/s (77 42.7 Gbit/s) WDM transmission over 1200 km fibre with 100 km repeater spacing using dual C- and L-band hybrid Raman/erbium-doped inline amplifiers", *Electronics Letters*, vol. 37, no. 13, pp. 844–845, 2001.

[2] Patrick Barnard, "Telecom Revenue to Grow an Average 9 Percent Annually Until 2009", *TMCnet, http://www.tmcnet.com/news/2006/02/17/1384506.htm*, vol. February 17, 2006.

[3] M Mandell, "120000 leagues under the sea", *IEEE Spectrum*, vol. 37, no. 4, pp. 50–54, 2000.

[4] P S Cho N Moulton E S Awad, C J K Richardson and J Goldhar, "Bi-directional coupling in nonlinear waveguides for absolute timing determination", *Lasers and Electro-Optics, CLEO Technical Digest.*, pp. 472–473, 2001.

[5] S G Evangelides Jr J P Gordon J R Simpson L F Mollenauer, M J Neubelt and L G Cohen, "Experimental study of soliton transmission over more than 10,000 km in dispersion-shifted fiber", *Optics Letters*, vol. 15, pp. 1203–1205, 1990.

2. DISPERSION, NONLINEARITIES AND PHOTONIC CRYSTALS

This chapter will introduce two concepts essential to understanding the fundamental behaviour of ultrashort optical pulses: dispersion and optical nonlinearities. The fundamental physics and origins of each will be discussed with particular emphasis paid to their effect on pulse integrity. The dispersion due to a dynamic free carrier population is also discussed, illustrating the need for experimental dispersion measurement techniques to complement theoretical modelling in the complicated regime of pulse propagation. Lastly a brief overview of photonic crystals will be presented. This class of devices is a leading contender for compact, tunable dispersion compensators, capable of being integrated in to semiconductor waveguide based devices.

2.1 Dispersion

The simplest definition for dispersion could be 'a frequency dependent speed of propagation', or in the case of propagation through an isotropic medium, 'a frequency dependent refractive index'.

Dispersion could then be defined by the following relation,

$$v(\omega) = \frac{\omega}{k(\omega)} = \frac{c}{n(\omega)} \tag{2.1}$$

where $v(\omega)$ is speed of light in the medium, ω is angular frequency of light, $k(\omega)$ is wavevector in the medium, c is the speed of light in vacuum and $n(\omega)$ is refractive index. To understand the origins of these observations it is necessary to start by considering the fundamental interactions between electric fields and dielectric media. Note, all media can be regarded as dielectric, even

if weakly so. While light consists of a coupled propagating electric and magnetic field, this chapter will treat all light-matter interactions as functions of the electric field only. Conceptually simpler, this is also a convenient and justifiable perspective. The strength of the electric and magnetic fields of a plane wave are proportional, with the proportionality constant a function of the medium through which the wave propagates. Consequently light matter interactions can be defined through either field. While it may be more convenient to choose to analyse a particular situation through one field (conventionally analysis of photonic crystal structures is carried out by considering the magnetic field for example), doing so will not effect any derived conclusions. The classical theory of light-matter interactions is founded on Maxwell's equations. With prior knowledge that the media of interest to optoelectronics research are non-magnetic semiconductors, these seminal equations can be simplified and written as below:

$$\nabla \times \mathbf{E} = -\mu_0 \frac{\partial \mathbf{H}}{\partial t} \tag{2.2}$$

$$\nabla \times \mathbf{H} = \epsilon_0 (\frac{\partial \mathbf{E}}{\partial t}) + \frac{\partial \mathbf{P}}{\partial t} \tag{2.3}$$

$$\nabla \cdot \mathbf{D} = \rho \tag{2.4}$$

$$\nabla \cdot \mathbf{H} = 0 \tag{2.5}$$

Here $\mathbf{E}$ and $\mathbf{H}$ represent electric and magnetic fields respectively, and $\mathbf{P}$ denotes polarisation. The fundamental constants μ_0 and ϵ_0 are magnetic permeability of, and electric permittivity in, free space respectively and ρ is conductivity. While $\mathbf{E}$ and $\mathbf{H}$ are properties of light, $\mathbf{P}$ quantifies the response of a medium to an electric field. In a classical sense $\mathbf{P}$ can be regarded as the induced displacement of loosely bound electrons from their default centre of gravity (which, in the absence of applied fields, will be the nucleus to which they are bound).

When light interacts with a material in the linear regime (i.e. for low or moderate light intensity), the induced polarisation is assumed to be directly proportional to the incident electric field. This assumption is justified within

experimental accuracy.

$$\mathbf{P}(\omega) = \epsilon_0 \chi(\omega) \mathbf{E}(\omega) \tag{2.6}$$

The constant of proportionality, χ is the electric susceptibility. It is a material property and a function of frequency. χ is complex and quantifies the response of a material to an applied electric field. Absorption, refraction and dispersion are all fully described by χ.

The electric susceptibility (which is sometimes defined through relative permittivity, $\epsilon_r = 1 + \chi$) is related to more physically intuitive quantities by:

$$\begin{aligned} \left(\frac{k(\omega)c}{\omega}\right)^2 &= 1 + \chi(\omega) = 1 + \chi^{'}(\omega) + \imath\chi^{''}(\omega) \\ \left(\frac{k(\omega)c}{\omega}\right) &= \eta(\omega) + \imath\kappa(\omega) = n(\omega) \end{aligned} \tag{2.7}$$

where $\eta(\omega) + \imath\kappa(\omega)$ is complex refractive index. The real component of refractive index, $\eta(\omega)$, determines the speed of light in the medium, and hence describes dispersion. The refractive index term in equation 2.1 is the real component of refractive index only and should be written explicitly as $\eta(\omega)$.

The imaginary component of refractive index, κ determines the magnitude of absorption (or gain in a biased medium).

$$\begin{aligned} \alpha &= \frac{2\omega\kappa(\omega)}{c} \\ I(z) &= I_0 \exp^{-\alpha z} \end{aligned}$$

where α is the absorption coefficient of the Beer-Lambert law as shown above, with I_0 and I initial intensity and intensity after propagation through distance z of the medium respectively.

Modelling χ either classically (equating the nucleus and an outer shell electron to a harmonic oscillator) or quantum mechanically allows the form of $\chi^{'}(\omega)$ and $\chi^{''}(\omega)$ to be found. For a resonance centred at ω_0, the real and imaginary components of χ will exhibit the form shown in figure 2.1.

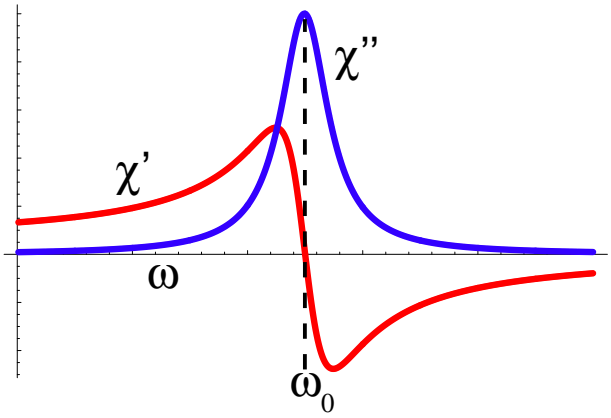

Fig. 2.1: Real ($\chi^{'}$) and imaginary ($\chi^{''}$) components of susceptibility around a resonance at ω_0

2.1.1 Kramers-Kronig relations

The slowing of light in a medium, which is quantified by refractive index (proportional to $\chi^{'}$), is attributed to the absorption and subsequent re-emission of light. It follows then that refractive index and absorption (or $\chi^{'}$ and $\chi^{''}$) should be related. This was established and formalised by Kramers[1] and Kronig[2] and can be expressed as in equation 2.8:

$$\chi^{'}(\omega) = \frac{P}{\pi}\int_0^\infty \frac{\chi^{''}(\omega')}{\omega' - \omega}d\omega' \qquad (2.8)$$
$$\chi^{''}(\omega) = -\frac{P}{\pi}\int_0^\infty \frac{\chi^{'}(\omega')}{\omega' - \omega}d\omega'$$

where P is the Cauchy principal value (the limit of an indefinite integral). The Kramers-Kronig (or KK) relations are a specific example of a relationship between the real and imaginary components of holomorphic functions (sometimes referred to as analytic functions). The interrelatedness of $\chi^{'}$ and $\chi^{''}$ stems from a causality argument, which states that no output can occur before the input. The physical significance of this is discussed in the context of optics by Toll [3]. He uses as an example an optical pulse incident on a perfect filter. The filter absorbs one frequency component of the pulse and if the remainder

of the pulse were to remain unaffected this would imply that components of the pulse would exit the filter before the pulse reached the filter. To satisfy causality the remaining frequency components of the pulse must experience a phase shift as the filter absorbs. This example demonstrates that absorption and refractive index are related (replacing $\chi^{'}$ with $(\eta - 1)$ and $\chi^{''}$ with κ in equation 2.8 gives the relevant relations).

The relationship defined by equation 2.8 between the real and imaginary components of χ (or n) is significant in that $\chi^{'}$ (or η) can be determined if $\chi^{''}$ (or κ) is known. This offers an alternative to the relatively complicated task of measuring $\eta(\omega)$ directly; measure the material absorption ($\kappa(\omega)$) as a function of frequency and apply the Kramers-Kronig relation to deduce dispersion ($\eta(\omega)$).

As the integral limits imply, to determine the real component of n from the imaginary component requires knowledge of absorption over all frequencies, which is not realistic experimentally. It is valid however to measure absorption over a large frequency range and to model the behaviour outside this range, although this must be done with care [4]. From a practical perspective this is very useful because it is often simpler to accurately measure absorption over a large frequency range than to measure refractive index.

There are limitations to the applicability of using the KK relations to determine the real part of n from the imaginary part. The most significant is the presence in the absorption spectra of 'zeros'; resonances due to the structure of the material or device under investigation (such as the frequency dependent reflectivity changes seen in interferometers [5]). These artifacts cause equation 2.8 to break down, and additional factors are needed to augment equation 2.8 to accurately describe the KK relations [6]. For a medium with no zeros, the Cauchy principle value will have a single term. Each resonance will correspond to a singularity in P and without compensatory terms, P would be undefined at the resonance. Compensating for zeros is possible even in complicated cases, such as when nonlinearities are present [7]. Unfortunately these compensatory factors require the precise location of the zeros, which can typically only be found through modelling (if the material / structure is well understood) or by measuring the real part of n. For the purposes of investigating materials

or structures which are not already well characterised, this renders the KK relations largely redundant, although they can still be valuable in verifying experimental data is valid. If the experimentally determined $\kappa(\omega)$ and $\eta(\omega)$ do not conform to KK relations then this indicates flawed measurements (since this implies causality is not obeyed).

2.2 Dispersion of Optical pulses

All optical pulses, regardless of the method of their creation, must satisfy one fundamental criterion: their minimum temporal length is inversely proportional to their bandwidth. Taking the Fourier transform of a pulse allows this to be shown clearly.

A simple example of an optical pulse could be modelled as the product of a plane wave and a pulse shape (i.e. amplitude function) which falls asymptotically to zero with increasing time from the pulse centre. For illustrative purposes choosing the simple case of a Gaussian amplitude profile and designating the centre of the pulse as t=0, the pulse will be described by;

$$\mathbf{E}(t) = A \exp^{-t^2/b^2} \exp^{-\imath\omega_0 t} \tag{2.9}$$

where A is a scalar amplitude, b is related to pulse width (the FWHM, $t_0 = b\sqrt{\ln 2}$) and ω_0 is the centre frequency of the pulse. Knowing $\mathbf{E}(t)$ inherently means knowing the electric field as a function of frequency; the two are equivalent and interchangeable through a Fourier transform.

$$\begin{aligned} \mathbf{E}(\omega) &= \textstyle\int_{-\infty}^{\infty} (A \exp^{-t^2/b^2} \exp^{-\imath\omega_0 t}) \exp^{\imath\omega t}\, dt \\ &= A\tfrac{b}{\sqrt{2}} \exp^{\frac{-b^2}{4}(\omega-\omega_0)^2} \end{aligned} \tag{2.10}$$

Comparison of the expressions for $\mathbf{E}(\omega)$ and $\mathbf{E}(t)$ show that pulse duration is inversely proportional to pulse bandwidth. The relationship between temporal and spectral width is not as simple for all pulses. The pulse described by equations 2.9 and 2.10 is a theoretical ideal; it has a linear phase with respect

to time (which corresponds to a linear phase with respect to frequency). This is equivalent to stating b is purely real; the imaginary component of b quantifies deviation from linear phase with respect to frequency (equation 2.10) and time (equation 2.9). Pulses of this type are referred to as "transform limited", because they are the most efficient pulses, i.e. possess the narrowest duration possible for a given bandwidth and pulse shape.

Were the variable, b in equation 2.9 to have a nonzero imaginary term, the pulse would not alter temporally, but the perturbation of phase would create additional frequency components, increasing the pulse bandwidth. The corollary of this is that if two pulses have identical bandwidths and pulse shapes, but one has purely real b and the other complex b, the former will have a smaller temporal duration than the latter.

Any redistribution of frequencies in a transform limited pulse (i.e. change of phase profile or creation of imaginary part of b) will correspond to temporal broadening. Similarly nonuniform absorption will, as predicted by the KK relations, temporally broaden a transform limited pulse.

Conversely to shorten a pulse requires reorganising the frequency distribution to be more uniform and/or creating additional frequency components. The former is termed *dispersion compensation* and the latter is usually carried out through the nonlinear process of self phase modulation (SPM).

The precise effect of dispersion on a pulse can be well understood by separating an induced phase change in to the terms of a Taylor series, as shown in equation 2.11, where L represents the path length of light in a medium. The speed with which a pulse will be compressed, expanded or broken up by a dispersive medium will depend on the initial phase profile of the pulse and the relative strength of the terms in equation 2.11.

$$\begin{aligned} \phi &= L\tfrac{\eta(\omega)\omega}{c} = L\tfrac{\omega}{c}\sqrt{1+\chi} \\ \phi &= \phi_0 + \tfrac{d\phi}{d\omega}[\omega-\omega_0] + \tfrac{d^2\phi}{d\omega^2}\tfrac{[\omega-\omega_0]}{2!} + \tfrac{d^3\phi}{d\omega^3}\tfrac{[\omega-\omega_0]}{3!} + ... \end{aligned} \tag{2.11}$$

The first term, ϕ_0, is a phase offset, which (being purely relative) has no meaning in the absence of interference with light from another source. The

second term (or more precisely the $\frac{d\phi}{d\omega}$ component) is termed *group delay*(τ_g), and is a measure of the overall time taken for a pulse to propagate through the medium. Obviously this depends on the length of the medium and so in practice group delay per unit length is sometimes quoted. Group delay quantifies the propagation time of the centre (i.e. spectrally weighted average) of the pulse, but does not indicate any changes in the pulse itself.

The second derivative, $\frac{d^2\phi}{d\omega^2}$, quantifies the spread in time between different frequency components of the pulse. It is termed *group delay dispersion* (GDD), and defines the speed with which a pulse will be dispersed and ultimately broken up in a given medium. Higher order terms act similarly to GDD but give rise to more complicated spectral phase. Higher order terms tend to be weaker than GDD, but remain significant for long propagation distances and broad bandwidth (i.e. temporally short) pulses.

Another popular convention is to define dispersion through group velocity (defined as $v_g = \frac{d\omega}{dk}$), and group velocity dispersion (GVD, defined as $\frac{dv_g}{d\omega} = \frac{d^2\omega}{dk^2}$,). If the propagation distance is known, both GDD (or τ_g) and GVD (or v_g) convey the same information; they are related through

$$\tau_g = \frac{d\phi}{d\omega} = L\frac{dk}{d\omega} = \frac{L}{v_g}$$

and

$$GDD = L\frac{d^2k}{d\omega^2} = L\frac{d}{d\omega}\frac{dk}{dv_g}\left(\frac{dv_g}{d\omega}\right) = L\frac{d}{dv_g}\frac{dk}{d\omega}GVD = -\frac{L}{v_g^2}GVD$$

It is important to remember that in addition to the frequency variation of $\eta(\omega)$ (i.e. $\chi^{'}(\omega)$), pulses will also be restructured/distorted by absorption or gain ($\kappa(\omega)$ or $\chi^{''}(\omega)$). In addition to these purely linear material characteristics, when considering ultrashort pulses, consideration must also be given to nonlinearities as will be discussed in section 2.3.

2.2.1 Dispersion compensation

Dispersion has been introduced as a material property, following from Maxwell's equations. The dispersion characteristics of isotropic media will be discussed

briefly before non material dispersion is introduced.

All media will exhibit dispersion and other than over narrow frequency ranges centred on resonances, this material dispersion will be *normal*, meaning that $\frac{d\tau_g}{d\omega} > 0$. *Anomalous* material dispersion (where $\frac{d\tau_g}{d\omega} < 0$) is exhibited around resonances, such as shown in figure 2.1. As the figure suggests this is always associated with an increase in absorption. It follows that pulses travelling through media with low absorption will always undergo normal dispersion. A transform limited pulse having propagated through high transmission media will undergo normal dispersion and spread temporally accordingly. Ideally to correct for this the pulse should be subjected (either before or after transmission) to anomalous dispersion of equal magnitude, restoring the pulse to its former transform limited state. As implied above, this dispersion compensation cannot be carried out in an isotropic medium without suffering absorptive losses.

To create a system with anomalous dispersion without significant absorption requires that instead of using isotropic media, structural or material changes must be used. Many examples of this exist, but all are based on the same concept; low frequency light should propagate over a longer optical path than high frequency light. A straightforward example of this is the use of prisms in ultrashort pulse lasers. Broad bandwidth pulses are created through nonlinearities, but these are highly *chirped*, meaning that the phase is far from linear with respect to time and frequency. Anomalous dispersion is provided by a set of prisms as shown in figure 2.2, first demonstrated by Fork et al[8].

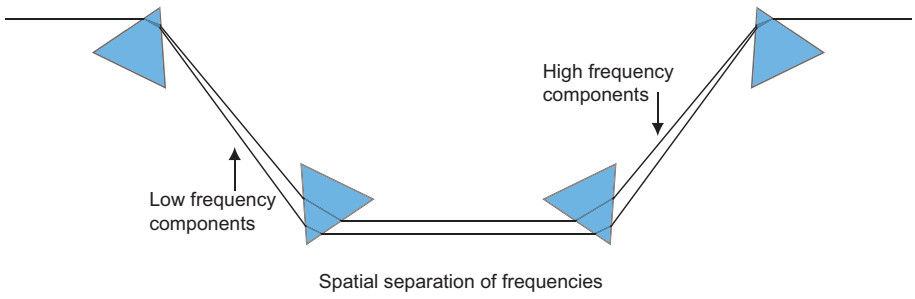

Fig. 2.2: A set of four prisms used for dispersion compensation of chirped pulses

Light incident on the first prism is spatially separated according to frequency. Low frequency components travel along a longer path, and thus accumulate greater phase than high frequency components. Gratings can similarly be used to spatially separate light according to frequency and apply anomalous dispersion, although prisms have the advantage of offering easily tuneable dispersion. The direction of light exiting the prisms is unaffected by the length of (normal dispersion) glass it propagates through and so the net anomalous dispersion generated by the prisms can be increased or decreased by decreasing or increasing the length of glass the beam passes through. Prisms can be used to fully compensate for GDD, but are rarely [9] able to counteract the effect of higher order dispersion terms.

Commonly used optical elements with anomalous dispersion include chirped Bragg reflectors (also called chirped mirrors) and other similar interferometers. In the case of the former, engineering these involves designing a structure which has greater reflectivity per repeating unit (i.e. pair of layers with different refractive indexes) for high frequencies than for low frequencies. To accumulate the same total reflectivity (of near unity) the low frequency components must travel further in to the Bragg reflector, accumulating greater phase. Other interferometers work similarly, taking advantage of the frequency dependent reflectivity to apply a larger (average) delay to low frequencies than to high frequencies.

Optical fibres (whether glass or plastic) can compensate for material dispersion by taking advantage of the boundary conditions of modes of light of different frequencies. Careful engineering can ensure that high frequencies have a greater proportion of their mode contained within the cladding layer (lower refractive index) than the guiding layer (higher refractive index) than is the case for lower frequencies. Essentially the material frequency dependent refractive index (or *material dispersion*) is compensated for by the frequency dependent mode index (*waveguide dispersion*). In common with prisms, dispersion compensated fibres cannot be used to offset dispersion of higher orders than GDD. This determines the maximum propagation distance before signal regeneration is needed.

Other systems of applying anomalous (and in some cases arbitrarily tune-

able [10]) dispersion exist, but like prisms, gratings and Bragg reflectors their use is typically confined to the production of ultrashort pulses rather than dispersion compensation of pulses being transmitted over significant distances.

2.3 Optical Nonlinearities

When generating and transmitting ultrashort optical pulses nonlinearities are unavoidable and essential. The creation of pulses with femtosecond durations is only possible through exploiting nonlinearities.

Some nonlinearities are inherently dispersive, while others will indirectly alter the dispersive properties of a given material or structure. Both of these cases can be useful, although nonlinearities are rarely associated with dispersion compensation. Nonlinearities are often associated with increased absorption, which limits the choice of materials suitable for transmitting and altering ultrashort pulses.

Nonlinearities are rarely present in isolation and the combined effects of multiple nonlinearities give rise to high order dispersion, which is complicated or impractical to compensate for.

In the previous sections of this chapter it has been assumed implicitly that the value of χ has no dependence on electric field strength. At sufficiently high intensities the linear relationship in equation 2.6 no longer holds and a Taylor series is used to model the dependence of polarisation on electric field:

$$\mathbf{P} = \epsilon_0(\mathbf{E}\chi_{(1)} + \mathbf{E}.\mathbf{E}\chi_{(2)} + \mathbf{E}.\mathbf{E}.\mathbf{E}\chi_{(3)} + ...) \tag{2.12}$$

The coefficients $\chi_{(n)}$ are tensors of rank $n+1$ (i.e. $\chi_{(1)}$ and $\chi_{(3)}$ have (3×3) and $(3 \times 3 \times 3 \times 3)$ components respectively). Successive χ terms fall significantly in magnitude; for non-centrosymmetric materials, $\chi_{(1)}$, $\chi_{(2)}$, $\chi_{(3)}$ have approximate relative strengths of 1, 10^{-8}, 10^{-16} (for centrosymmetric materials $\chi_{(n)}$=0 for all even n). Because of this, deviations from the linear approximation are only seen for high intensities. Each order of χ describes different optical processes; $\chi_{(1)}$ quantifies dispersion and linear absorption (i.e. $\chi_{(1)}$ fully defines η and κ if it is assumed that the complex refractive index is

intensity independent), $\chi_{(2)}$ governs the strongest nonlinear processes; second harmonic generation (SHG), sum and difference frequency mixing (SFM and DFM), optical parametric generation, and the Pockels effect. The $\chi_{(3)}$ processes of two photon absorption (TPA), self phase modulation (SPM), and the optical Kerr effect may be weaker than $\chi_{(2)}$ based nonlinearities, but are more prevalent because even orders of χ are zero for centrosymmetric media, due to their inversion symmetry, unlike odd orders of χ. Consequently, semiconductors such as Si and Ge exhibit no $\chi_{(2)}$ nonlinearities, but other common semiconductors (such as the zinc-blende structured GaAs and InAs) will.

Crystal symmetry not only effects which nonlinearities a material can exhibit, but also the complexity (i.e the number of independent, non-zero tensor components) of each χ term. Modelling nonlinearities is made considerably easier by applying symmetry arguments, such as piezoelectric contraction, Kleinman symmetry and crystal symmetry. Kleinman symmetry allows the 27 components of $\chi_{(2)}$ to be reduced to a maximum of 10 (dependent on crystal structure), but this reduction is based on the assumption that the nonlinearities under consideration exhibit no frequency dependence. When considering transmission through dispersive media, this is invalid and modelling $\chi_{(2)}$ nonlinearities becomes more complex. The tensor components of $\chi_{(n>1)}$, sometimes referred to as nonlinear coefficients, have been well catalogued for the majority of direct bandgap semiconductors (accepting variation dependent on material quality) and commercially available fibres. When considering isotropic or bulk media, these values allow accurate modelling of nonlinearities. For nanostructures or media of variable composition or unknown growth quality, experimental measurements are required to determine the effective nonlinear coefficients. Indeed, several papers have been published demonstrating experimental enhancement of nonlinearities due to photonic crystal fabrication [11].

While the behaviour of individual nonlinearities is well understood, their interplay and the variances of material and device quality require experimental measurements paired with theoretical modelling to iteratively improve fundamental understanding and device performance.

It is informative to briefly describe nonlinearities which are used in the

creation of ultrashort pulses and also those which are responsible for pulse dispersion.

2.3.1 $\chi_{(2)}$ Nonlinearities - Three wave mixing

The $\chi_{(2)}$ processes of optical parametric generation, second harmonic generation (SHG) and sum and difference frequency mixing (SFM and DFM) are responsible for the transfer of energy between light of different frequencies. Occurring in non-centrosymmetric media only, the magnitude of these nonlinear processes is determined by intensity and phase matching. Phase matching refers to the situation where different frequencies of light remain in phase over a significant interaction length, i.e the phase velocities at the different frequencies are equal. By considering conservation of energy and momentum, the prerequisite conditions for three wave mixing can be expressed as

$$\eta_1\omega_1 = \eta_2\omega_2 \pm \eta_3\omega_3 \tag{2.13}$$

where η_1 is the refractive index at frequency ω_1 etc. Optical parametric generation corresponds to incident light at frequency ω_1 being converted to light at frequencies ω_2 and ω_3 (often referred to as signal and idler). SFM and DFM correspond to the generation of light at $\omega_1 = |\omega_2 - \omega_3|$ and $\omega_1 = \omega_2 + \omega_3$ respectively (for incident light at ω_2 and ω_3, where $\omega_2 \neq \omega_3$). SHG is the degenerate case where light of frequency $\omega_2 = \omega_3$ is converted to light of frequency $\omega_1 = 2\omega_2$.

Unlike $\chi_{(3)}$ nonlinearities, which will always occur given sufficient optical intensity, the phase matching requirement prevents three wave mixing from being significant except in rare cases (i.e. when phase matching is deliberately applied).

Dispersion causes the group velocity of light of different frequencies to separate temporally, or *walk off* as they propagate through a medium. To achieve high conversion efficiencies (which implies long interaction lengths) the different frequencies must be both in phase and temporally overlapped. Such phase matching can be achieved through birefringence or quasi phase matching (QPM).

The polarisation dependence of refractive index in a birefringent crystal can be used to compensate for material dispersion. By polarising the high frequency light (which sees a higher material $\eta(\omega)$) such that it experiences a lower refractive index due to birefringence than lower frequency light (of orthogonal polarisation), the condition specified by equation 2.13 can be met. By varying the polarisation of both frequencies (while maintaining their orthogonality), the magnitude of relative offset in η between high frequency and low frequency can be increased or decreased. In this way, the light generated by three wave mixing in birefringent crystals can be frequency tuned. Despite offering theoretically higher conversion efficiencies than QPM, birefringent phase matching suffers from two significant drawbacks. Firstly, the phase matching condition is only truly met for single frequencies; birefringent phase matching is not suitable for broad bandwidth sources (such as ultrashort pulses). Secondly, the wavevector, k and the Poynting vector (i.e. direction of energy flow) do not typically coincide (this occurs only if the beam propagates perpendicular to the optic axis). This limits the interaction length, even well before full separation of the incident and generated beams.

QPM does not rely on an intrinsic material property, but rather uses an engineered periodic structure. The phase matching condition (equation 2.13 corresponds to,

$$\Delta k = k_1 - k_2 - k_3 = 0 \tag{2.14}$$

where $k_1 = \frac{\omega_1 \eta_1}{c}$ etc. In an isotropic material Δk cannot be zero, and will be finite. Physically this corresponds to a conversion efficiency which oscillates with interaction length; over a period of $l = \frac{2\pi}{\Delta k}$ the phase difference between frequencies will rise from zero to π and fall back to zero again. Accordingly, the conversion efficiency rises during the first half period, and returns to zero in the second half period. If a correction to phase is made each time the phase difference reaches π then conversion efficiency will continue to rise (although in an oscillatory manner). The phase difference between the interacting frequencies cannot be corrected, but the polarity of nonlinear coefficient (which is determined by crystal orientation) can. By alternating the crystal orientation every $\frac{l}{2}$, the correct phase relationship between crystal and different frequencies

for increasing conversion efficiency will be maintained. QPM is theoretically limited to lower conversion efficiency than birefringent phase matching (by a factor of $\frac{\pi}{2}$, but does not suffer from the latter's limitations. Without restrictions on polarisation of incident light, or useable wavelength range (within the crystal's transmission window), QPM can more efficiently perform three wave mixing over broader spectral ranges in materials with higher optical damage thresholds. It should be noted that although separate frequencies can be phase matched (through birefringence or QPM), this does not imply that the group velocities at each frequency are the same. Phase matching corresponds to matching phases spatially only. A limiting factor when performing three wave mixing with pulses is that the incident and generated light will overlap only for a limited time (i.e. distance). Conversion efficiency can be increased by increasing light intensity and the limiting factor in this case is the damage threshold of the crystal. Despite phase matching limitations, for the generation of ultrashort, high power pulses, tunable over a wide spectral range, three wave mixing is an extremely useful and versatile tool.

In cases where a medium has not been carefully engineered, dispersion has the desirable effect of drastically reducing the efficiency of three wave mixing. Consequently the $\chi_{(2)}$ nonlinearities very rarely diminish pulse energy, unless deliberately chosen to.

The optical parametric oscillators (OPO)s used as sources of ultrashort pulses in this thesis relied on QPM. For the first OPO used, a combination of a commercial frequency doubled Nd:YVO_4 CW laser and a Kerr lens modelocked Ti:Al_2O_3 laser produced pulses of ∼800fs at $\lambda \simeq$840nm with a repetition rate of ∼80MHz. These pulses then underwent optical parametric generation in a synchronously pumped (i.e. the pump pulse and generated signal propagated together through the nonlinear crystal), singly resonant (i.e. only the signal frequency was preserved, the idler experienced no feedback) OPO. The OPO used a periodically poled Li:NbO_3 (PPLN)[12] crystal to generate pulses tunable from ∼1310nm to ∼1360nm and ∼1500nm to ∼1570nm (dependent on the PPLN grating period). A series of prisms, as shown in 2.2, provided dispersion compensation, resulting in transform limited Gaussian pulses ∼700fs in duration.

The second OPO was based on identical components, other than that the Ti:Al_2O_3 laser produced pulses of ∼100fs. This source was used because it offered a slightly different, and therefore complementary, range of wavelength tunability (∼1290nm to ∼1350nm). The increased bandwidth also allowed for measurement over larger spectral ranges between wavelength tuning.

2.3.2 $\chi_{(3)}$ Nonlinearities - Optical Kerr effect

One of the most useful and fundamental nonlinearities is the optical Kerr effect, quantified by $\chi_{(3)}$. The optical Kerr effect is the intensity dependent change in refractive index, often approximated by

$$\eta_{eff} = \eta_0 + \eta_2 \mathbf{E}^2 = \eta_0 + \left(2\sqrt{\frac{\mu}{\epsilon}}\right)\eta_2 I \tag{2.15}$$

where

$$\eta_2 = \frac{3\chi_{1111}}{2\epsilon_0 \eta_0}$$

η_{eff} is the effective refractive index experienced by light of intensity I and η_0 is the real component of refractive index (η) at zero intensity. The coefficient η_2 is positive and material dependent (generally materials with high η have large values of η_2).

The simple relationship between I and η_{eff} is the basis for several observed nonlinear effects, namely self-focussing (the basis of Kerr lens modelocking), self phase modulation and self steepening.

2.3.2.1 Self focussing

Self focussing refers to the spatial dependence of refractive index, brought about by the spatial intensity dependence (i.e. beam profile) of a pulse. Perpendicular to propagation direction, an ultrashort pulse will have an intensity profile, typically a central maximum which falls symmetrically with distance in all directions (i.e. a TEM_{00} mode),

$$\eta(r) = \eta_0 + \left(2\sqrt{\frac{\mu}{\epsilon}}\right)\eta_2 I(r) \tag{2.16}$$

where r is the radial distance from the centre of the beam. The intensity induced refractive index change brought about by such a pulse will form a lens, with focal length inversely proportional to pulse intensity.

Self focussing behaviour is used to provide an intensity dependent loss mechanism in lasers, through the addition of an aperture (either a physical aperture or a spatial variation of gain). Modelocking then occurs because low intensity pulses (being more spatially diffuse) experience greater net loss than high intensity pulses. This is termed Kerr lens modelocking [13][14], and is the basis for the majority of high power, ultrashort pulse lasers.

2.3.2.2 Self Phase Modulation

Self phase modulation (SPM) is an analogue of self focussing, it is the name given to the change to a pulse in frequency resulting from the optical Kerr effect. Just as the spatial intensity profile of a pulse is responsible for self focussing, the temporal intensity profile of a pulse is responsible for SPM.

$$\eta(t) = \eta_0 + \left(2\sqrt{\frac{\mu}{\epsilon}}\right)\eta_2 I(t) \tag{2.17}$$

Figure 2.3 shows a pulse with a Gaussian intensity profile in time. Because intensity is directly proportional to change in refractive index, the Gaussian profile quantifies the nonlinearity's contribution to phase. The derivative of this phase is an intensity dependent frequency shift experienced by the pulse ($\Delta\omega = -\frac{d\phi}{dt}$).

It can clearly be seen in figure 2.3 that the leading edge of the pulse will be red shifted (i.e. it will be reduced in frequency), while the trailing edge will be blue shifted. SPM gives rise to new frequency components, broadening pulse bandwidth, but does not alter the temporal intensity profile. Although the pulse has an increased bandwidth, its distribution of frequencies will be skewed by SPM; the pulse will effectively experience normal dispersion. New frequencies are created by SPM, but this does not correspond to energy being added to the pulse. The frequency distribution of energy under the pulse is altered, but the total energy remains constant. Mathematically SPM is equivalent to adding an imaginary term to b in equations 2.9 and 2.10.

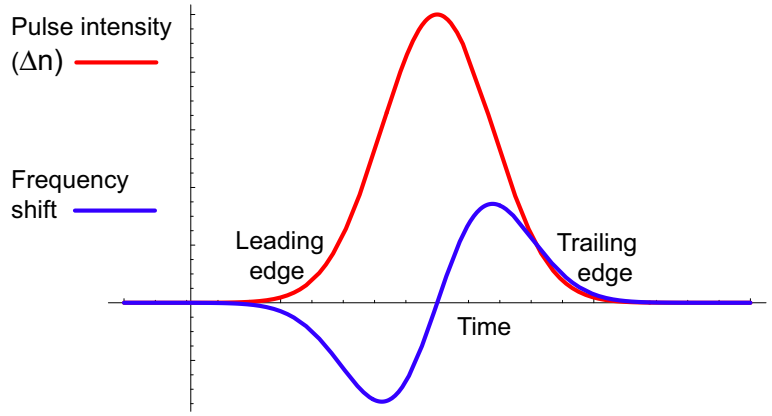

Fig. 2.3: Temporal pulse intensity and induced frequency shift due to Kerr effect

Any pulse (even transform limited) undergoing SPM can be shortened temporally by dispersion compensation, due to the increase in bandwidth. The intensity profile will determine the dispersion of the resultant pulse and so complicated or asymmetric intensity profiles will require correspondingly more complex (i.e. higher order) dispersion compensation.

2.3.2.3 Self steepening

Having briefly described the ramifications of equation 2.15 on the frequency and spatial profiles of pulses, it may be unsurprising to realise that the Kerr effect alters pulses temporally also. The temporal distortion of pulses due to the relationship in equation 2.17 is named self steepening.

Consider the effects on temporal pulse profile of figure 2.3; the refractive index at the pulse centre will be higher than in the wings of the pulse: group velocity is intensity dependent. As the pulse propagates, the trailing edge of the pulse will catch the peak, becoming steeper and effectively shifting the peak backwards in time. The leading edge will travel faster than the peak and move further from it, giving rise to an asymmetric pulse shape.

The asymmetry in pulse shape corresponds to an asymmetry in spectrum. Also, due to the increased intensity, SPM will cause greater spectral broadening in the trailing edge of the pulse. SPM will broaden the pulse spectrum more

at the high frequency side than at the low frequency side. Equal energy will be redistributed to additional high and low frequency components; the trailing edge will contain a broad range of lower intensity high frequency components than the leading edge.

The distinction between self steepening and SPM is somewhat artificial, one cannot be present without the other. In many situations however, SPM will have a significant magnitude and self steepening will be nearly negligible. Self steepening becomes increasingly obvious (and problematic) for short pulse durations, typically < 100fs.

2.3.3 Nonlinear absorption

Photons with lower energy than a material bandgap (E_g)are usually regarded as suffering no absorption. In the case of low intensities this assumption is negligibly far from true. A single photon may excite an electron from the valence band, but will have insufficient energy to promote the electron to the conduction band. The electron cannot occupy any state in the bandgap and so falls back to the valence band, releasing a photon of identical energy to the original in the process. The net effect is a delay for the photon (quantified by η) and no absorption.

There is a chance that two photons will strike the same electron nearly instantaneously (i.e. before the electron can relax). In this case if $2\hbar\omega \geq E_g$ the electron will be promoted to the conduction band, by two photon absorption (TPA). Intuitively it can be realised that this process will be dependent on photon density (i.e. intensity) and the ratio of $\hbar\omega : E_g$. TPA is quantified by $\chi_{(3)}$, and as such occurs in any medium where $\frac{E_g}{2} \leq \hbar\omega \leq E_g$ (regardless of symmetry), given sufficient optical intensity. Indirect bandgap semiconductors such as Si, tend to have low (temperature dependent) TPA coefficients because two photons and a phonon are required by the absorption mechanism. Absorption with three or more photons (where n is the number of photons and $\frac{E_g}{n} \leq \hbar\omega \leq \frac{E_g}{n-1}$) is possible also, but the coefficients for higher order multiple photon absorption are orders of magnitude smaller than for TPA. Consequently, even when considering ultrashort pulses, very high pulse energies

must be used before three (or higher) photon absorption become significant.

Concentrating a pulse spatially (due to waveguiding or focussing with a lens) or temporally (as can occur in interferometers or photonic crystals) exacerbates TPA. When using a nonlinear detector or attempting to increase the strength of a nonlinearity (in wavelength conversion or switching[15] for example), this is beneficial, providing an increase in sensitivity. If absorption is undesirable, and high peak power pulses are used, TPA can only be avoided if $\hbar\omega < \frac{E_g}{2}$. As well as a loss mechanism, multi-photon absorption (as with any absorption) will increase the number of free carriers and hence alter refractive index.

2.4 Carrier dependence of refractive index

Whether due to an applied field across a junction, dopants in a semiconductor or optical absorption, the density of free carriers will always effect absorption and refractive index in a medium. As free carrier density changes, both the slope of absorption curve and effective bandgap change also. The change in absorption and refractive index with free carriers is due to the combination of three distinct effects, band filling, intraband absorption (most significantly free carrier absorption (FCA)) and bandgap shrinkage. Each of these increase with carrier density, and their individual contributions to refractive index are discussed below (for $\hbar\omega < E_g$ only; high absorption/gain regimes are not reviewed here).

Band filling is commonly referred to as the Burstein-Moss shift, and refers to the effective increase in bandgap brought about by the presence of free carriers. For low temperatures, as the number of free carriers increases, the states at the bottom of the conduction band and/or the top of the valence band fill. Consequently to excite an electron from the valence band to the conduction band requires more energy than if no free carriers were present. Band filling lowers absorption (it shifts the absorption edge to higher frequencies) and decreases refractive index, except for photon energies very near the bandgap. Here the change in gradient of the absorption edge shifts such that the refractive index increases due to band filling.

Intraband absorption primarily consists of two processes, interconduction band absorption and FCA. The mechanisms for both processes are phonon assisted. Interconduction band absorption starts with an electron in the lowest conduction band minimum absorbing a photon and being promoted to a higher conduction band. The electron subsequently relaxes, via phonon emission to an intermediate state. The intermediate state is in the band originally occupied by the electron, but separated from the lowest conduction band minimum by the phonon's momentum (i.e. another minimum in the same band offset in both momentum and energy). An analogue of this process (intervalence band absorption) exists where holes are the majority free carriers, but its absorption coefficients disappear if electrons are the majority carriers. Consequently intervalence band absorption need only be considered for p-doped materials. Free carrier absorption also starts with an electron in the conduction band absorbing a photon. In FCA however the electron is promoted to a higher energy within the same band, and after emission or absorption of a phonon, finishes in an excited state. Interconduction band absorption and FCA both cause a decrease in refractive index, despite increasing absorption.

Band gap shrinkage results from the perturbation to lattice potential experienced by electrons due to the presence of other electrons. The mutual repulsion between electrons, particularly at high densities results in a screening of lattice potential (i.e. a decrease in effective binding energy for each electron). Overlap of individual electron wavefunctions, which occurs at high electron densities, causes the behaviour of each electron to be determined by all other electrons. This effect, known as correlation, also contributes to a reduction in effective bandgap. Bandgap shrinkage is responsible for an increase in both absorption and refractive index.

A similar interaction leading to a decrease in effective bandgap occurs between free carriers and impurities. This is sometimes referred to as band tailing, because it creates a tail in the absorption edge. As well as screening lattice potential as with bandgap shrinkage, if a material is doped heavily enough the dopants (impurities) will form a band which will merge with either the conduction or valence band. Band tailing increases refractive index and absorption, but will only be present in heavily doped materials.

The relative strengths of band filling, intraband absorption and bandgap shrinkage and their combined effect on η are determined by the ratio $\hbar\omega : E_g$ and carrier density. Band filling is proportional to carrier density, N. FCA is proportional to N and λ^x, where x varies between 1.5 and 3 (dependent on scattering mechanism). Bandgap shrinkage is proportional to $N^{\frac{1}{3}}$. A paper by Huang et al [16] quantifies and experimentally validates the free carrier induced changes in refractive index for a n-doped GaAs sample. Close to the bandgap, refractive index increased with carrier density due to band shrinkage (and possibly band tailing), until very high densities at which point band filling and FCA began to dominate, reducing and then making $\Delta\eta$ increasingly negative. For more moderate photon energies, the trend is similar, but the relative strengths of band filling and FCA are stronger, causing $\Delta\eta$ to turn negative at lower electron densities.

When considering pulse propagation, it is important to recognise the many variables which determine dispersion, namely nonlinearities, absorption and free carrier effects. The presence of nanostructures further complicates modelling dispersion, as reported by Mondia [17]. She found that in addition to expected changes in the behaviour of a planar photonic crystal waveguide through the Kerr effect, below bandgap excitation also gave rise to free carrier effects. Attempting to quantify dispersion theoretically in such a structure would likely yield specious results. Material quality (i.e. density of defect states), nanostructure fabrication quality and temporal, spectral and spatial pulse characteristics combined with uncertainty over the nanostructure's effect on nonlinearities and free carrier processes would all contribute to dispersion. Quantifying the device's behaviour would require experimental measurements. Existing dispersion measurement techniques are reviewed in the following chapter.

2.5 Photonic crystal waveguides

Photonic crystals are a class of structures which prevent the propagation of light over a given frequency range in either one, two or three dimensions. The frequency range for which light is forbidden to propagate is often termed the

photonic bandgap (PBG), because it is analogous to the electronic bandgap seen in semiconductors. The photonic bandgap arises due to periodic modulation of permittivity on the order of wavelength. Photonic crystals can be thought of as multidimensional gratings; Bragg scattering can result in total reflection of all incident light over a given frequency range.

The 1-D case corresponds to the familiar Bragg stack or Distributed Bragg Reflector (DBR). The dispersion of such structures can be engineered by varying layer thicknesses. Light is not confined in 1-D PBG structures and while their current role in ultrafast optics (as dispersion compensators in pulsed lasers) is valued, they offer little opportunity for new classes of device.

At the other end of the spectrum, 3-D PBG structures offer unparalleled confinement of light. Current fabrication methods (colloidal growth or forming woodpile structures) do not yet allow waveguiding structures to be constructed with the same level of freedom as conventional semiconductor ridge waveguides.

The intermediate 2-D planar photonic crystal waveguides are an amalgam of conventional waveguides and photonic crystals. If a regular lattice of holes is etched in to a semiconductor, light can not propagate perpendicularly to those holes over a given frequency range (the PBG). By etching an incomplete lattice of holes (for example leaving one row unetched) a waveguide is formed. Waveguides of this type are commonly referred to as WX waveguides, where X is the number of rows of holes absent from the lattice, e.g. a W3 waveguide has three rows of holes un-etched. While the in plane confinement of 2-D photonic crystals is carried out by the waveguide, lateral (i.e. vertical) confinement is provided by refractive index changes; either the photonic crystal is surrounded by air on both sides (known as a membrane structure) or a conventional heterostructure is used. In both cases light is lost out of plane as a result of no lateral confinement in the low refractive index regions, the holes[18]; this effect can be minimised by reducing hole diameter. The roughness of hole edges also contributes to out of plane losses, proportional to $\Delta\eta^3 = (\eta_{semiconductor} - \eta_{air})^3$ (reducing these losses requires improved fabrication or oxidation around hole surfaces). The dispersion of the waveguide will be determined by material, waveguide geometry and photonic crystal geometry. Lossless structures with arbitrary dispersion can be envisaged, given the

parameters available to engineer. At the millimetre scale, such structures have been demonstrated experimentally (using metal rods in air [19]).

Photonic crystals do not have an absolute length scale, their behaviour is described through Maxwell's equations, which contain no fundamental constant of length. The consequence of this is that the properties of a given photonic crystal structure for light at a given wavelength will be identical to the properties of the structure scaled by any factor for light whose wavelength is similarly scaled. Experiments carried out in the millimetre regime can act as a perfect model for the behaviour of a photonic crystal in the optical regime. Optical scale PBG structures are limited by growth and fabrication rather than fundamental physics. As technological advances are made (such as moving from standard lithography to deep UV lithography[20]), smaller and more precise structures can be constructed. An initial limitation to the usefulness of photonic crystal waveguides was their high losses. Advances in fabrication technology however are resulting in reductions in propagation losses [21].

At the photonic band edge, light propagates slowly (group velocities several hundred times lower than c [22] have been demonstrated), which corresponds to very significant dispersion[23]. The gradient of group delay (or group velocity) versus frequency can be engineered as can the position of band edges. Varying hole diameter and position allows the equivalent of impedence matching between different sections of waveguide (for instance at a Y junction or even the interface between a conventional ridge waveguide and a photonic crystal waveguide). Structures with negligible loss due to the photonic crystal interfaces can be constructed in this way [24].

For the purposes of designing dispersion compensators, photonic crystals with separate cavities (i.e. a series of point defects) offer greater control than those with one or more rows of holes left un-etched (i.e. row defects). Waveguides formed from a series of cavities are referred to as coupled cavity waveguides (CCWs) or coupled resonator optical waveguides (CROWs)[25]. The group delay of light propagating from cavity to adjacent cavity is determined by the quality factor, or Q, of each cavity and intercavity spacing (which determines the overlap of evanescent fields from each cavity). These additional degrees of freedom theoretically allow even greater control over the dispersion

imparted to pulses [26]. Hybrid structures which take advantage of both point defect and row defect behaviour are now being designed [27].

Because the dispersion from photonic crystals can be orders of magnitude greater than material dispersion, a photonic crystal millimetres in length could compensate for the dispersion from many kilometres of fibre. The ideal situation of being able to measure the dispersion (to all orders) imparted by a fibre and design and fabricate a structure with exactly opposite dispersion is moved closer with each advance in photonic crystal research. Understandably research in to photonic crystals is of great interest to the telecommunications industry.

2.5.1 Conclusions

This chapter has reviewed the phenomenon of dispersion, highlighting its causes. The dispersive properties of materials, nonlinearities and free carriers have been discussed. Photonic crystals and their potential role as dispersion compensators have been introduced. It is clear that experimental dispersion measurement techniques are necessary to complement modelling and fabrication of potential dispersion compensators. Ideally high sensitivity dispersion measurements could be carried out with a range of sources (ultrashort pulses, broadband incoherent sources etc), for a range of devices (passive and active, with variable transmission) and without the need for overly complicated, expensive or slow equipment.

REFERENCE LIST

[1] H A Kramers, "La diffusion de la lumiere par les atomes", *Atti del Congresso Internazionale dei Fisici, Como*, vol. 2, pp. 545–557, 1927.

[2] R De L Kronig, "On the theory of dispersion of x'rays", *J. Opt. Soc. Am.*, vol. 12, pp. 547, 1926.

[3] J S Toll, "Causality and the Dispersion Relation: Logical Foundations", *Phys. Rev.*, vol. 104, no. 6, pp. 1760–1770, 1956.

[4] K-E Peiponen and E M Vartiainen, "Kramers-Kronig relations in optical data inversion", *Phys. Rev. B*, vol. 44, pp. 8301–8303, 1991.

[5] R Sprik R H J Kop, P de Vries and A Lagendijk, "Kramers-Kronig relations for an interferometer", *Optics Communications*, vol. 138, pp. 118–126, 1997.

[6] K-E Peiponen V Lucarini, F Bassani and J J Saarinen, "Dispersion theory and sum rules in linear and nonlinear optics", *Rivista Del Nuovo Cimento*, vol. 26, no. 12, pp. 1–120, 2003.

[7] D J Hagan D C Hutchings, M Sheik-Bahae and E W Van Stryland, "Kramers-Kronig relations in nonlinear optics", *Optical and Quantum Electronics*, vol. 24, no. 1, 1992.

[8] O E Martinez R L Fork and J P Gordon, "Negative dispersion using pairs of prisms", *Opt. Lett.*, vol. 9, no. 5, pp. 150–152, 1984.

[9] B E Lemoff and C P J Barty, "Cubic-phase-free dispersion compensation in solid-state ultrashort-pulse lasers", *Optics Letters*, vol. 18, no. 1, pp. 57–59, 1993.

[10] V Binjrajka, D E Leaird C-C Chang, A W R Emanuel, and A M Weiner, "Pulse shaping of incoherent light by use of a liquid-crystal modulator array", *Opt. Lett.*, vol. 21, no. 21, pp. 1756–1758, 1996.

[11] S Fan M Ibanescu E Ippen M Soljacic, S G Johnson and J D Joannopoulos, "Photonic-crystal slow-light enhancement of nonlinear phase sensitivity", *Journal Opt. Soc. Am. B*, vol. 19, no. 9, pp. 2052–2059, 2002.

[12] A Miller W Sibbett D T Reid, G T Kennedy and M Ebrahimzadeh, "Widely tunable, near- to mid-infrared femtosecond and picosecond optical parametric oscillators using periodically poled $LiNbO_3$ and $RbTiOAsO_4$", *IEEE J. Select. Topic. Quantum Electron.*, vol. 4, no. 2, pp. 238–248, 1998.

[13] P N Kean D E Spence and W Sibbett, "60-fsec pulse generation from a self-mode-locked Ti:sapphire laser", *Optics Letters*, vol. 16, no. 1, pp. 42–44, 1991.

[14] S De Silvestri G Cerullo and V Magni, "Self-starting Kerr-lens mode locking of a Ti:sapphire laser", *Optics Letters*, vol. 19, no. 14, pp. 1040–1042, 1994.

[15] A Haché and M Bourgeois, "Ultrafast all-optical switching in a silicon-based photonic crystal", *Appl. Phys. Lett.*, vol. 77, no. 25, pp. 4089–4091, 2000.

[16] Yee Huang and Soma, "The carrier effects on the change of refractive index for n-type GaAs at lambda =1.06, 1.3, and 1.55μm", *J. Appl. Phys.*, vol. 67, no. 3, pp. 1497–1503, 1990.

[17] S Linden H M van Driel J P Mondia, H W Tan and J F Young, "Ultrafast tuning of two-dimensional planar photonic-crystal waveguides via freecarrier injection and the optical Kerr effect", *J. Opt. Soc. Am. B*, vol. 22, no. 11, pp. 2480–2486, 2005.

[18] L C Andreani and M Agio, "Intrinsic diffraction losses in photonic crystal waveguides with line defects", *Appl. Phys. Lett.*, vol. 82, no. 13, pp. 2011–2013, 2003.

[19] T J Gmitter E Yablonovitch and K M Leung, "Photonic band structure: The face-centered-cubic case employing nonspherical atoms", *Phys. Rev. Lett.*, vol. 67, no. 17, pp. 2295–2298, 1991.

[20] D Taillaert S Beckx B Luyssaert P Bienstman W Bogaerts, V Wiaux and R Baets, "Fabrication of photonic crystals in silicon-on-insulator using 248-nm deep UV lithography", *IEEE Selec. Topic. Quantum Elec.*, vol. 8, no. 4, pp. 928–934, 2002.

[21] N Ikeda Y Nakamura K Asakawa Y Sugimoto, Y Tanaka and K Inoue, "Low propagation loss of 0.76 dB/mm in GaAsbased single-line-defect two-dimensional photonic crystal slab waveguides up to 1 cm in length", *Optics Express*, vol. 12, no. 6, pp. 1090–1096, 2004.

[22] D Mori and T Baba, "Dispersion-controlled optical group delay device by chirped photonic crystal waveguides", *App. Phys. Lett.*, vol. 85, no. 7, pp. 1101–1103, 2004.

[23] A Shinya J Takahashi C Takahashi M Notomi, K Yamada and I Yokohama, "Extremely Large Group-Velocity Dispersion of Line-Defect Waveguides in Photonic Crystal Slabs", *Phys. Rev. Lett.*, vol. 87, no. 25, pp. 253902, 2001.

[24] M Midrio S Boscolo, C Conti and C G Someda, "Numerical analysis of propagation and impedance matching in 2D photonic crystal waveguides with finite length", *Journal of Lightwave Technology*, vol. 20, no. 2, pp. 304–310, 2002.

[25] S Mookherjea and A Yariv, "Pulse propagation in a coupled resonator optical waveguide to all orders of dispersion", *Phys. Rev. E*, vol. 65, no. 056601, 2002.

[26] W Kuang W J Kim and J D OBrien, "Dispersion characteristics of photonic crystal coupled resonator optical waveguides", *Optics Express*, vol. 11, no. 25, pp. 3431–3437, 2003.

[27] S Lan and H Ishikawa, "Broadband waveguide intersections with low cross talk in photonic crystal circuits", *Optics Letters*, vol. 27, no. 17, pp. 1567–1569, 2002.

3. COMPARISON OF EXISTING TECHNIQUES

3.1 Introduction

Few articles discuss experimental pulse dispersion in semiconductor devices, although many publications demonstrate measurements carried out using broadband incoherent sources[1][2], or of optics designed for use in ultrashort pulse production or manipulation[3].

In attempting to construct (or adapt) a new technique with which to accurately measure pulse dispersion in semiconductor nanostructures, it is logical to critically analyse existing techniques. This chapter will review techniques used to investigate the dispersive properties of optical elements and techniques used to measure pulses. The former group rely predominantly on linear detection, while the latter require a nonlinear measurement. Both linear and nonlinear detection are required to measure the phase profile of an isolated pulse (i.e. one without a reference) and full pulse characterisation techniques that do this will also be described. The operational requirements, limitations and the information offered in each case will be discussed. This chapter will be biased towards techniques offering high sensitivity and high spectral and phase resolution.

3.2 Linear pulse characterisation

All linear stationary optical components, such as mirrors, prisms, beamsplitters and lenses, or combinations of such components, can be fully described by a transfer function. A transfer function relates an incident electric field to the resultant electric field, as the equation below describes[4];

$$E(t) = \int R(t-t')E_0(t')dt' = \int R(\omega)E_0(\omega)exp(-i\omega t)\frac{d\omega}{2\pi} \qquad (3.1)$$

Here $E_0(t)$ is the original electric field, which is transformed in to $E(t)$ by the linear element(s), represented by the transfer function $R(t-t')$. The transfer function is complex, it describes both the absorption of the device and its dispersion. Often the absorption is seen as unimportant and so the term "phase transfer function" is used to describe the dispersive portion of the transfer function only.

If the field from equation 3.1 were to be measured by an integrating detector (which applies to most detectors in the case of ultrashort pulses), then the signal would correspond to the time or frequency integration of the electric field.

$$S = \int |E(t)|^2 dt = \int |E(\omega)|^2 \frac{d\omega}{2\pi} = \int |R(\omega)|^2 |E_0(\omega)|^2 \frac{d\omega}{2\pi} \qquad (3.2)$$

The detected signal relates only to the integral of the power spectrum, through either $E(\omega)$, or $E(t)$, and the magnitude of the transfer function. As this shows no phase data can be extracted from the pulse with such a measurement. Phase defines the manner in which fields are related to each other: it follows that there can be no such thing as an absolute phase. A reference of some kind is needed before any phase information can be extracted from a field.

If the power spectrum of an electric field is desired, and this can be a very useful complementary measurement as will be discussed later, then two common experimental setups to achieve this are the scanning Michelson interferometer and the monochromator. The monochromator would yield the power spectrum directly. With the interferometer the length of one of the arms would be scanned to give a temporal interference pattern, which must be Fourier transformed to yield the power spectrum.

The interference pattern given by a Michelson interferometer is sometimes referred to as a first order autocorrelation, because a linear detector is used. Using a nonlinear detector gives a second order autocorrelation, typically simply referred to as an autocorrelation (see section 3.4.1).

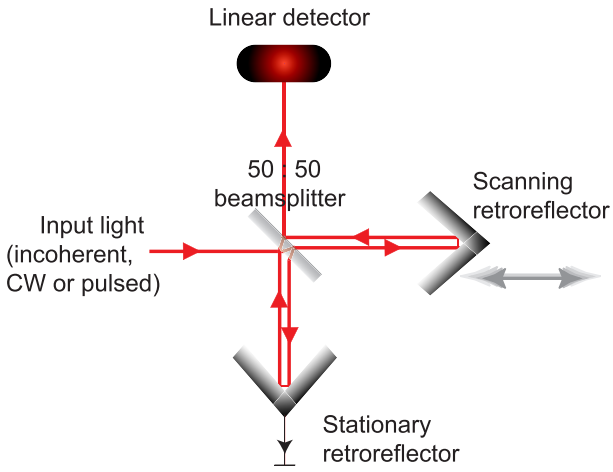

Fig. 3.1: A scanning Michelson-Morley interferometer

If a pulse enters the interferometer and is split evenly in to two lower intensity pulses, then the electric field of the pulse in the static arm will be dependent on time only ($E_0(t)$), while the field in the scanning arm will also have a delay dependence ($E_0(t+\tau)$). After these two pulses recombine at the beamsplitter, the measured signal will be:

$$S = \int |E_0(t) + E_0(t+\tau)|^2 dt \tag{3.3}$$
$$= \int |E_0(t)|^2 dt + \int |E_0(t+\tau)|^2 dt + \int E_0(t) E_0^*(t+\tau) dt + \int E_0^*(t) E_0(t+\tau) dt$$

The measured signal does have phase dependent terms. From this it may be expected that some phase information could be extracted from this setup. Although $E_0(t)$ and $E_0(t+\tau)$ do carry phase information about the original pulse, when they interfere the first two terms have no phase dependence (they are the intensities of the two pulses respectively), and the remaining two terms are complex conjugates of each other, leaving a purely real value of intensity. As this shows, no phase information can be extracted from a pulse with a linear measurement only. To use the power spectrum of a pulse to determine its phase requires interaction with a reference of known phase.

3.3 Linear phase measurements

In some cases the phase profile of the incident pulse is not relevant, but only the induced dispersion from a device is sought. In these cases interferometric techniques can be used to find the phase transfer function of the device. Most spectral interferometry techniques typically yield group delay, but because knowledge of any one of phase transfer function, group delay or GDD (group delay dispersion) allows the others to be calculated, these can all be referred to as phase measurement techniques. All linear phase measurements rely on the interference of light with an altered version of itself, whether the light is pulsed, CW or incoherent; these techniques are all based on first order cross correlations.

Where a device with transfer function $R(t-\tau)$ is placed in one arm of an interferometer the measured signal given will (as in equation 3.3) contain intensity only terms and phase dependent terms. The phase dependent terms are now directly related to the complex transfer function.

$$\int E_0(t)E^*(t+\tau)dt+\int E^*(t)E_0(t+\tau)dt = \int R(t-\tau)|E_0(t)|^2dt+\int R^*(t-\tau)|E_0(t)|^2dt \tag{3.4}$$

Although equation 3.4 relates to a measurement where one arm of an interferometer is scanned (varying τ) and no spectral filtering takes place, this is directly analogous to keeping both arms stationary and spectrally scanning the light entering the interferometer. The equations given in both cases only differ in the variable which each term is a function of (either τ or ω). Consequently the following argument applies equally to both cases.

As before a simple time integrated measurement of this signal yields only information about a purely real intensity value, the power spectrum. However, the complex transfer function, $R(t-\tau)$ or $R(\omega)$, can be extracted from these terms by examining the response of R as a function of delay or frequency. Consequently delay, frequency or both must be scanned while recording the resultant signal. Phase is determined by either relating a delay to every frequency (necessitating the measurement of both variables) or using a Fourier

transform to de-convolve a measured interference pattern. In the latter case either the delay will remain stationary while varying measured frequency or the spectrum will remain constant as the delay is scanned.

Several classes of linear phase measurement exist, although they all exploit the fact that the phase transfer function can be extracted from a measure of the interference of light passing through each arm of an interferometer. The techniques are distinguished by whether they rely on scanning delay, frequency or both, which factors determine accuracy and resolution and what properties (if any) are required of the light source and device to ensure valid results are given.

3.3.1 Compensated Michelson interferometer

This is a conceptually simple phase measurement technique[5] and was the first used to characterise optical components. Light from a broadband source is spectrally filtered before passing in to a Michelson interferometer (as in Figure 3.1). The length of one arm is scanned allowing an interference pattern to be detected. The position of the maximium of the interference pattern is noted and serves as a reference. This position corresponds to the the interferometer being 'balanced', i.e. the optical path length (OPL) in each arm being the same. The device under investigation is inserted in to one arm of the interferometer (or replaces one of the mirrors of the interferometer if the device is reflective). The scanned arm is moved to produce another interference pattern and the new position corresponding to the detected maximum is noted. The difference between the new maximum position and the previous one can be equated to the group delay of the device at the filtered wavelength. The group delay, τ, is simply the difference in arm length divided by the speed of light, c. The spectral filter selects a new wavelength and again the arm is scanned to establish the position which corresponds to the maximum in detected signal. In this way the group delay is quantified for each filtered wavelength.

The simplicity of this technique is advantageous in that there is no ambiguity associated with the group delay values generated. The spectral resolution is determined solely by the filter. The factor that compromises this technique

most is the need to measure delay with a very high accuracy (without sufficient points per wavelength only the envelope of the interferogram will be recorded, giving an incomplete characterisation of the field[6]). Thermally induced fluctuations in the length of both arms are not compensated for, which causes the fringes of the interference pattern to drift with respect to measured delay, reducing fringe contrast.

Despite the limited accuracy of this technique it is suited to measuring highly dispersive samples such as chirped Bragg mirrors, and benefits from being quick to arrange and yield data.

It should be noted that this is very similar to 'time of flight' measurement techniques (see section 3.4.3), which use a non linear detector to find the position corresponding to maximum temporal pulse overlap.

3.3.2 Phase locked interferometery

Phase locked interferometry is closely related to the compensated Michelson arrangement; it could be thought of as the successor to the older technique. The technique again uses a Michelson interferometer, but with the addition of a feedback loop and external reference eliminates the latter method's problem of measuring delay correctly.

Developed by Beck et al[1], this technique uses a feedback loop to track the peak of an interference fringe in real time as a spectral filter placed immediately before the detector scans through a range of wavelengths. As the filtered wavelength is scanned any change in group delay is compensated for by lengthening or shortening the reference arm (i.e. the arm which does not contain the device) accordingly. This is carried out by driving a piezo transducer so that the mirror in the reference arm oscillates continuously over a fixed, narrow delay range (a fraction of a single wavelength). The detector is sampled at the same frequency used to drive the piezo and its signal is fed back to the piezo, altering the offset in its driving voltage. In this manner the piezo's oscillations are continuously centred around the peak of one interference fringe.

To avoid the uncertainty in delay caused by thermal drift that the compensated Michelson suffers from, the delay is not measured using the voltage

offset. Instead a separate external reference is used to quantify delay for each filtered wavelength.

The reference chosen by Beck et al was a He-Ne laser which passed through the interferometer parallel to the beam used to characterise the sample. A second detector monitors the intensity of the He-Ne exiting the interferometer. Because it follows a parallel path to the filtered beam it can be used to quantify changes in relative path length between the two interferometer arms. As the delay between arms changes, both as a function of wavelength (due to the dispersive properties of the device) and of thermally induced drift, the signal resulting from the interference of the He-Ne beams from both arms of the interferometer will rise and fall accordingly. With its nearly monochromatic output, the fluctuations in detected intensity can be equated to distance and hence time (the peak to peak distance will always correspond to λ_{He-Ne}=632.8nm).

Using the He-Ne to quantify delay changes as the spectral filter tunes through its range gives a very accurate (<1fs) measure of spectrally dependent group delay. Importantly using this method also eliminates the effect of any drift in OPL of the interferometer arms, because the He-Ne signal and spectrally filtered signal will be effected equally by any such drift.

Although more complicated, expensive and taxing to align than the compensated Michelson, this is a far more robust and accurate technique. Indeed this technique's method of referencing delay has been incorporated in to commercial Fourier transform spectrometers, such as mentioned in a paper by Kop and Sprik[7].

3.3.3 Frequency domain interferometry

First proposed by Froehly [8], this technique (sometimes referred to as Fourier transform spectral interferometry or FTSI) determines the spectral dependence of phase through Fourier transforming a spectral interferogram, rather than scanning delay. Because of this the experimental setup is the simplest of any of the linear phase measurement techniques, as shown in figure 3.2.

The interferometer is not balanced in this case; the (nominally static) delay, τ, between arms must be chosen carefully. To yield valid results the measured

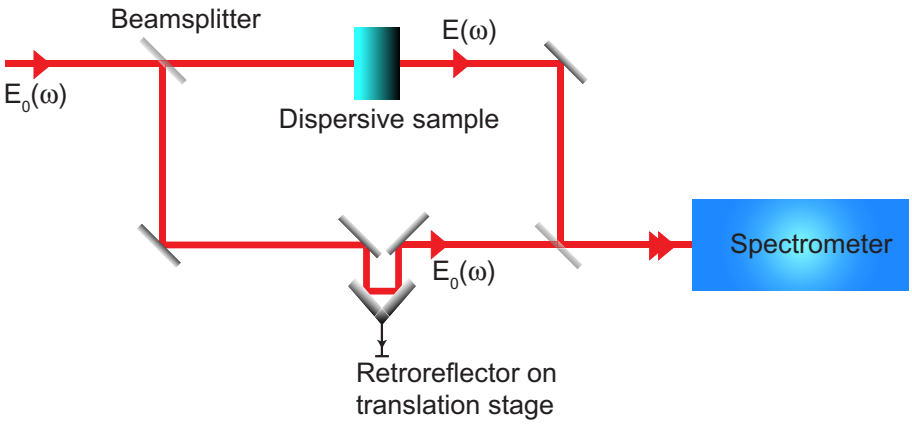

Fig. 3.2: Frequency domain interferometer, as conceived by Froehly et al

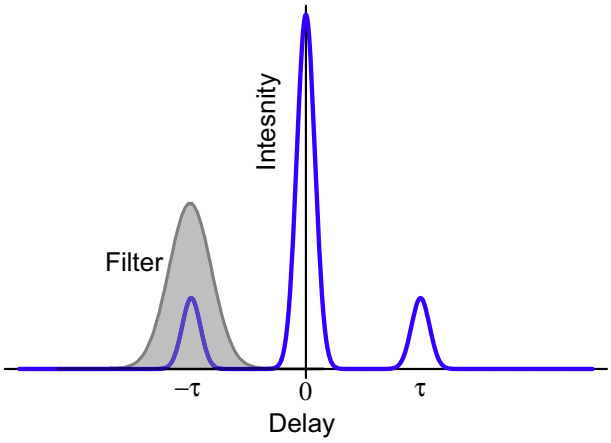

Fig. 3.3: Delay vs intensity trace resulting from Fourier transform of spectral interferogram

interference pattern must both contain sufficient fringes to be able to determine phase and have the fringes spaced sufficiently far apart to be well resolved. The upper limits of delay corresponding to these requisites are wholly dependent on the scanned spectral range and the resolution of the spectral filter respectively. By defining the largest value of delay for which a valid interferogram can still be measured, these two factors determine the maximum spectral resolution of phase for this technique.

The measured interference pattern is Fourier transformed to give intensity

as a function of time, as shown in figure 3.3. Two peaks are symmetrically placed around $t = 0$, centred at $t = \tau$ and $t = -\tau$. These peaks correspond to the terms in equation 3.4. A large peak centred at $t = 0$ represents the sum of the intensity terms (i.e the envelope of the interferogram) and carries no phase information. Filtering out all of the transform other than one of the satellite peaks leaves data corresponding only to intensity and the complex transfer function, $R(t - \tau)$. The central peak can be eliminated from the transformed trace by chopping the beams in both arms of the interferometer at different frequencies and using lock-in detection to measure the signal at the difference frequency. The resultant interferogram then corresponds only to the interference between beams and not to their intensities. Removing the central peak allows a broader filter to be applied to the transformed trace. Having filtered one of the satellite peaks, it is necessary to compensate for the finite delay, τ (remembering that a constant delay corresponds to a slope of constant spectral gradient). This is done by centering the filtered peak at $t = 0$, before the inverse transform is applied to it. The argument of the resultant complex frequency plot gives the phase transfer function of the device. The resolution of the complex frequency plot is proportional to the width of the filter applied before performing the inverse transform. The maximum filter width is proportional to the chosen delay, τ.

The interferogram can also be treated without using Fourier transforms; the position of fringe peaks can be noted and the separation of adjacent peaks equated to a phase of 2π radians[9]. If this is done, phase resolution is again proportional to delay (assuming spectral resolution is sufficient to adequately resolve the interferogram), but is independent of the scanned spectral range. Increasing delay decreases the ratio of signal to noise however and for this reason, Fourier transforms offer greater accuracy at high spectral phase resolution.

This technique relies on the accuracy of the spectral filter and to ensure a smooth trace, measurements are taken at intervals smaller than the resolution of the spectral filter. This 'running average' removes sharp features from the delay trace and so eliminates the corresponding high frequency components that would otherwise be seen in the Fourier transform.

Other than the simplicity of the experiment, there are several appealing features to this method. The calculated phase is given as a continuous function of frequency, as opposed to the discrete points given by methods which do not use Fourier transforms. This is ideal when consideration is given to the fact that the phase transfer function will be derivated once to give group delay and again to yield GDD. No fitting functions have to be applied, which avoids the associated problems of smoothing. When applying a fitting function to discrete data it is common to smooth and so eliminate spurious sharp features caused by noise, but this also lessens the accuracy of the fit.

A significant advantage of FTSI is that all spectral components can be measured concurrently. The spectral components can be measured simultaneously by an optical spectrum analyser or spatially separated and detected by a CCD array. For any significant acquisition time (in comparison to the rate of drift in delay), variations in delay will cause the measured interferogram to lose contrast (as with the compensated Michelson). Despite the unavoidable decrease in signal to noise ratio this causes, data taken this way is still valid. Such measurements are also complicated by the frequency step between pixels changing over the width of a CCD, although compensation can be made for this[10].

3.3.4 Time-domain interferometry

Also referred to as dispersive FTSI, this technique, as its name suggests, determines phase through measurement of a time domain interference pattern. Nagamuma[11] demonstrated that a spectral filter was not needed to resolve dispersion: he proposed and demonstrated a modified version of FTSI in which all frequency components were measured simultaneously.

To provide the necessary interference pattern from which to extract a transfer function, the signal from all wavelengths entering the interferometer is measured as a function of delay. Fourier transforming this gives a complex function in the frequency domain. As with FTSI, three peaks are evident in the transform; a central peak corresponding to the time envelope of the interferogram and two symmetrically placed peaks (at $\Omega = -\omega$ and $\Omega = \omega$, where ω is the

central frequency of the source light) which are mutual complex conjugates. The phase transfer function can be found by taking the argument of either of the satellite peaks.

The conceptual differences between this technique and FTSI are minimal; in experimental setup, delay is scanned instead of wavelength and in data treatment, a single Fourier transform is performed rather than both a transform and its inverse.

In terms of applicability to measuring a given device with a given source, this technique is distinguished from FTSI by the factors that determine the accuracy and resolution of measured transfer functions.

A fundamental property of discrete Fourier transforms is that the range over which a measurement is taken in one domain is directly proportional to the resolution in the transformed domain. Specifically; in frequency domain interferometry, the width of the filtered complex delay function is proportional to the resolution of the transfer function given after performing the inverse transform. The maximum filter width is proportional to τ, the upper limit of which is determined by the resolution of the spectral filter. Hence the limiting factor to resolution in frequency domain interferometry is (indirectly) the resolution of the spectral filter.

In time domain interferometry the equivalent requirement (interferograms taken over a large delay range) is easily met; delay ranges of several thousand wavelengths can be scanned with piezo-transducers and / or automated stages. The interferogram width has no fundamental limit, but each of the large number of fringes must be well resolved, which requires many discrete measurements. Consequently high phase resolution is associated with long measurement times and large data files. As with PLI, accurate measurement of delay is crucial in time domain interferometry (any uncompensated fluctuations in optical path length would invalidate the measurement). Therefore it is essential to use a reference beam passing through the interferometer to provide this. Recording the signal from the reference at each delay further increases the required time and volume of data needed to provide a given phase resolution. Another result of requiring so many data points is that this technique can not realistically be adapted for use with CCD arrays (which have a limited

number of pixels).

A novel variation of time domain interferometry has been demonstrated by Sprik and Kop [7]. Their method is designed exclusively for pulsed sources (all other linear techniques listed in this chapter were designed or can be applied to CW or incoherent light sources). In their experiment pulses pass in to an unbalanced, static Mach Zehnder interferometer, one arm of which contains the device under investigation. The interferometer emits collinearly one pulse which has been altered by transmission through the device and a temporally delayed, unaltered pulse. These two pulses then enter a commercial Fourier transform spectrometer, previously referred to in section 3.3.2. The spectrometer acts as a scanning Michelson interferometer (i.e. a first order autocorrelator) and measures a time domain interferogram. This consists of three peaks, a central peak which is the sum of the autocorrelations of both pulses and two symmetrically positioned peaks which are the correlations of the unaltered pulse with the altered pulse and vice versa. Similar to the Fourier transforms of frequency or time domain interferograms, the central peak contains intensity information only and the two satellite peaks (again complex conjugates of each other) each fully describe the phase transfer function.

The combination of two interferometers distinguishes this technique from other time domain FTSI arrangements, but does not offer any comparative advantages other than ease of use. An issue left unaddressed in this technique is the inevitable fluctuations in delay between pulses; the delay imparted by the interferometer is not referenced. During the course of a single measurement, unnoticed changes in delay could invalidate the interferogram.

3.3.5 Single-shot spectral interferometry

This is not a specific technique, instead it encompasses several examples of spectral interferometry which all share a common feature. Instead of scanning delay, frequency or both and taking a measurement at each scanned position, the entire interference pattern is measured simultaneously with stationary optical components only. To achieve this, light is separated spatially according to its delay, frequency, or both. Refractive and diffractive elements such as prisms

and gratings are routinely used to separate light according to frequency. In the case of frequency domain interferometry such an element combined with a CCD array is enough to provide a single shot measurement from which the transfer function can be extracted.

Meshulach [12] and Kovacs [13] have both demonstrated single shot techniques which yield plots of delay versus frequency. The experimental arrangement for the former is shown in figure 3.4. A retroreflector in the empty arm of the balanced interferometer is used to vertically offset the unaltered beam with respect to the beam which passes through the device under investigation. Both beams are refracted by passage through a prism, spreading the frequency components in parallel horizontal planes. To provide a variable delay a cylindrical lens is used to focus the two beams vertically but leave them unchanged horizontally. The relative delay between the interfering beams is a function of the height at which they are detected. The beams interfere on a two dimensional CCD array, with the horizontal axis corresponding to wavelength and the vertical axis corresponding to delay. The horizontal axis must be calibrated because it is not directly proportional to either frequency or wavelength [10]. The delay axis is scaled directly from the plot; the derivative of a given fringe is equated to delay at the corresponding frequency. Both of these operations can be carried out very quickly, and because acquisition time is short, this technique can yield real-time phase measurements. Because thermal drift occurs over a longer time scale than the acquisition time (seconds and fractions of a second respectively) a reference is not needed to calibrate delay. The delay vs. frequency plots are intuitive and unambiguous, the direction of fringe curvature indicates whether delay is increasing or decreasing with wavelength.

The earlier technique described by Kovacs is very similar to that of Meshulach. Designed to investigate optics for use in femtosecond laser systems, the variable delay is provided by slightly tilting (around a horizontal axis) a mirror in one arm of the interferometer. The higher the point at which light is reflected from the mirror, the further it has to travel to exit the interferometer. This provides delay dependence in the vertical axis of the CCD array. In Kovacs' publication[13], the tilted mirror is the dispersive element under investigation, but if the device were transmissive tilting one of the standard

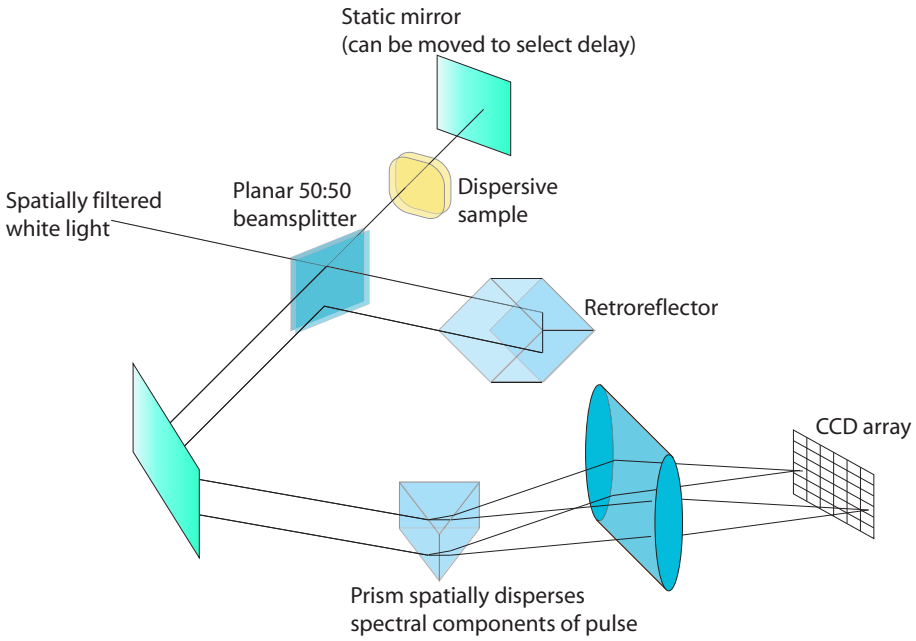

Fig. 3.4: Single shot technique as demonstrated by Meshulach[12] et al

metal coated mirrors would provide the desired variable delay. To horizontally spectrally disperse the interfering beams a grating and lens are placed before the CCD array. The delay axis of the CCD array is again calibrated directly from the delay vs. frequency plot. The frequency axis is calibrated by noting the positions at which light of known wavelengths is incident on the CCD array (spectral lines from a caesium lamp were used by Kovacs et al).

Both techniques quickly (Kovacs states a 200ms acquisition time) and unambiguously yield high resolution phase transfer functions (Kovacs estimates group delay resolution to be 0.2fs). The absence of moving components is a distinct advantage over other techniques; despite scanning delay, no reference is needed to calibrate the delay scale.

The advantages of these techniques are tempered by two significant limitations; spectral resolution and sensitivity are inferior to other linear techniques discussed previously. While Meshulach does not declare a value of resolution, Kovacs states a spectral resolution of $\sim$1nm. This value is at least an order of magnitude poorer than would be expected from an optical spectrum analyser

or scanning monochromator. Resolution is proportional to the number of pixels in the CCD array, but sensitivity is proportional to the intensity incident on each pixel. Because of this, signal to noise ratio limits both resolution and sensitivity. Both of these techniques were demonstrated using "white light" (i.e. broadband, incoherent) sources, which necessitated the use of spatial filters to improve spatial coherence (Kovacs used slits, Meshulach used a single mode fibre). This significantly reduced the intensity of light incident on the CCD arrays, as evidenced by the 250W quartz tungsten halogen lamp used as the light source by Meshulach.

Although not a truly single-shot technique, Tarhan[14] describes an experiment where a series of single-shot time domain interferograms are taken at successive wavelengths. A balanced Mach Zehnder interferometer is used to provide temporally overlapped, but spatially separated beams. Pentaprisms in either arm allow for convenient manipulation of both parallel, horizontally offset beams. A lens focuses the beams on to a 1D CCD array, exactly as in the Meshulach arrangement, which allows a single-shot time domain interferogram to be recorded. As an intermediate between true single-shot techniques and time domain interferometry, this technique has qualities of both. No reference is needed to calibrate and compensate for changes in delay, but spectral resolution is governed by the spectral filter applied to the source before the interferometer. The time required to acquire and process data prevents this from being implemented in real time. Because the source light is both spectrally filtered and spatially dispersed, sensitivity is low.

3.4 Nonlinear characterisation

The signal from a detector relying on one of the lowest order nonlinearities (such as TPA or SHG) will be directly proportional to the square of intensity, (i.e.$|E|^4$). At low intensities, the signal will be lower than for a linear detector due to the relative inefficiency of the nonlinearity (represented by $\chi_2 \simeq 10^{-8}\chi_1$). Only at high intensities (where the electric field due to incident light is comparable to the intra-atomic field binding the electron to the nucleus), will the nonlinear signal be stronger than the linear signal. For this

reason nonlinear measurements are best suited to characterising pulses before they are significantly attenuated or temporally dispersed.

A significant factor arising from the detected signal being proportional to the square (or cube etc, depending on what order of χ the nonlinearity corresponds to) of electric field is the corresponding increase in sensitivity. Small fluctuations (such as noise) in electric field translate to significant variations in detected signal, with the result that nonlinear signals are typically noisier than their linear counterparts.

As can be seen in equation 3.3 if interference effects are neglected (i.e. if the detector has a slow response time) the detected linear signal from two spatially overlapped pulses will be independent of their temporal overlap. If interference effects are included then the detected signal will describe the bandwidth of the pulses, but in either case the linear signal cannot yield the pulse duration.

Conversely for nonlinear detection, measuring intensity as a function of delay will yield information about the degree of temporal overlap; as the overlap between two pulses increases so does peak intensity and hence the detected signal. It is this relationship that allows nonlinearities to be used to resolve the temporal overlap of pulses and in doing so allows the phase and amplitude of an unreferenced pulse to be determined.

3.4.1 Autocorrelation

The simplest and most common way to extract information about an optical pulse, other than measuring its power spectrum, is to perform an autocorrelation measurement, or more accurately a second order autocorrelation. An autocorrelation trace, with prior knowledge or assumption of the pulse shape, allows the temporal width of a pulse to be determined. The associated error can be relatively large (of the order of 10%) but these traces can be taken in real time and require only a relatively simple apparatus as shown in figure 3.5.

The autocorrelation technique uses a Michelson-Morley interferometer to split each incident pulse in to two equal halves and then rejoins them after travelling along separate arms. Both arms have nominally equal length, but one stays fixed while the other oscillates. This results in the two halves of the

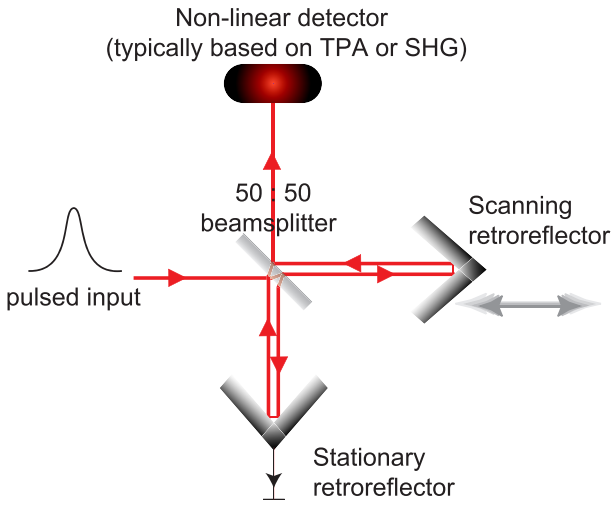

Fig. 3.5: Schematic of a (2nd order) autocorrelator.

original pulse travelling collinearly after meeting at the beamsplitter, but with a time varying degree of overlap. The degree of overlap determines the peak intensity incident on the detector.

The detected signal will be directly proportional to the time integral of the squared intensity.

$$S = \int_{-\infty}^{\infty} |(E(t) + E(t+\tau))^2|^2 dt \tag{3.5}$$

Because of this as the pulses move through a cycle of no overlap through to complete overlap and back again the nonlinear signal traces a flat background value (for $\tau \gg \tau_p$, where τ_p is pulse width),

$$g = 2 \int_{-\infty}^{\infty} |E(t)^2|^2$$

which rises to a maximum (at $\tau = 0$) of

$$g = 16 \int_{-\infty}^{\infty} |E(t)^2|^2$$

and back again.

The width of the trace's peak corresponds to the width of the pulses and

can be calculated exactly, assuming the shape of the pulses is known. The non linear detectors are typically either a photomultiplier tube (PMT) and non linear crystal combination or a two photon absorbing (TPA) material. The PMT, crystal combination offers superior sensitivity, but is relatively bulky, polarisation dependent, more costly and alignment of the crystal is of critical importance. TPA detectors are often chosen because of their ease of use if sensitivity is not a problem. TPA is determined by intensity only and as such polarisation need not be considered.

Whichever type of detector is chosen, the alignment of the overlapping beams is extremely important. The pulse shape is normally assumed from the shape of the spectrum (a Gaussian spectrum corresponds to a Gaussian temporal profile when Fourier transformed, and so this is often the assumed temporal profile). Only if this assumption is true and if the two beams incident on the detector are collinear over the full length of the scanned delay will the calculated pulse width be correct. To aid the alignment of the two beams an ideal ratio of background signal to peak signal should be achieved before measuring the width of the peak. Fortunately the trace can be observed in real time while the beams are being adjusted and this becomes straightforward in practice. The optimal background to peak ratio depends on the response time of detection. From the equations immediately above the optimal peak to background ratio should be 8:1. This is the case for a detection system of sufficiently fast time response. The autocorrelation is called an interferometric autocorrelation because the individual fringes under the pulse envelope can be resolved. Knowing the central wavelength of the pulse allows the spacing of these fringes to provide the time scale of the trace.

If either the detector or connected electronics are too slow then a time averaged, or intensity, autocorrelation results. This has a 3:1 optimal ratio, and exhibits no fine structure. In an intensity autocorrelation the delay scale must be established by scrolling the static delay arm and equating the distance scrolled to the distance the peak of the trace is seen to move.

The interferometric autocorrelation contains phase information, but only qualitatively so. Second order autocorrelations are ambiguous in that they always have symmetrical profiles, which can be realised intuitively by consid-

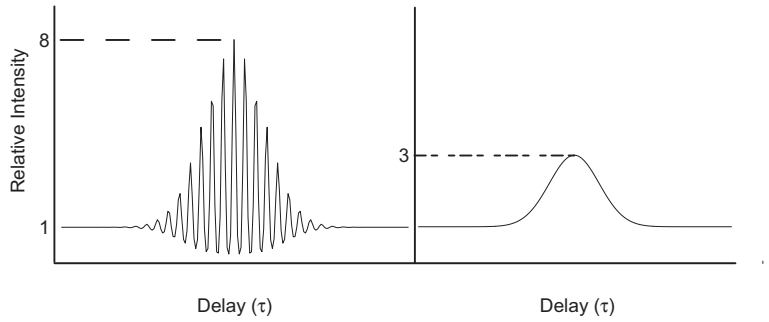

Fig. 3.6: Interferometric (left) and intensity (right) autocorrelations of a Gaussian pulse.

ering that the autocorrelation is constructed by scrolling one pulse through a duplicate pulse. Because the two pulses are identical, it is not possible to state (from the autocorrelation trace alone) which pulse is ahead of the other temporally. The intensity measured for a finite degree of overlap could correspond to the trailing edge of pulse $'A'$ overlapping with the leading edge of pulse $'B'$, or vice versa. From the practical perspective of measuring pulse width this is immaterial, only a measure of the width of the autocorrelation peak is needed (FWHM is the standard measure).

The FWHM, taken in conjunction with an assumed pulse shape can be used to calculate temporal pulse width. Ideally this value would correspond to transform limited pulses, (i.e. the temporal pulse width given by measuring the frequency spectrum of the pulse, arbitrarily applying a constant spectral phase and Fourier transforming the result). If this occurs then the assumed pulse shape is validated, otherwise little can be concluded other than that the pulse is "chirped".

3.4.2 Cross correlation

The cross correlation technique is equivalent to an autocorrelation in which the environments in each arm are different. For example, placing a piece of glass in one arm would alter the pulse in that arm only. The resultant correlation trace would be a function of both the original pulse and the properties of the glass. If the original pulse width were known, the width of the pulse

that passed through the glass could be calculated. Uncertainties in the pulse width found through an autocorrelation would be exacerbated by taking a cross correlation measurement, but there is a trade-off in measurable signal level. In an autocorrelator the pulses in each arm are by definition identical, including their powers. In cases where it is undesirable to use pulses with high peak intensity (due to the operational requirements of the sample under investigation), autocorrelations may not offer sufficient sensitivity. In a cross correlator, this can be circumvented by attenuating power in the sample arm only. The detected signal will then be a function of the intensity in both arms:

$$S = \sqrt{I_{arm1} I_{arm2}}$$

The cross correlation technique benefits from being simple and intuitive, but like autocorrelation can only offer low resolution information about the temporal expansion or contraction of a pulse.

3.4.3 Time of flight measurements

When investigating propagation speeds of pulses through structures, one common technique used is to measure the time of flight[15]. This is similar to a cross correlation, except that only the peak of the signal vs. delay trace is sought. Again a nonlinear detector is used in conjunction with a scanning arm, but only the position of the centre of the peak is noted. The width of the correlation peak and the spacing of fringes (if any are present) are ignored. This is the nonlinear version of the compensated Michelson arrangement, only the absorption mechanism employed by the detector separates these techniques. Due to nonlinear detection the resolution offered by this technique is determined by the width of pulses used. Typically this is inferior to the resolution offered by linear techniques, and as such time of flight measurements are usually performed on devices expected to exhibit significant changes in group velocity over a narrow spectral range (e.g. colloidal photonic crystals[16]).

3.5 Full pulse characterisation

Full pulse characterisation is a term used to specify a range of techniques which are all capable of defining the amplitude and phase of an unknown (and unreferenced) ultrashort pulse. Although it could be argued that the spatial properties of pulses should be considered also, the majority of these techniques do not (although exceptions exist[17]).

As discussed previously all full pulse characterisation techniques require both a linear and nonlinear measurement. The nonlinear optical processes used in detection are second harmonic generation (SHG), sum and difference frequency generation (SFG, DFG), two photon absorption (TPA) and the optical Kerr effect. Of these SHG, SFM and DFM are governed by $\chi_{(2)}$ and therefore offer greater sensitivity to techniques that use them than to those that rely on the weaker $\chi_{(3)}$ based nonlinearities, TPA and the Kerr effect. Because sensitivity is crucial when considering measurements of potential optical network components, techniques which use $\chi_{(2)}$ will be given priority here.

Since the advent of FROG [18] many similar techniques have been developed [19][20]. As well as their shared reliance on both linear and nonlinear processes, all full pulse characterisation techniques require some form of algorithm to iteratively arrive at an amplitude and intensity profile from measured data. While there are many competing algorithms, from an experimentalists perspective it is necessary only to realise that given a good quality (i.e. low noise, free from systematic errors) measurement it will be possible to calculate phase and amplitude correctly.

3.6 FROG

Frequency resolved optical gating (FROG) is the most well established full characterisation technique, whether in its original guise of SHG FROG or one of the more complex variants[21]. The concept of FROG is shared with the majority of other full characterisation techniques; to measure a nonlinearity-dependent intensity as a function of both frequency and delay. The traces produced in this manner are then subjected to a retrieval algorithm, to de-

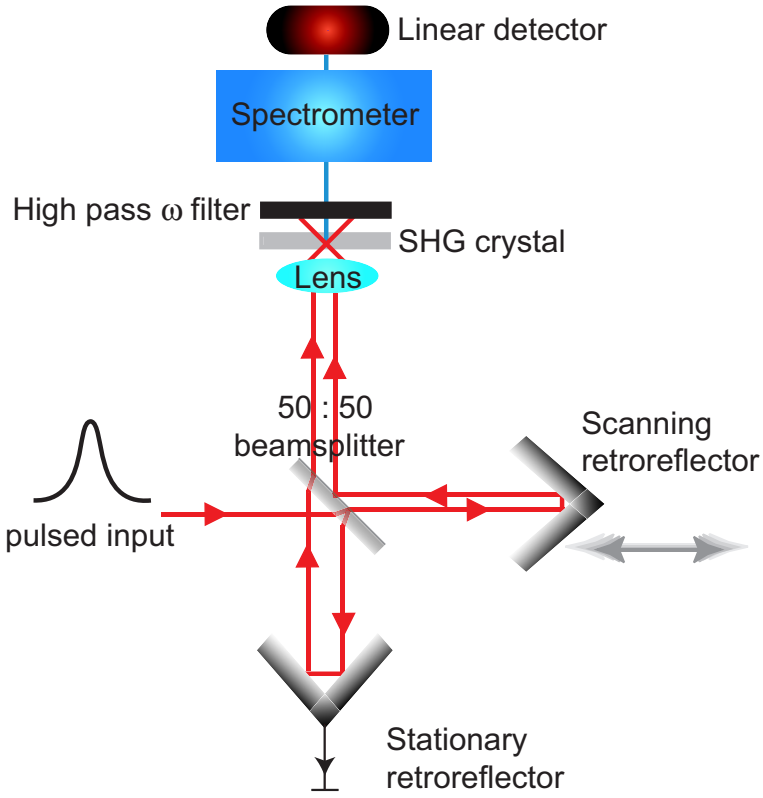

Fig. 3.7: The experimental layout of a SHG FROG

termine the phase and amplitude of electric field (i.e complex $E(\omega)$ or $E(t)$) from the purely real values measured. The algorithm is designed to iteratively determine a (hopefully unique) complex electric field function.

SHG FROG is the most sensitive and widely used of this class of techniques. The other FROG techniques utilise $\chi_{(3)}$ based nonlinearities (most commonly the optical Kerr effect). The experimental setup for SHG FROG reveals that it is only slightly different from a standard autocorrelation.

The addition of a spectrometer before detection spectrally resolves the second harmonic signal and allows a time (i.e. delay) vs frequency plot to be constructed by taking spectra at many discrete delays.

It is important to scan over a larger frequency and delay range than the pulse occupies, i.e. the pulse should be completely contained within the FROG trace. To verify (or reassure the experimentalist of) the validity of the resultant

complex electric field, integrals of the trace with respect to either frequency or delay can be compared to autocorrelations or spectra taken simultaneously. Although this implies large data sets, many FROG experiments are performed quickly, in some cases in real time [22].

The resolution offered by FROG is excellent. This stems from the retrieval algorithm having access to data over the entire frequency and delay range of the pulse. In a strict sense the pulse amplitude will fall asymptotically to zero at $\pm\infty$ delay and $\pm\infty$ frequency only, but the portion of pulse amplitude outside FROG traces is regarded as negligible [23].

The limiting factor to the ubiquity of FROG is sensitivity, SHG FROG is sensitive enough to measure pulse energies of the order of 1pJ (for a 100fs pulse). While this is sensitive this would correspond to an average power of >1mW for the OPO used at St. Andrews ($\tau_p \sim$ 750fs, 80MHz repetition rate). Although not a limiting factor, it should be noted that SHG FROG traces are symmetrical in time, just as autocorrelations are. To determine the direction of any measured chirp requires prior knowledge of the pulse or sample or a separate measurement with a material of known dispersion inserted in to one arm of the setup.

As mentioned above SHG FROG is the most sensitive of FROG techniques and so is the only FROG variant to be considered here. Similarly the single shot variation of SHG FROG, GRENOUILLE [22] is less sensitive by at least three orders of magnitude and likewise will not be considered.

3.7 PICASO

Phase and Intensity from Correlation And Spectrum Only (PICASO)[24], differs from FROG and similar techniques (such as the less sensitive Sonogram[20]) in that it does not attempt to construct a time-frequency plot. Instead, a conventional pulse spectrum and correlation measurement are required for the algorithm to determine the complex electric field of the pulse. The additional information that makes this possible is the presence of a sample of known transfer function in one arm of the cross correlator. In this case, the interference terms are functions of both the transfer function and the phase profile of

the pulse. This serves to remove the symmetry exhibited in autocorrelations and indicate the direction of time in the correlation.

The accuracy (both spectrally and temporally) and noise induced errors of PICASO are comparable but generally slightly poorer than SPIDER or FROG[25], which may be due to less well developed retrieval algorithms. PICASO uses a TPA detector, which offers the comparative advantage over SHG detectors of making no demands on pulse width; when choosing a SHG crystal the material must allow phase matching over the entire pulse bandwidth. For ultrashort pulses centred at awkward wavelengths or with very large bandwidths there may be no suitable material. The associated disadvantage of TPA as a detection mechanism is that it is a $\chi_{(3)}$ based nonlinearity and so is less sensitive than SHG. PICASO could be adapted to use SHG instead of TPA, its only advantage over FROG would then be its simplicity.

3.8 SPIDER

Spectral Phase Interferometry for Direct Electric field Reconstruction (SPIDER), as its name suggest relies on spectral interferometry to determine the phase across a pulse. The SPIDER technique relies on frequency mixing; two temporally delayed identical copies of an unknown pulse interact nonlinearly with a highly chirped (and therefore temporally longer) version of the unknown pulse. The temporally delayed pulses both undergo frequency mixing with the chirped pulse in a nonlinear medium. Because there is a delay between the two identical pulses they interact with different temporal slices of the chirped pulse (which have different average frequencies). After frequency mixing the temporally separated pulses are frequency shifted by different amounts. Performing spectral interferometry now allows a comparison of the phase corresponding to one frequency of the original pulse with comparison of phase at a separate frequency of the original pulse.

The SPIDER technique is complicated experimentally and perhaps less intuitive than other techniques, but offers high sensitivity (being based on SFG, a $\chi_{(2)}$ nonlinearity) and compares well with SHG FROG. The nonlinear interaction occurs collinearly, there are no moving parts and no ambiguities in

measured data.

3.9 Discussion and Conclusions

Each technique described in this chapter has been outlined and its suitability with respect to the measurement of dispersion in challenging samples. The pertinent question is which of these techniques are best suited to the investigation of low transmission, highly dispersive devices using a pulse source.

Full pulse characterisation techniques are designed to analyse unreferenced pulses and as such are suited to high incident peak intensities. Their sole disadvantage,namely comparatively low sensitivity is not a factor when consideration is given to their optimal role; to fully define a pulse which will later serve as a reference for a spectral interferometry technique. Whether the complexity, cost and time required to arrange such a setup are justified depend on the priorities of the experimentalist.

Autocorrelations are ubiquitous for a good reason; while they offer limited information, they are easy (and comparatively inexpensive) to arrange, can deliver data in real time and are compact. As a precursor to a more advanced technique and as everyday tools for diagnosis of a laser system they are extremely valuable, but can never measure dispersion quantitatively.

Spectral interferometry can boast the highest sensitivities and quickly deliver detailed phase information, fully characterising a transfer function. No pulse information can be extracted by SI alone, either it can be used as a complementary technique to full pulse characterisation or pulses can be partly defined by simple measurements such as autocorrelation and spectrum.

REFERENCE LIST

[1] I A Walmsley M Beck and J D Kafka, "Group Delay Measurements of Optical Components Near 800 nm", *IEEE J. Quantum Electron.*, vol. 27, no. 8, pp. 2074–2081, 1991.

[2] F Marabelli L C Andreani L Pavesi M Galli, D Bajoni and G Pucker, "Photonic bands and group-velocity dispersion in $SiOSiO_2$ photonic crystals from white-light interferometry", *Phys. Rev. B*, vol. 69, no. 115107, 2004.

[3] T J Yu T Imran, K-H Hong and C H Nam, "Measurement of the group-delay dispersion of femtosecond optics using white-light interferometry", *Rev. Sci. Instr.*, vol. 75, no. 7, pp. 2266–2270, 2004.

[4] M Joffre C Dorrer, "Characterization of the spectral phase of ultrashort light pulses", *C. R. Acad. Sci. Paris*, vol. 2, no. IV, pp. 1415–1426, 2001.

[5] L Waxer I Walmsley and C Dorrer, "The role of dispersion in ultrafast optics", *Rev. Sci. Instr.*, vol. 72, no. 1, pp. 1–29, 2001.

[6] V Binjrajka, D E Leaird C-C Chang, A W R Emanuel, and A M Weiner, "Pulse shaping of incoherent light by use of a liquid-crystal modulator array", *Opt. Lett.*, vol. 21, no. 21, pp. 1756–1758, 1996.

[7] H J Kop and R Sprik, "Phase-sensitive interferometry with ultrashort optical pulses", *Rev. Sci. Instrum.*, vol. 66, no. 12, pp. 5459–5463, 1995.

[8] A Lacourt C Froehly and J-C Vienot, "Time impulse response and time frequency response of optical pupils. Experimental confirmations and applications", *J. Opt.(Paris)*, vol. 4, pp. 183–196, 1973.

[9] G Cheriaux L Lepetit and M Joffre, "Linear techniques of phase measurement by femtosecond spectral interferometry for applications in spectroscopy", *J. Opt. Soc. Am. B*, vol. 12, no. 12, pp. 2467–2474, 1995.

[10] J P Likforman M Joffre C Dorrer, N Belabas, "Spectral resolution and sampling issues in Fourier-transform spectral interferometry", *J. Opt. Soc. Am. B*, vol. 17, pp. 1795–1802, 2000.

[11] K Mogi K Naganuma and H Yamada, "General method for ultrashort light pulse chirp measurement", *IEEE J. Quantum Electron.*, vol. 25, no. 6, pp. 1225–1233, 1989.

[12] D Yelin D Meshulach and Y Silberberg, "White Light Dispersion Measurements by One and Two-Dimensional Spectral Interference", *IEEE J. Quantum Electron.*, vol. 33, no. 11, pp. 1969–1974, 1997.

[13] Z Bor A P Kovacs, K Osvay and R Szipocs, "Group-delay measurement on laser mirrors by spectrally resolved white-light interferometry", *Opt. Lett.*, vol. 20, no. 7, pp. 788–790, 1995.

[14] M P Zinkin I I Tarhan and G H Watson, "Interferometric technique for the measurement of photonic band structure in colloidal crystals", *Opt. Lett.*, vol. 20, no. 14, pp. 1571–1573, 1995.

[15] N M B Perney J J Baumberg M C Netti M E Zoorob M D B Charlton C E Finlayson, F Cattaneo and G J Parker, "Slow light and chromatic temporal dispersion in photonic crystal waveguides using femtosecond time of flight", *Phys. Rev. E*, vol. 73, pp. 016619, 2006.

[16] G Klein B Honerlage Y A Vlasov, S Petit and C Hirlimann, "Femtosecond measurements of the time of flight of photons in a three-dimensional photonic crystal", *Phys. Rev. E*, vol. 60, no. 1, pp. 1030–1035, 1999.

[17] D H Sutter T Rupp L Gallmann G Steinmeyer D H Sutter T Rupp L Gallmann, G Steinmeyer and U Keller, "Spatially resolved amplitude and phase characterization of femtosecond optical pulses", *Opt. Lett.*, vol. 26, no. 2, pp. 96–98, 2001.

[18] D N Fittinghoff J N Sweetser M A Krumbugel B A Richman R Trebino, K W DeLong and D J Kane, "Measuring ultrashort laser pulses in the time-frequency domain using frequency-resolved optical gating", *Rev. Sci. Instrum.*, vol. 68, no. 9, pp. 3277–3295, 1997.

[19] C Iaconis and I A Walmsley, "Self-Referencing Spectral Interferometry for Measuring Ultrashort Optical Pulses", *IEEE J. Quantum Electron.*, vol. 35, no. 4, pp. 501–509, 1999.

[20] I G Cormack, W Sibbett, and D T Reid, "Rapid measurement of ultrashort-pulse amplitude and phase from a two-photon absorption sonogram trace", *J. Opt. Soc. Am. B*, vol. 18, no. 9, pp. 1377–1382, 2001.

[21] 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*, vol. 10, no. 5, pp. 1101–1111, 1993.

[22] P OShea S Akturk, M Kimmel and R Trebino, "Highly simplified ultrashort pulse measurement", *Opt. Lett.*, vol. 26, pp. 932–934, 2001.

[23] 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.*, vol. 18, no. 10, pp. 823, 1993.

[24] J Jasapara J W Nicholson, W Rudolph, F G Omenetto, and A J Taylor, "Full-field characterization of femtosecond pulses by spectrum and cross-correlation measurements", *Opt. Lett.*, vol. 24, no. 23, pp. 1774–1776, 1999.

[25] J W Nicholson and W Rudolph, "Noise sensitivity and accuracy of femtosecond pulse retrieval by phase and intensity from correlation and spectrum only (PICASO)", *J. Opt. Soc. Am. B*, vol. 19, no. 2, pp. 330–339, 2002.

4. EVOLUTION OF DISPERSION MEASUREMENT

4.1 Introduction

The initial attempts to meet the challenge of quantifying the dispersive properties of available samples consisted of reviewing current measurement techniques and applying those that could be employed quickly and would give useful preliminary data. The implications of this data and the limitations of the first measurements then led to the use of more sophisticated measurement techniques. Iteratively the quality of data improved as refinements were made to the investigative methods until a final, accurate and reliable technique was found. This chapter details this process and explains why each technique was chosen and states their comparative qualities with respect to the photonic crystal waveguide samples under investigation. Concurrent with this work, efforts were made by the photonic crystal group in the University of St.Andrews to otherwise characterise the samples they fabricated and to improve their properties, most notably their transmission.

4.2 Intensity characterisation

The first sample available for investigation was an array of W1 photonic crystal waveguides, fabricated in GaAs. These samples consisted of arrays of ridge waveguides leading to, and from, variable lengths of photonic crystal sections. The arrays were arranged according to their lattice constants, i.e. the position of their bandgaps. Measurements were carried out by the photonic crystal group using a tunable CW source over the range of available OPO emission (namely 1260nm to 1370nm). The CW transmission spectra indicated that the samples were highly transmissive except in regions close to or within the

photonic bandgap. The CW spectral transmission experiments were followed by autocorrelations of pulses having passed through the sample. The autocorrelator built previously to analyse the OPO output pulses showed the pulses to be 700fs long and transform limited (assuming a Gaussian pulse shape). It was expected that this could be compared to the pulses exiting the PCs to determine whether an appreciable chirp had been applied to the pulses. Ideally little or no chirp should be noticed when the pulse was centred at wavelengths far from the bandedge, and noticeable chirp would be seen otherwise.

Before the autocorrelator could be constructed, it was necessary to efficiently couple light in to and out of the sample. A waveguide manipulator stage was used to hold and precisely position two lenses on either side of the sample. Both lenses and the sample could be translated along three orthogonal axes, and the sample could be rotated in the plane of the beam. This allowed fine manipulation of the train of pulses emitted from the OPO to maximise the light coupled in to and out of the waveguides. Fibres and fibre couplers were available but free space coupling was preferred because this allowed more of the OPO output power to be preserved and eliminated any dispersion that may have resulted from transmission through fibre.

The divergence of the beam from the OPO was large enough to give a beam cross-section significantly broader than the coupling lenses however. To rectify this a small telescope, consisting of two antireflection coated lenses was added after the OPO, which collimated and reduced the diameter of the beam. Similarly coated x40 lenses were used in the waveguide manipulator stage. These offered a large numerical aperture (0.55) but also focussed the pulse train to a spot size of $\sim 3\mu$m, approximately half the input waveguide diameter. An infra red camera was used to position the focussing lens and input waveguide correctly. Unfortunately no light was seen exiting the waveguide. After repeated attempts to realign the lenses and sample, light was detected on a biased Ge photodiode. Approximately 2mW of pulsed light was incident on the input waveguide, and $<10\mu$W was detected by the photodiode. Surprisingly a weak red spot was sometimes seen to be scattered from the mouth of the waveguide. This was too weak and diffuse to be detected and spectrally resolved, but using a series of glass filters, it was established that this was the second harmonic of

the OPO output wavelength. This was evidence that nonlinearities were being strongly excited in the GaAs. Because CW measurements showed far higher transmission than with pulses it was concluded that nonlinear absorption was responsible for the low pulse transmission. Two photon absorption is usually defined by a TPA coefficient for a given material. This coefficient, β relates absorption to light intensity.

$$\frac{dI}{dz} = \alpha I + \beta I^2 + ... \tag{4.1}$$

Here I is light intensity, z is the distance travelled through the absorbing medium and α is the linear absorption coefficient. β is a function of the ratio of bandgap to photon energy (TPA is highest when the photon energy is $\sim$70% of the bandgap).

To verify that TPA was the main cause of the exceptionally low transmission, measurements of input power vs. output power were taken. As can be seen from equation 4.1 transmission should decrease as a function of intensity if TPA is present and this was seen to occur. This result should not have been surprising, considering that the detectors used in the TPA autocorrelators (Hamamatsu G1115 and G1735 GaAsP photodiodes, as used previously by Reid et al[1]) were only marginally different in composition from the GaAs in which the photonic crystal waveguides were fabricated. It should not be overlooked that the effect of TPA is exacerbated when using narrow waveguides (such as the W1 waveguides). The behaviour of photonic crystal waveguides is also dominated by their ability to slow light through multiple reflections, which increases absorption and out of plane scattering concomitantly.

4.2.1 Eliminating Two Photon Absorption

Having established TPA as the cause of the unacceptably low transmission efforts were made to rectify this. Their are no material requirements for photonic crystals, although a large refractive index (i.e. dielectric constant) is preferable. Consequently moving from GaAs to another semiconductor with a lower β value would alleviate the problem of TPA. Silicon has been attributed with a low value of $\beta(0.45\text{cmGW}^{-1})$[2], and was the first material considered

as a replacement for GaAs. Unfortunately the first generation of samples fabricated in silicon exhibited similarly weak transmission, although it was later confirmed that this was due partly to difficulties encountered during fabrication. Also, although transmission was higher, TPA was still present and its effect was significant.

The solution to the problem of low transmission was to fabricate the sample in AlGaAs and to use pulses with central wavelength of $\sim$1.55μm. The PC waveguides were also broadened to W3, i.e. three rows of missing holes, in an attempt to improve coupling efficiency. By using a high proportion of aluminium in the AlGaAs the bandgap was broadened, while still keeping the bandgap direct. The heterostructure consisted of a 250nm thick $Al_{0.25}Ga_{0.75}As$ cladding guiding layer, a 488nm thick $Al_{0.2}Ga_{0.8}As$ guiding layer, a 55nm thick $Al_{0.3}Ga_{0.7}As$ cladding layer and a 2μm thick $Al_{0.3}Ga_{0.7}As$ layer to provide lateral confinement. Several other groups have used this material system to avoid TPA, namely[3]. AlGaAs is very similar to GaAs; other than possessing a wider bandgap its only notable difference is that fabrication is moderately more time consuming and awkward. With respect to integrating a photonic crystal waveguide with one or more active optical components, AlGaAs is an entirely suitable choice of material.

The material change and the move to operate around 1.55μm meant that the photon energy ($\sim$0.8eV) was now less than half of the bandgap(the narrowest bandgap was $\sim$1.711eV[4] in the guiding layer), removing TPA as a loss mechanism. Higher order absorption (such as three photon absorption) was still possible, but the coefficients of these were so small as to be considered negligible.

Comparing the transmission of one of the W3 AlGaAs waveguides at 1265nm (well inside the TPA regime) and at 1570nm highlighted the strength of TPA in this case. The graphs shown in figure 4.1 differ only in central pulse wavelength.

In accordance with Tsang et al[5], a graph of reciprocal transmission vs. peak input power was plotted for the λ=1265nm graph. The constant gradient in figure 4.2 allowed the TPA coefficient β to be quantified.

The curvature of the left hand graph is indicative of TPA and this is quan-

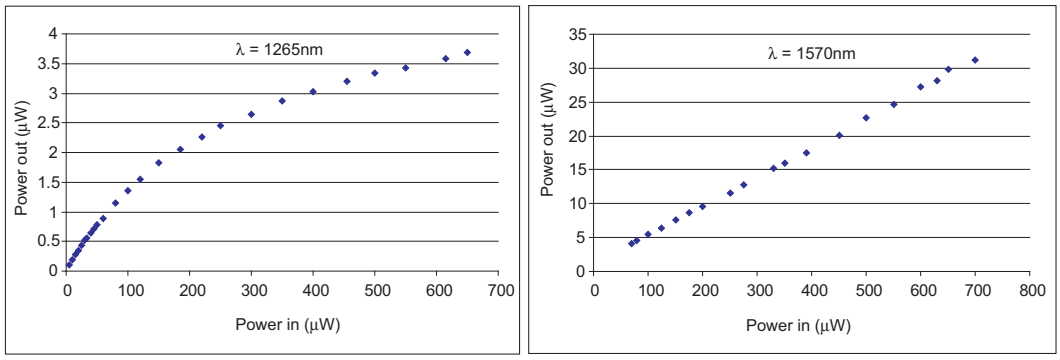

Fig. 4.1: Transmission through a waveguide at 1265 and at 1570nm

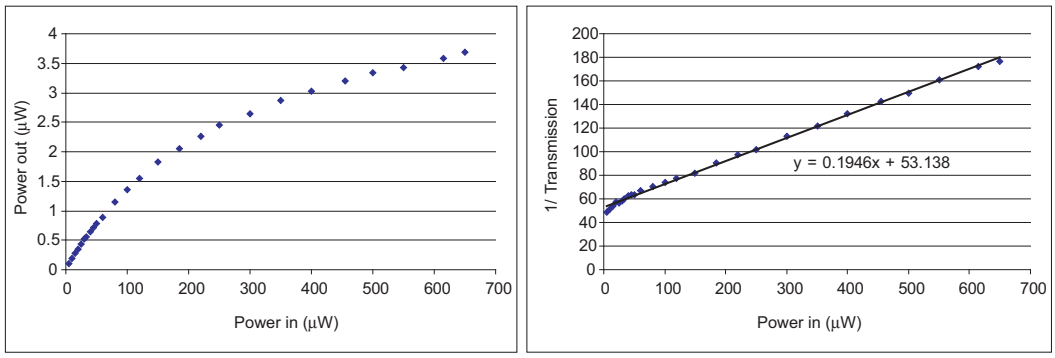

Fig. 4.2: Graphs of Transmission vs. Intensity and 1/T vs. Intensity taken in AlGaAs with pulses centred at 1265nm

tified from the gradient of the right hand graph, indicating a value of between 41 cm/GW and 128 cm/GW for β. The large range of possible β values is due to having to estimate the effective area of the waveguide mode (which is made difficult by the presence of long tapered ridge waveguides).

After adopting AlGaAs and shifting to longer wavelength pulses, the transmission increased markedly, with most waveguides having transmissions of 10 to 15 % (normalised against light passing through the coupling lenses without a sample in place).

4.2.2 Autocorrelation and Cross Correlation

With sufficiently highly transmissive samples a TPA autocorrelator was built to attempt to quantify pulse width after transmission through the sample.

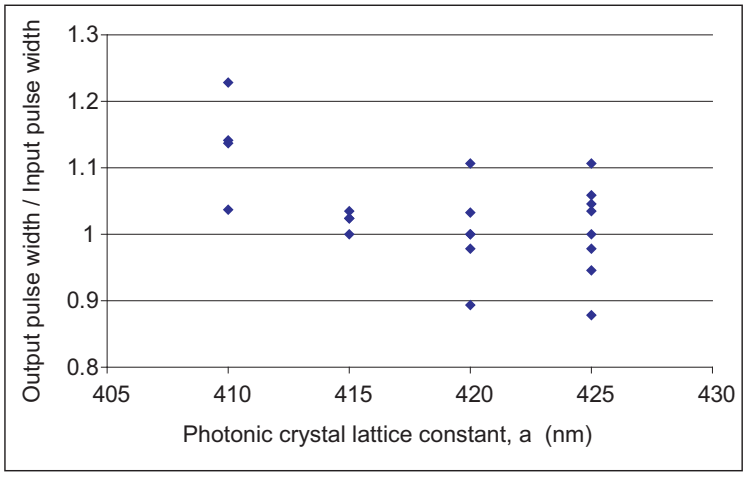

Fig. 4.3: Graph of measured pulse width/original pulse width versus a, the lattice constant of the photonic crystal

The above results were all taken using pulses centred around 1505nm, while the position of the bandedge was varied by investigating waveguides with different values of lattice constant, a. The results implied that the pulses were indeed being broadened by variable amounts depending on their proximity to the bandedge, but the associated error of the technique meant that it became difficult to quantify by how much. Essentially noise and low resolution prevented a meaningful measurement to be taken.

Considering that the limiting factor in the autocorrelation measurements was signal to noise ratio, cross correlation was an easy to implement potential improvement. By superimposing the pulses passing through the sample with unperturbed pulses the resultant detected signal intensity is higher than in an autocorrelation.

$$S_{detected} = \sqrt{I_{arm1}}\sqrt{I_{arm2}} \tag{4.2}$$

Another advantage of the cross correlation was that it was an interme-

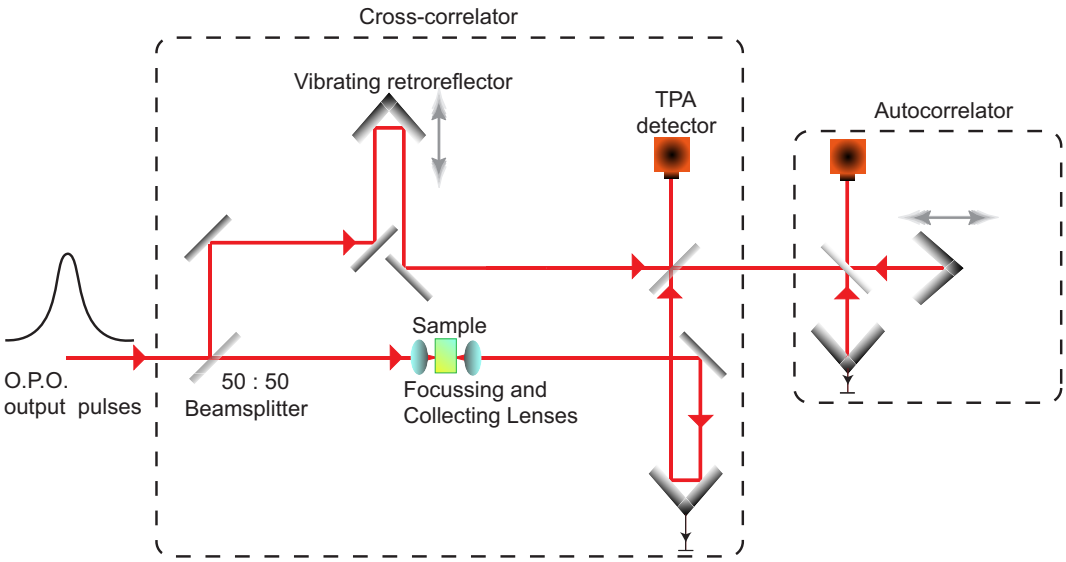

Fig. 4.4: Schematic of cross correlator

diate step towards more complex measurement techniques involving spectral interferometry. It has been shown that full pulse characterisation (PICASO) can be carried out with only a cross correlation and a measurement of power spectrum[6].

By blocking the sample arm of the interferometer, the autocorrelator measured the width of pulses exiting the OPO. A subsequent cross correlation would give a "pulse width" resulting from the correlation of these pulses with pulses of unknown duration (the pulses having been transmitted through the sample). An accurate autocorrelation and corresponding cross correlation were needed to be able to calculate the width of pulses exiting the sample. As this might imply, the effect of compounding uncertainties by needing to calculate pulse width counteracted the increased signal to noise ratio given by the cross correlator. Both autocorrelation and cross correlation gave values of pulse width with too great an uncertainty to provide a consistent and accurate measure of the behaviour of the sample.

This indicated that the variation in group delay across the bandwidth of the pulse was small in comparison to the average group delay of the pulse. In turn this implied that time of flight measurements were likely to suffer from similar

uncertainties, and so non-linear detection based techniques were abandoned.

4.3 Fourier transform spectral interferometry

Nonlinear detection was insufficiently sensitive to provide a useful measure of pulse broadening and so techniques utilising linear detection were then considered. The alternatives were to attempt to set up a full pulse characterisation technique, and a complementary spectral interferometry technique (such as SPIDER etc[7]) or to abandon characterising the pulse. Many different methods of characterising the dispersion induced by reflective samples through spectral interferometry have been established, see for example the references listed in a review article by Dorrer and Joffre[8]. These have been used previously with CW or incoherent sources, but can be applied to pulsed sources also. The advantage of using a pulsed source in the investigation of devices designed to manipulate pulses is that no speculation of the effect of nonlinearities need occur. The information given using a pulsed source describes all of the real effects (and any interaction between them) under the conditions for which the devices were designed. Currently nonlinear effects are typically estimated (in FDTD based models[9] for example) or ignored (in the case of CW characterisation[10]).

Whilst it could be argued that time of flight measurements are an exception to this, their inherently low time resolution (see for example[11], which boasts of $\sim$250fs resolution) and reliance on high powered pulses (in the μJ regime) explains why the information they yield is of limited usefulness and hence why these measurements are not commonplace.

Full pulse characterisation offered 'before' and 'after' time vs. frequency plots of pulses passing through the sample under investigation. The induced changes to the pulses' frequency distribution (i.e. the phase transfer function of the sample) could be deduced from this. Full pulse characterisation (assuming any fitting algorithms work as expected) offered a complete picture of the pulses and the changes effected to them, but required a far more complicated setup (remembering that a linear and nonlinear measurement are needed) and was likely to be slower than the alternative. Much of the information given by

full pulse characterisation could be regarded as redundant; the spectral phase profile of the pulse was unlikely to effect the phase transfer function of the sample. Only the optical Kerr effect (and its resultant shifting of the photonic bandgap) should alter the phase transfer function, and this effect is dependent on intensity (i.e. pulse energy and duration) rather than the relative phase of the pulses' frequency components. Using solely a SI technique would give no pulse information, but autocorrelations and power measurements could quickly and easily verify transform limited (or near transform limited) pulses were being produced and establish the energy per pulse. The facility of full pulse characterisation with a complementary SI technique to quantify temporal shortening of pulses by dispersion compensation was the only distinction between the two alternatives. Because of this, the relative simplicity to set up and the faster data acquisition times offered, a fully linear technique was employed. Fourier transform spectral interferometry (FTSI) was chosen as an appropriately quick and accurate technique requiring a simple setup with few components requiring critical alignment.

An appealing aspect of FTSI was that no moving parts were needed (with the exception of the gratings in the monochromator used as a spectral filter). An uncomplicated example of a FTSI setup was constructed as shown in figure 4.5:

An automated loop, written and controlled by Agilent HPVEE, accepted values for initial and final wavelength, then moved the monochromator filter in steps of equal frequency through the specified range. At each of these steps a sufficiently long pause allowed the lock-in amplifier to integrate over several ($\geq$5) time constants (typically either 100ms or 300ms).

This technique extracts the phase transfer function of the sample from the spectrum of the temporally offset, collinear pulses entering the monochromator.

$$\begin{aligned} S(\omega) &= |E_0(\omega) + E(\omega)e^{i\omega\tau}|^2 \\ &= |E_0(\omega)|^2 + |E(\omega)|^2 + E(\omega)E_0^*(\omega)e^{i\omega\tau} + E^*(\omega)E_0(\omega)e^{-i\omega\tau} \\ &= |E_0(\omega)|^2 + |E(\omega)|^2 + 2|E_0(\omega)E(\omega)|cos(\Delta\phi(\omega) + \omega\tau) \end{aligned} \tag{4.3}$$

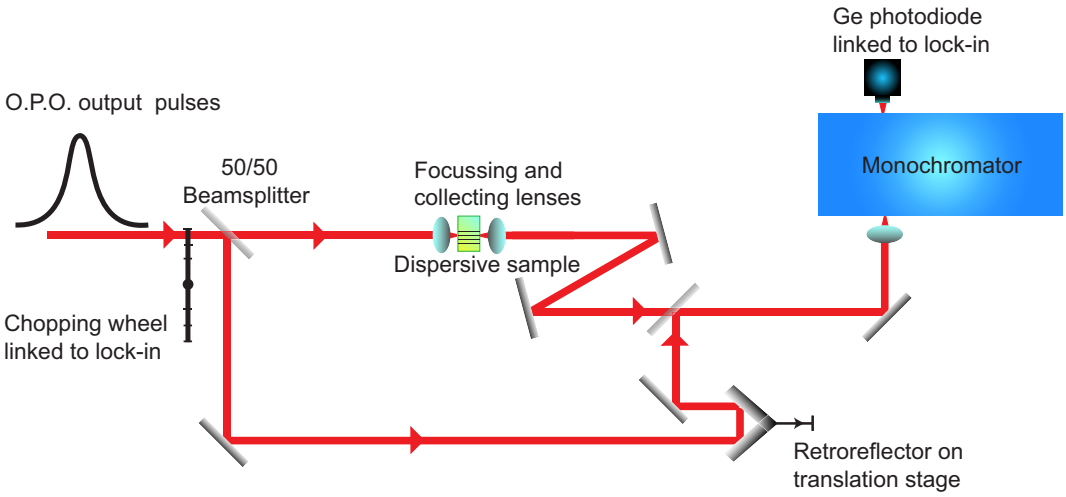

Fig. 4.5: Schematic of the setup used to perform FTSI

The x-axis of the spectrum was converted from wavelength to frequency before being Fourier transformed. The spacing between adjacent points in the transform was uniform, which was required to preserve the integrity of the measurement[12]. Non-uniformly spaced points would skew the scale of the resultant phase transfer function. The transform contains three distinct features.

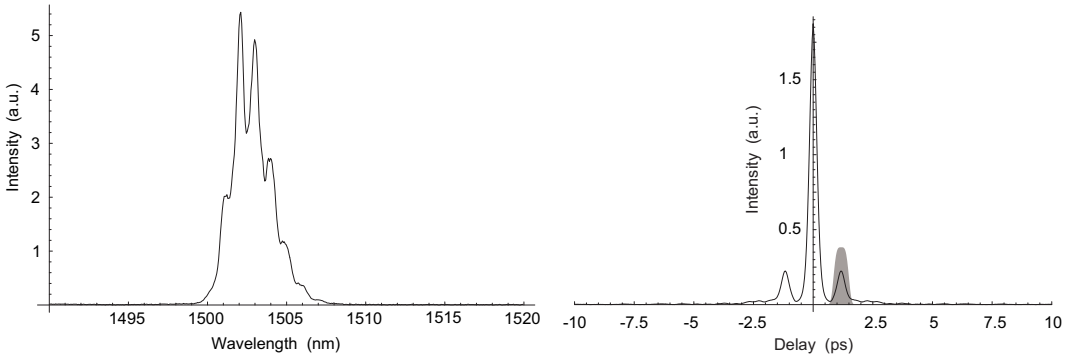

Fig. 4.6: Spectrum and resulting Fourier transform of interfering recombined pulses exiting interferometer. The shadowed area represents the filter applied to the transform.

The large central peak corresponds to the sum of the first-order autocor-

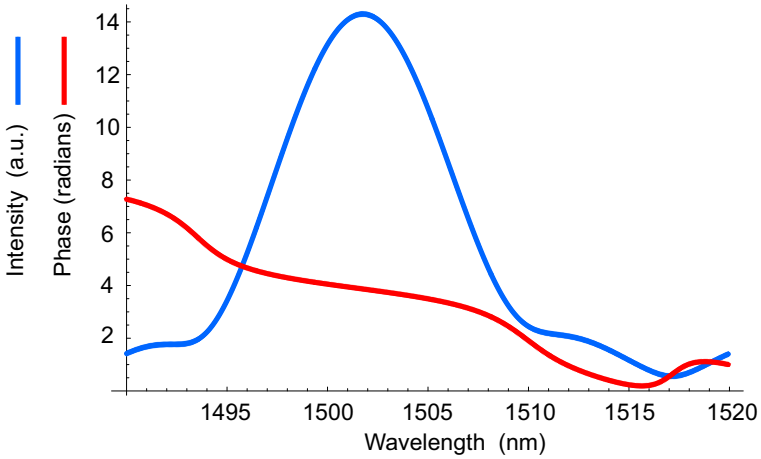

Fig. 4.7: Phase transform function derived from Fourier transformed spectrum. Phase values are only valid for non-negligible intensity.

relation functions of the pulses from both arms of the interferometer. This describes the envelope of the power spectrum (i.e. the spectrum that would be given if the pulses were measured without any interference).

The satellite peaks are mutual complex conjugates symmetrically spaced around zero delay and each contains the information required to fully describe the spectrally dependent phase that this experiment is designed to yield. Placing a filter over one of these satellite peaks and suppressing the remainder of the trace, then shifting the filtered peak to zero frequency and performing the inverse transform gives a complex function. The argument of this function is the phase transfer function of the sample, as first demonstrated by Froehly[13]. It should be remembered that for a finite delay a straight line with gradient proportional to group delay will be added to the phase transfer function. Upon derivation this translates to a spectrally constant addition to the group delay graph. Examples of the data resulting from these operations is shown in figure 4.7:

The variable delay stage was used to set the OPL in the arms of the interferometer. Two competing factors dictated the working range of OPL offset (the difference in OPL between the two arms). The first of these factors was that in order for fringes to be seen in the recorded spectrum, the pulses in

each arm had to be temporally offset when they coincided at the recombining beamsplitter. In the case of the two pulses being maximally overlapped the resulting spectrum would correspond to the spectral envelope of the pulses, with no interference fringes present. If the temporal separation of recombining pulses were increased, fringes would appear in the spectrum, with the number of fringes being proportional to the temporal separation. The second factor is the resolution of the interference fringes; with decreasing temporal overlap the pulses will interfere less strongly with each other, decreasing the magnitude of the interference fringes. More significantly decreasing temporal overlap requires a corresponding increase in spectral resolution, which is set here by the monochromator.

Consequently at too large an OPL offset, no fringes can be clearly resolved (they have similar magnitudes to noise) but at too small an OPL offset not enough fringes will be present to enable a phase to be extracted from the spectrum. This is clearly seen in the transformed spectrum, where the x-axis corresponds to time delay. Reducing the time delay (i.e temporal separation) between pulses moves the satellite peaks closer to zero and when they begin to overlap with the large peak due to the spectrum envelope they can no longer be filtered from the rest of the trace.

Repeated measurements were made with this setup, but it was concluded that the two limits imposed on OPL offset overlapped. This meant no value of offset allowed meaningful values of phase to be measured, because either the spectral resolution of the monochromator was insufficient or the sidebands were inseparable from the pulse envelope peak, for any given OPL offset.

4.3.1 Double-chopped FTSI

Two factors limit the sensitivity of the FTSI described above, namely temporal pulse overlap and sideband resolution in the Fourier transform. The first can only be compensated for with improved equipment and alignment, and only to a limited extent. The second can be improved by altering the experimental setup slightly. By chopping both arms of the interferometer at different frequencies and measuring at the difference of those frequencies, the

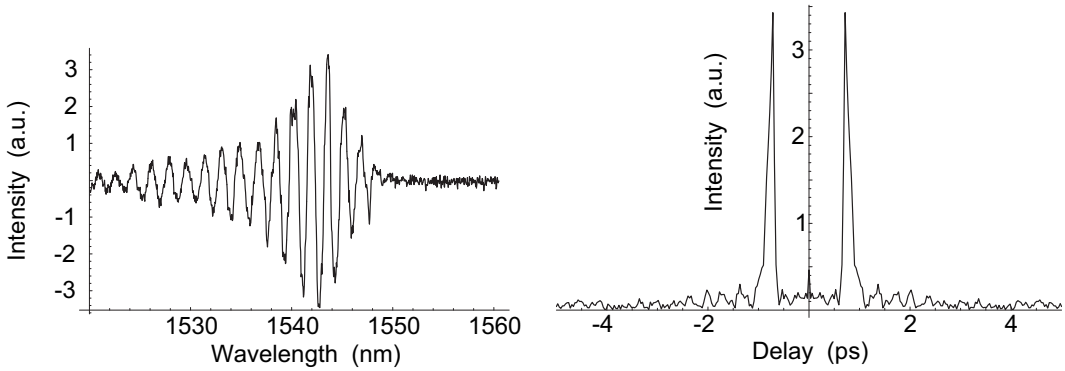

Fig. 4.8: Spectrum and Fourier transform given by double chopped FTSI.

lock-in amplifier would only record a signal when pulses from each arm of the interferometer were both incident (and interfering with each other) at the detector.

Because only the signal corresponding to the interference between the pulses in each arm was recorded, no information about the envelope of the pulse spectrum was present in the measured spectrum. In the Fourier transform this corresponds to the absence of a central peak around zero delay. The only features present would be the two complex conjugate peaks, placed symmetrically about zero delay.

This in turn allowed greater temporal pulse overlap, because the lower limit to OPL offset was now governed by proximity of the satellite peaks to the ends of the transform. Fourier transforms are symmetric and wrap around, so for a small enough delay components of each satellite peak will "jump" to the opposite end of the transform and superimpose upon part of the other satellite peak. The satellite peak to be filtered must be clearly distinguishable and any overlap prevents the extraction of correct phase. Easing the restriction of viable values of offset allowed this experiment to produce traces as shown in figure 4.9.

Unfortunately repeating this measurement immediately afterwards gave a different result. No two scans taken under the 'same conditions' gave the same result; some factor was vitiating the reliability and reproducibility of the measurements. This factor was a thermally induced drift in OPL offset. Even

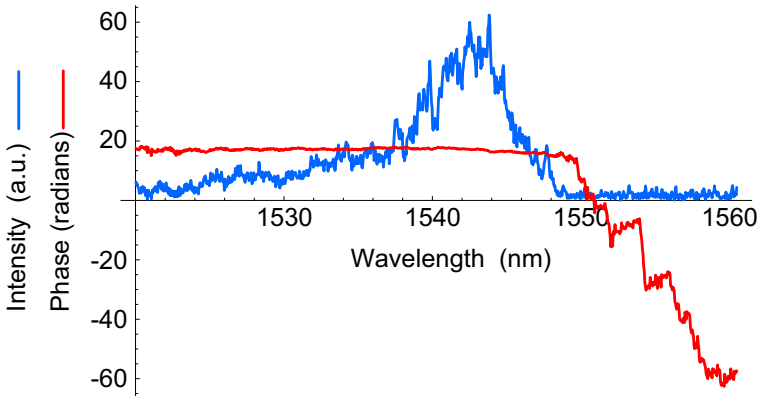

Fig. 4.9: Phase transform function derived from Fourier transformed spectrum. Phase values are only valid for non-negligible intensity.

over the course of a single measurement (∼1minute) the OPL of each arm varied by significant fractions of a wavelength. This is proportionally a tiny movement (hundreds of nanometres in comparison to the OPL of each arm ∼80cm), but invalidates the experimentally derived phase transfer functions. Attempts to correct for this included bolting all optical components of the interferometer to a breadboard, minimising vibrations from electronic devices in the lab and shortening the length of the arms in the interferometer. Each of these should have reduced the jitter in OPL offset, but even together the error persisted and was significant and persistent enough to render this setup unusable.

4.4 Compensation for drift in optical path length offset

To completely remove the effect of OPL jitter, rather than merely reduce it several options were available. Perhaps the most obvious option was to measure the spectra nearly instantaneously. The value of group delay (i.e. OPL offset) added to the phase transfer function was of no importance in investigating the dispersion properties of a sample. The requirement for retrieving a valid phase transfer function was that the OPL offset did not change during the measurement of interference spectra. Several techniques used for full pulse characterisation (such as GRENOUILLE[14]) and spectral interferometry[15][16] do

not use moveable gratings or other spectral filters that are scanned actively, but instead spectrally disperse interfering beams so that their position can be related to their wavelength. A prism made of a highly refractive material is commonly used to separate spectral components of interfering pulses along an axis.

In the GRENOUILLE arrangement a cylindrical lens, aligned perpendicularly to the prism gives a variable delay in the axis perpendicular to this. The light is then incident on a CCD array, and this allows the interference pattern at all wavelengths to be measured simultaneously, giving a 3D plot of intensity vs. wavelength and delay. No measure of delay is needed in this case but the inclusion of a static spectrally dispersive element, such as a grating or prism would allow the simultaneous measurement of all wavelengths within the pulse spectrum.

The advantages to this type of spectral measurement are that any change in OPL offset would effect all wavelengths equally and that with no moving parts only one measurement would be required per phase transfer function (as opposed to the many discrete measurements that occur when using a monochromator).

The associated disadvantages include that because the light intensity on any given pixel of the CCD array is low (given that one axis of the array is separated by delay and the other separated by wavelength), a longer acquisition time is needed than with a single detector. CCD arrays are more sensitive to thermal noise and dark current than standard semiconductor detectors and so to counter such effects an average of several measurements may need to be taken. CCD arrays can have high resolution (i.e. a large number of pixels), but with fringes being spaced only one wavelength apart attempting to measure an entire time domain interference pattern is not realistic.

These disadvantages mean that although the effect of drift in OPL offset should be removed, it will still influence the measured phase transfer functions because a single measurement requires a finite acquisition time. As the OPL offset increases and decreases arbitrarily through drift during the required integration time of the CCD array, the resultant measurement will be compiled from the weighted sum of interference spectra at different delays, lessening

fringe contrast. To resolve this the CCD array could be allowed to integrate for longer, averaging out the effect of OPL jitter. This returns the problem to its original position: either resolution and accuracy or acquisition time must be compromised. This is not specific to CCD arrays, an alternative would be to take the spectral measurement using an optical spectrum analyser and the same conclusion results. Due to limited intensity, either resolution or speed of measurement must be compromised to some extent. Indeed this applies to the FTSI setup discussed above; with sufficient acquisition time (which may well correspond to taking several hundred or thousand spectra per measurement), the effect of OPL offset drift (as well as many other sources of error) could be factored out.

A second approach to eliminating the OPL drift error would be to add an external reference to the setup, as is present in phase locked interferometry[17]. This reference could then be used to quantify and remove the effect of drift in OPL.

4.5 Time domain interferometry without Fourier transforms

Several significant changes were made to the FTSI setup above to arrive at a technique that satisfied the criteria of offering high resolution, accurate and fast measurements. Another criterion that was seen as important was that minimal assumptions about the sample were made; so long as the sample was transmissive and capable of fitting in the interferometer it should be able to be investigated.

The need for a reference to compensate for drift in OPL was acknowledged, but aligning another source parallel (or collinear) to the OPO output was problematic. The reference should ideally travel the same distance in each arm as the OPO output and pass through / reflect from the same optical components. To be affected in the same manner as the OPO output at the beamsplitters the reference would need to be of a similar wavelength. Passing through the focussing and collecting lenses was not viable; their numerical aperture was too small to allow a beam to pass by the sample and the alternative of having collinear beams could have led to interference between the reference and pulsed

beams.

A far less convoluted and more elegant solution was to take advantage of the fact that Mach-Zehnder interferometers output light in two directions. With a Michelson arrangement incoming light either leaves the interferometer in one direction or is reflected back towards its source. With a Mach-Zehnder arrangement light will not be reflected back towards its source but will instead pass through or be reflected from the final beamsplitting cube. Which direction and to what extent light exits is governed by the OPLs of each arm. Consider the situation where one arm oscillates over several wavelengths while the other arm is fixed. The intensity of light exiting the interferometer in both directions will oscillate accordingly, although they will necessarily be mutually out of phase in their response to changing OPL offset. The total optical power exiting the interferometer will equal the optical power entering it (assuming no losses through absorption, etc occur). In the FTSI setup described above half of the light leaving the interferometer is unused. By inserting a detector to collect this light a perfect reference is created. The outputs from the interferometer are identical in every way other than being mutually out of phase. Therefore if one of these outputs is altered in any way (such as being spectrally filtered) and then the two are compared, any observed difference relates only to the way in which one output was altered. This type of differential measurement means that anything that alters one output, alters the other identically and so directly comparing them eliminates the effect of whatever gave rise to the alteration. Such an arrangement is shown in Figure 4.10.

Due to the addition of a second detector, the planar beamsplitters used in the interferometer were replaced by beamsplitting cubes. Both arms in the interferometer should ideally be identical in terms of the lengths of air and glass they contain. Looking at the interferometer and considering that light exiting from both directions of the second beamsplitter will be detected, it is clear that were planar beamsplitters to be kept in place this condition could not be met. The interferometer is 'balanced' for one exit direction; light that passes through the first beamsplitter will reflect from the second and exit while light that is reflected from the first beamsplitter will pass through the second and exit. Looking at the other possible exit direction; light in one arm passes through

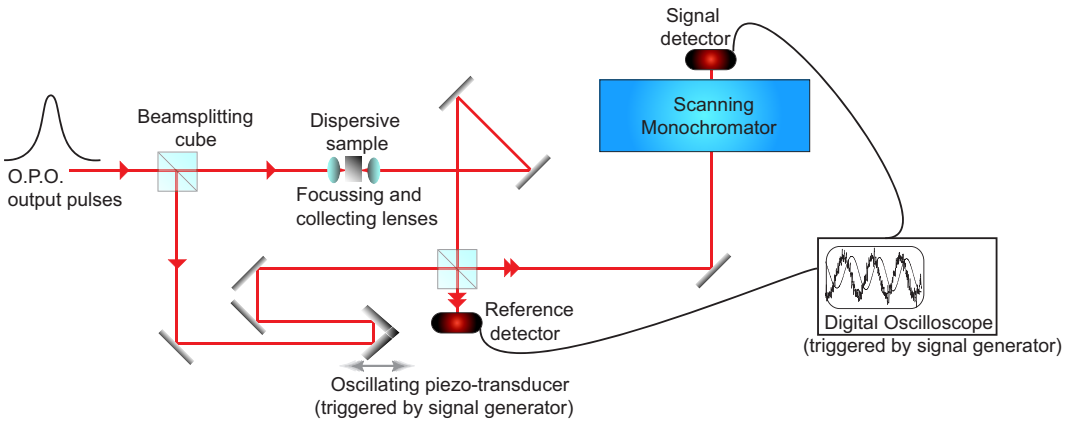

Fig. 4.10: Schematic of time domain spectral interferometry

both beamsplitters while light in the other reflects from both beamsplitters. In the latter case light in one of the arms will pass through more glass than light in the other. While the glass is unlikely to be highly dispersive, this still represents an unnecessary (and easily removed) error. Light passing through beamsplitting cubes (assuming correct alignment) travels through an equal length of glass irrespective of whether it is fully reflected, partially reflected or fully transmitted. Unfortunately one factor prevents the interferometer from being completely balanced. No compensation was made for the focusing and collecting lenses in the sample arm.

The relationship between the signals from the two detectors can be considered differently. The 'signal' detector carries information about the phase changes that have occurred to the filtered wavelength only. The 'reference' detector carries spectrally weighted information about the phase changes that have occurred to the entire pulse bandwidth. Comparing the two allows the difference in accumulated phase between one spectral component and the averaged pulse spectrum to be found. This also highlights the one requirement for this reference to be valid. If, during a measurement (or series of measurements) the spectral distribution of the pulses changes, then the reference can not be guaranteed to be valid. Consider a pulse with a Gaussian spectral distribution, it has a spectrally weighted average wavelength (i.e. an "effective" wavelength) corresponding to the central wavelength of its spectrum. If this pulse moves

spectrally, or its spectrum becomes asymmetric its "effective wavelength" will have moved. Part of the measurement (or series of measurements) will have used one reference, but later parts will have used another reference. A spectral filter could be placed before the reference detector to force the reference to be constant. However, the OPO used in this experiment exhibited spectrally stable output over many hours, while the highest resolution measurement took a maximum of 15 minutes and so this was deemed unnecessary.

Monitoring the signal from the reference detector would give an intensity trace whose amplitude was indicative of the OPL offset. Without a variable delay to provide a minimum and maximum and hence indicate the relationship between intensity and phase explicitly the reference trace would be difficult to treat.

Using Fourier transforms to extract phase data from a measured spectrum is a proven technique that is appropriate to many situations. One fundamental feature of Fourier transforms is that the resolution in the transformed domain (i.e. the time domain if the measurement was a spectrum) is proportional to the range over which the measurement was made.

To achieve high temporal resolution, it is necessary to scan over a large spectral range. For a narrowband source this may mean that the intensity is zero over the majority of the scanned spectrum. The bandwidth of the source determines the spectral range over which phase is measured and so determining the phase transfer function over a significant range with a narrow bandwidth source may involve a large number of measurements (with the source centred on a different frequency for each). Due to this redundancy, frequency domain Fourier transform spectral interferometry is a slow process if a broadband source is not used.

Without altering the physical setup of the experiment, the same data (a phase transfer function) can be extracted by taking time domain interference traces at different wavelengths instead of performing a standard FTSI measurement. Time domain interferometry typically uses Fourier transforms also, with the requirement for high resolution traces being a sufficiently large delay range and well resolved fringes.

Changing to time domain interferometry involved stepping the monochro-

mator through a chosen wavelength range and at each step scrolling the delay stage over a short distance, while recording the resulting interference pattern. Each time domain interference pattern could then be treated to give a phase value. The automated stage (with a precision of $\sim 1\mu$m) did not offer suitably fine adjustment of the delay and so was replaced by a piezo-transducer (PZT), which was capable of far finer movement (quoted precision was <50nm). The PZT also had the advantage that it responded nearly instantly to an applied voltage and could therefore be used to alter the OPL of an interferometer arm quickly.

Any change in OPL offset, due to either the PZT moving or thermal drift, will alter any time domain interference traces. As discussed above if the reference and signal detectors are measured simultaneously they will each experience the same change due to this. The scale of delay cannot be established from the voltage applied to the PZT only due to thermal drift, but the scale of delay in the reference and signal interference traces will be identical. This shared scale means that comparisons between the reference and signal traces are valid. In this respect, treating the reference is straightforward; the phase value given by the reference trace need only be subtracted from the phase value given by the signal trace to validate the measurement at a given wavelength.

Because the delay scale could not be established with this setup, it was not necessary to have the PZT move in discrete steps. A sinusoidal (or other continuously varying) voltage can be applied to the PZT and so long as the electronics used to sample the detectors are triggered by the voltage applied to the PZT, two intensity versus delay traces can be recorded. To this end, the lock-in amplifier was replaced by a digital oscilloscope. This was capable of monitoring both detectors and was linked via USB and controlled by HPVEE in the same manner as the lock-in amplifier. There were several advantages that the oscilloscope had over the lock-in amplifier. Firstly, it could be operated far more quickly (but still retained the ability to integrate over many traces if required). Also, the oscilloscope screen showed in real time the interference between the light in each arm of the MZ interferometer, which allowed beam alignment to be carried out easily and with great accuracy. Perhaps most useful was the ability to look at the interference traces over any chosen delay

range. The oscilloscope took 2000 points per trace (for each channel), and by varying the position and width of the oscilloscope window this could be chosen to correspond to any number of interference fringes (accepting an absolute limit due to the range of movement of the PZT). Instead of oscillating the PZT over a large range (such as 1mm), a low amplitude sinusoidal voltage (0.2V, 80Hz) moved the PZT through a distance of $\sim$10 wavelengths.

Conventional time domain interferometry techniques rely on having a very large number of fringes in their traces and again rely on Fourier transforms to provide the measure of phase. A requirement for Fourier transforms to be valid is that the spacing between points in the trace to be transformed is constant, which is complicated by thermal OPL drift in the interferometer arms.

Some researchers[18] suggest that a large delay range is needed to provide high resolution in a time domain interferometry measurement. The reason for this is that typically the position of each interference maxima is taken as one data point, effectively limiting resolution of phase to 2π radians (if OPL offset is zero). Resolution can be increased at the expense of a proportional increase in noise by operating with a large OPL offset. In this experiment the PZT was chosen to oscillate by $\sim 10\mu$m only, centred at zero delay, but there is no reason that the measurements should suffer from low resolution. Setting the oscilloscope window to be several wavelengths long gives hundreds of points per wavelength (i.e. per 2π radians).

Extracting a single value of phase from such data (assuming noise is not catastrophically large) is readily achievable to a high degree of accuracy.

The reason a lock-in amplifier was chosen initially was because of its ability to minimise noise by removing spurious signals at frequencies away from the chosen chopping frequency.

By looking at the interference of the two beams directly on the oscilloscope and having a reference a similar noise reduction takes place; only the amplitude variation due to interference is measured and any noise which effects this should be removed when the reference and signal traces are compared.

4.6 Conclusion

The steps taken towards constructing a sensitive and accurate dispersion measurement setup have been described. Careful consideration of the characteristics of both the ultrashort pulse sources and dispersive samples for which this experiment was designed led to a satisfactorily robust and accurate experimental setup. Exploiting the sensitivity offered by spectral interferometry while rejecting the use of Fourier transforms because of their inherent bandwidth dependent resolution provides an appropriate measurement technique capable of investigating dispersive samples under their expected operating conditions.

REFERENCE LIST

[1] J M Dudley L P Barry B Thomsen D T Reid, W Sibbett and J D Harvey, "Commercial Semiconductor Devices for Two Photon Absorption Autocorrelation of Ultrashort Light Pulses", *Applied Optics*, vol. 37, no. 34, pp. 8142–8144, 1998.

[2] T K Liang I E Day S W Roberts A Harpin J Drake H K Tsang, C S Wong and M Asghari, "Optical dispersion, two-photon absorption and self-phase modulation in silicon waveguides at 1.5 μm wavelength", *Appl. Phys. Lett.*, vol. 80, no. 3, pp. 416–418, 2002.

[3] N Carlsson K Asakawa N Kawai Y Sugimoto, N Ikeda and K Inoue, "Fabrication and characterization of different types of two-dimensional AlGaAs photonic crystal slabs", *J. Appl. Phys.*, vol. 91, no. 3, pp. 922–929, 2002.

[4] S Giugnit and T L Tansley, "Comment on the compositional dependence of bandgap in AIGaAs and band-edge discontinuities in AlGaAs-GaAs heterostructures", *Semicond. Sci. Technol.*, vol. 7, pp. 1113–1116, 1992.

[5] I H White R S Grant W Sibbett J B D Soole H K Tsang, R V Penty, "Two-photon absorption and self-phase modulation in InGaAsP/lnP multi-quantum-well waveguides", *J. Appl. Phys.*, vol. 70, no. 7, pp. 3992–3994, 1991.

[6] J Jasapara J W Nicholson, W Rudolph, F G Omenetto, and A J Taylor, "Full-field characterization of femtosecond pulses by spectrum and cross-correlation measurements", *Opt. Lett.*, vol. 24, no. 23, pp. 1774–1776, 1999.

[7] C Iaconis and I A Walmsley, "Self-Referencing Spectral Interferometry for Measuring Ultrashort Optical Pulses", *IEEE J. Quantum Electron.*, vol. 35, no. 4, pp. 501–509, 1999.

[8] M Joffre C Dorrer, "Characterization of the spectral phase of ultrashort light pulses", *C. R. Acad. Sci. Paris*, vol. 2, no. IV, pp. 1415–1426, 2001.

[9] A Chutinan and S Noda, "Waveguides and waveguide bends in two-dimensional photonic crystal slabs", *Phys. Rev. B*, vol. 62, pp. 4488–4492, 2000.

[10] F Marabelli L C Andreani L Pavesi M Galli, D Bajoni and G Pucker, "Photonic bands and group-velocity dispersion in SiOSiO2 photonic crystals from white-light interferometry", *Phys. Rev. B.*, vol. 69, pp. 115107, 2004.

[11] N M B Perney J J Baumberg M C Netti M E Zoorob M D B Charlton C E Finlayson, F Cattaneo and G J Parker, "Slow light and chromatic temporal dispersion in photonic crystal waveguides using femtosecond time of flight", *Phys. Rev. E*, vol. 73, pp. 016619, 2006.

[12] J P Likforman M Joffre C Dorrer, N Belabas, "Spectral resolution and sampling issues in Fourier-transform spectral interferometry", *J. Opt. Soc. Am. B*, vol. 17, pp. 1795–1802, 2000.

[13] A Lacourt C Froehly and J-C Vienot, "Time impulse response and time frequency response of optical pupils. Experimental confirmations and applications", *J. Opt.(Paris)*, vol. 4, pp. 183–196, 1973.

[14] P OShea S Akturk, M Kimmel and R Trebino, "Highly simplified ultrashort pulse measurement", *Opt. Lett.*, vol. 26, pp. 932–934, 2001.

[15] D Yelin D Meshulach and Y Silberberg, "White Light Dispersion Measurements by One and Two-Dimensional Spectral Interference", *IEEE J. Quantum Electron.*, vol. 33, no. 11, pp. 1969–1974, 1997.

[16] Z Bor A P Kovacs, K Osvay and R Szipocs, "Group-delay measurement on laser mirrors by spectrally resolved white-light interferometry", *Opt. Lett.*, vol. 20, no. 7, pp. 788–790, 1995.

[17] I A Walmsley M Beck and J D Kafka, "Group Delay Measurements of Optical Components Near 800 nm", *IEEE J. Quantum Electron.*, vol. 27, no. 8, pp. 2074–2081, 1991.

[18] G Cheriaux L Lepetit and M Joffre, "Linear techniques of phase measurement by femtosecond spectral interferometry for applications in spectroscopy", *J. Opt. Soc. Am. B*, vol. 12, no. 12, pp. 2467–2474, 1995.

5. DATA TREATMENT AND RESULTS

5.1 Introduction and motivation

In this chapter the results from investigation of three types of waveguide based devices are outlined. Each device has been designed to have distinctive dispersive features and operates in either of the telecom communication windows. One of the prerequisites for all-optical long haul transmission is some form of dispersion compensation. Optical amplification and retiming provide a limited capacity to preserve signal integrity, but dispersion compensation is essential to allowing the fastest transmission rates over arbitrary distances. All components in an optical network, whether they are sources, fibres, multiplexors, switches, amplifiers etc will all apply some level of dispersion to pulses that pass through them. Many devices exist that can be used to compensate for a given dispersion, such as Spatial Light Modulators[1], Prism pairs[2], Bragg gratings and "chirped" dielectric mirrors amongst others. Almost all of these were developed and applied as intracavity components to allow the production of temporally short pulses. Ultrashort pulse production has highlighted dispersion and its uses and associated problems, but only now is serious effort being made to apply this to commercial optical communications. None of the devices mentioned above are capable of being integrated on-chip and much interest and research has been instigated by the potential of devices capable of being fabricated in common semiconductors. Chief amongst these are photonic crystals, which offer several appealing features including compactness, the potential for active tunability and ease of integration. The ability to fabricate in either the III-V compounds used for optical devices or Si may be the most significant advantage of photonic crystals. The potential for growing and fabricating multiple devices within one tiny volume of semiconductor, with no

limitations on composition dramatically increases the range of potential uses for photonic crystals.

Several of the methods of dispersion or group delay measurement discussed in the previous chapter have been used to investigate photonic crystals. Their range/limitation of application, sensitivities and operational requirements have been discussed, with particular emphasis paid to investigating waveguide based devices exhibiting low transmission and spectrally sharp dispersive features.

The devices described in this chapter have all been designed as prospective waveguide based dispersion compensators. Each offers the potential to engineer their phase transfer function which, combined with accurate descriptions of their behaviour, offers the opportunity to iteratively move towards a commercially viable device. Of equal importance is the opportunity to quantify the performance of the measurement technique itself, both in terms of the physical setup chosen and the methods of data treatment used to extract raw data and convert it to useful and meaningful values.

5.2 Methods for extracting phase transfer functions from raw data plots

5.2.1 Following interference maxima

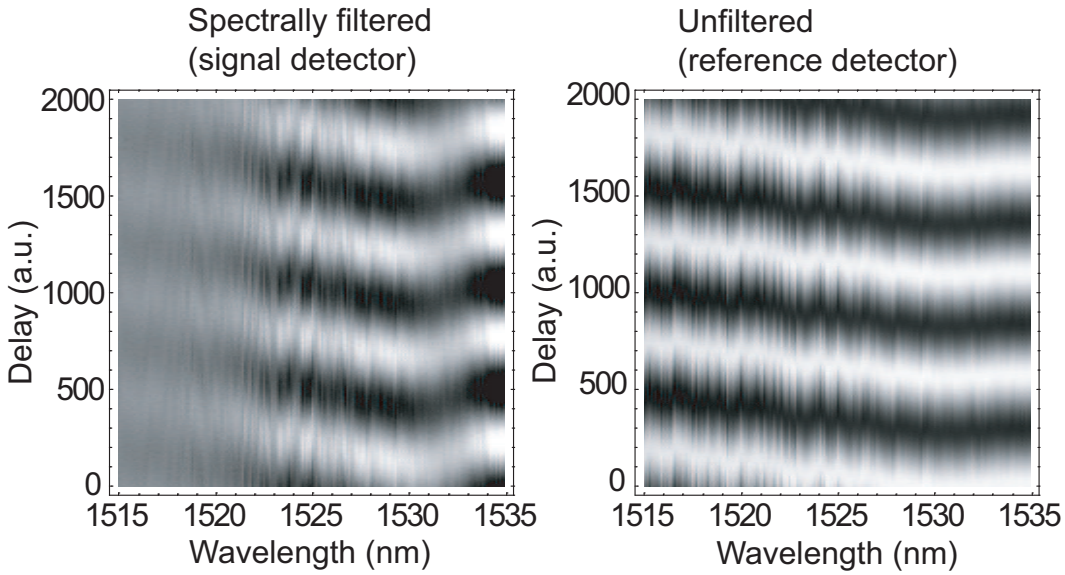

Fig. 5.1: Typical raw data plots from signal and reference detectors, taken using setup shown in Figure 4.10.

The final technique described in the previous chapter yields two 3-D plots, with axes corresponding to intensity, wavelength and delay. The plots generated by the experiment allow the presence or absence of dispersive features to be seen, but this must then be quantified. It should be noted that the wavelength axis refers to the wavelength transmitted by the filter (i.e monochromator) and so the plot corresponding to the unfiltered detector contains no spectrally dependent information. The wavelength scale is included in this plot merely to emphasise that each intensity vs. delay slice in the reference plot was taken simultaneously with the corresponding slice in the filtered plot. The unfiltered plot serves as a reference only and the filtered plot contains all of the spectrally dependent data.

Extracting a single 2-D trace describing a phase transfer function from these

requires some method of comparing spatial interferograms and quantifying their phase difference. Previous techniques [3] have fitted a sinusoidal function to the interferograms and so give a value of phase directly. In this investigation using such a method to retrieve phase is complicated by the nature of the piezotransducer. The movement of a piezo is not linearly related to the applied voltage, even when driving the piezo well below resonance and with a small amplitude, causing the scale to vary across the delay axis. In this case, knowing the amplitude and frequency of the sinusoidal driving voltage (0.2V and 80Hz respectively) does not allow the delay axis to be correctly scaled. One solution to this would be to measure the response of the piezo over the working range of driving voltages and then calibrate each trace based on this. Doing so would give traces with unequal spacing between adjacent points. This is difficult to process and would effectively weight some points in the trace more heavily than others. Also, as mentioned in the previous chapter thermal drift will alter the delay scale and this cannot be measured without an external reference. Alternatively taking traces over a range of delays for which the piezo's response is close to linear (but still large enough to produce several fringes) allows a function of the form:

$$Y = a_0 + a_1 Cos(a_2 Cos(a_3 z + a_4) + a_5) \tag{5.1}$$

to be used to determine the phase. Here, z is delay and variables a_0 to a_5 require fitting. a_4 is the desired phase, defined with respect to the start of the oscilloscope window. Fitting this function to traces, even with carefully chosen initial values of amplitude and frequencies, often gave poor fits and was time consuming due to the required number of iterations in each fit.

A simpler method of extracting phase was to quantify the delay scale through the spatial interferogram traces taken at each wavelength. Projecting the position of a maxima in the reference trace on this scale would give the phase transfer function. This has already been done by Beck and Walmsley[4]. To calibrate the movement of their piezo they monitored the intensity of a reference beam (in their case a HeNe laser) passing through the Michelson interferometer used in their experiment. As the piezo changed position (a

feedback loop caused the piezo to track an interference maxima as the spectral filter was scanned) the resulting intensity variation of the reference beam allowed its phase to be quantified. Knowing the HeNe's emission wavelength, this then gave an accurate measure of delay. Mapping the piezo's position on to the delay axis then gave the phase transfer function. As mentioned in the previous chapter this also eliminated the effect of OPL drift in the arms of the interferometer.

With no separate reference beam in our experiment an equivalent method can be implemented by following the position of interference maxima (or minima), similar to the technique used by Kovacs et al[5]. The separation of adjacent maxima (minima) in a trace corresponds to one wavelength. Comparing the position of a maximum (minimum) in the reference (i.e. unfiltered) trace with that of a maximum (minimum) in the signal (i.e. filtered) trace as a function of wavelength gives the phase difference between them. This eliminates any relative OPL drift between the two arms of the interferometer. Quantifying this difference can be achieved by normalising the difference in position by the separation of adjacent maxima (minima).

This simplifies the problem of extracting phase, although approximations are implicitly made by doing this. A significant approximation is that the maximum to maximum spacing changes negligibly over the width of the trace. Other limitations include having a high enough signal to noise ratio that the position of a peak/trough can be found reliably: small variations in intensity can cause the position of a maximum to seem to jitter. In each case manual selection of the delay range in which to search for maxima was needed, preventing a fully automated measurement. Despite these problems, representative phase transfer functions were obtained using this method. Assuming that the most valid judge of this method is the accuracy and sensitivity of the resultant phase transfer functions it produces, a sample with known, distinctive dispersion was analysed. A Fabry Perot etalon was chosen because of its simplicity to construct and model and its clear dispersion features. The mirrors used in its construction were antireflection coated 85 % reflective broadband dielectric mirrors. The mirror separation, approximately 350 nm, was chosen to give several distinctly visible fringes over the spectral width of the pulses. Mod-

elling the etalon, assuming perfect alignment, gave a transmission and phase profile as shown in Figure 5.2.

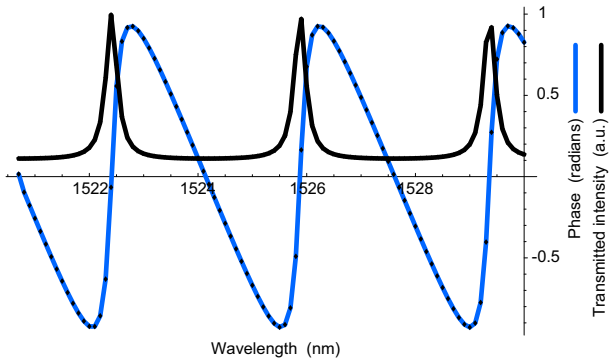

Fig. 5.2: Modelled transmission (black line) and phase transfer function (blue line) for Fabry Perot etalon.

Taking a high spectral resolution scan gave results as shown in Figure 5.3.

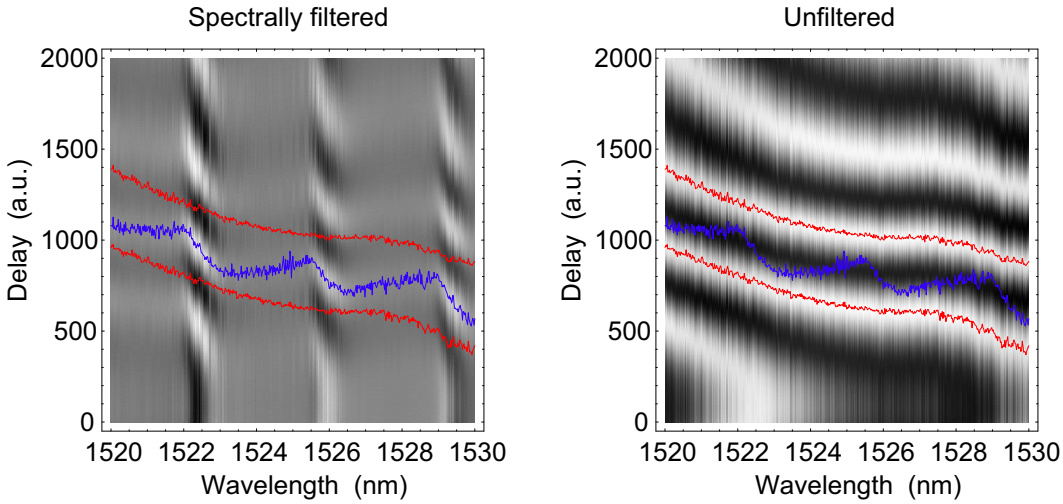

Fig. 5.3: Spectrally filtered (reference detector) and unfiltered (signal detector) plots for Fabry Perot etalon. The red lines trace adjacent maxima in the unfiltered plot and the blue line traces a maxima in the filtered plot.

The 'washed out' sections of the spectrally filtered trace suffer from such poor contrast because the transmission through the device is low at these

wavelengths. Despite this, the detected signal is well resolved and repeat measurements showed negligible variation in phase.

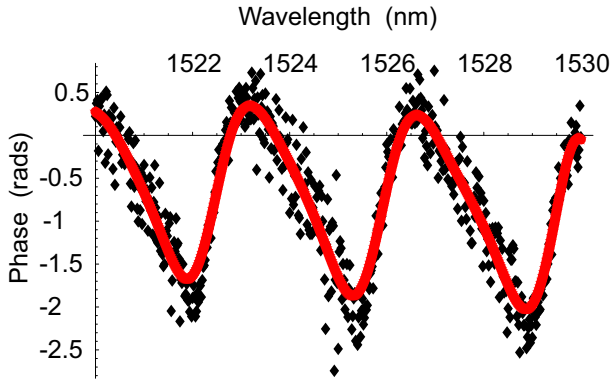

Fig. 5.4: Experimental phase (black points) of Fabry Perot etalon extracted by comparing maxima positions in filtered and unfiltered plots (red and blue lines in Figure 5.3). The red line in this figure is a fit given by a series of weighted polynomials.

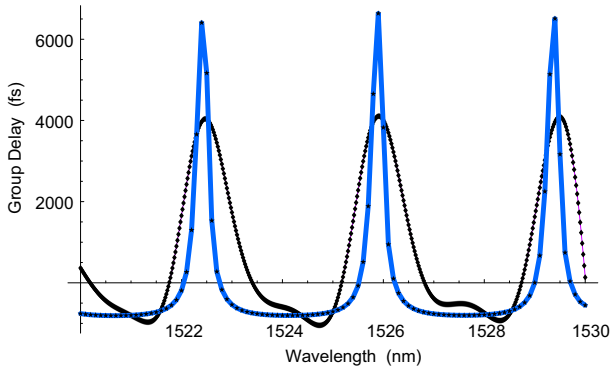

Fig. 5.5: Experimental group delay (black line, given by differentiating fit in Figure 5.4) and modelled group delay (blue line) for Fabry Perot etalon.

The raw data plots clearly show the changing phase difference between the reference and data scans, and this does translate clearly to the spectral phase. A significant level of noise can be seen in the phase transfer function however, especially when the signal level is low (such as between fringes in this case). With high noise levels fitting a function to the phase becomes crit-

ically important. Taking the first and second derivatives of the fitted phase places a significant emphasis on the method of fit. Any fit with high frequency components will lead to strong features in the derivated group delay and GDD. To remedy this some measure of smoothing is needed, however all of the important features of the phase trace must be preserved to maintain the integrity of the measurement. Techniques used previously include Fourier transforming the trace, applying a low pass filter and inverse transforming the trace[6]. When this was attempted with our data, frequently the fits matched poorly with the raw data. Fitting polynomials around each point guaranteed a close match to the raw data, but did not remove noise significantly. When a weighting factor was also included the fit showed a significant reduction in noise and matched the raw data well. The weighting function was chosen to be a smooth function, rather than one with abrupt edges. If a delta function (or some similar function) was chosen the fitted gradient around adjacent points was often found to be discontinuous. When differentiating such a fit, this resulted in strong, erroneous features being seen in the group delay and GDD. The weighting function used in this technique was chosen to be a Gaussian, centred on the point around which the polynomial was being fitted. Having a smooth weighting function such as this ensured that the gradient of the fit was smooth and represented the raw data well, but also removed shot noise (which was by its nature high frequency).

One of the reasons that the experimentally determined phases were so noisy was that the vast majority of the raw data taken was redundant; only one (or possibly two) points were used from any given intensity vs. delay trace. Ideally every data point taken should contribute to a cleaner (i.e. less noisy) phase transfer function. Unfortunately attempts to fit the spatial interferograms directly (as mentioned above) were very time consuming and often failed.

Tracking peak positions does give convincing results, but this method can only be used for a limited range of samples. If a peak being tracked moves out of the range of delay of the raw data plots, it becomes very awkward and time consuming to try and compensate for this. Because of this and the need to manually specify initial conditions this method is relatively slow and lacks robustness, but works well in some cases.

5.2.2 Projecting data

A method of comparison between two traces was needed that didn't involve fitting functions to the filtered and unfiltered traces, but following maxima suffers from problems, as discussed above. A more thorough comparison of reference and data traces would ideally use every point in each pair of traces. A 'correlation' of the two traces would effectively move one trace over the other and find the point at which they were most similar. This would provide improved accuracy, because the noise of one data point (the maximum) would no longer dictate the noise of the resultant phase transfer function. The reference and data traces are both finite in width however and so performing a correlation would involve arbitrarily selecting a 'window' in one of the traces and scrolling it over the other. The width and position of this window then may effect the derived spectral phase. This was found to occur, and while careful selection of window position and width did yield phase profiles as expected this method of data treatment was found to be inconsistent and lacking rigour.

A method was needed which would use all of the available data, without having to specify initial conditions or having to correct for artifacts (such as sudden jumps in phase caused by a maximum moving outside the width of the trace). This method should be applicable to any data taken, without making assumptions about the signal trace, the reference trace or their relatedness.

The mathematical operation which allows a comparison between two plots, without cutting or moving, is a projection. Applying a projection between two similar plots can give their phase difference. "Pearson's correlation coefficient" is a projection used to quantify linear correlation. It is defined as,

$$C_{Pearson} = \frac{\sum_i (x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\sum_i (x_i - \bar{x})^2}\sqrt{\sum_i (y_i - \bar{y})^2}} \tag{5.2}$$

where x_i is the ith term of a one dimensional list of values of the variable x, y_i is the corresponding term for variable y and $\bar{x}$ and $\bar{y}$ are the mean values of each list. The values from this function always lie between 1 and -1. The value of the function indicates whether the variables are positively correlated ($C_{Pearson} > 0$), negatively correlated ($C_{Pearson} < 0$) or not linearly correlated ($C_{Pearson} = 0$). The spectrally filtered and unfiltered traces should only differ

in their amplitudes (which does not effect the correlation operation) and their relative phase. Substituting $x_i....x_n$ with a function of the form shown in equation 5.1 and $y_i....y_n$ with a similar equation but with a different phase and graphing $C_{Pearson}$ against $Cos^{-1}(\theta)$ gives the graph shown in Figure 5.6.

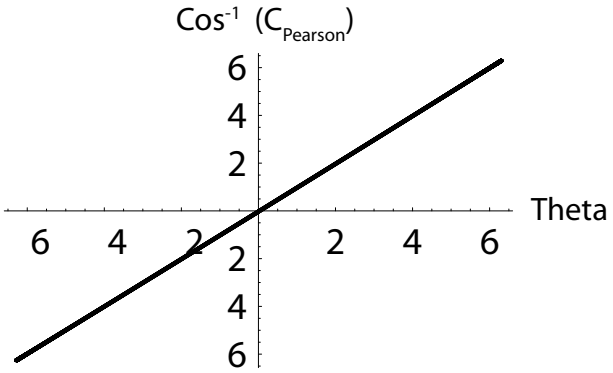

Fig. 5.6: Graph of theta calculated through $C_{Pearson}$ against difference in theta between two functions of the form shown in equation 5.1.

Proving rigourously that $C_{Pearson} = Cos^{-1}(\theta)$ is made difficult by $C_{Pearson}$ being a statistical measure, in that it deals with lists of discrete points, not continuous functions. However Figure 5.6 shows that (for the sinusoidal function corresponding to the oscilloscope traces at least), this is a reliable measurement method.

Comparing the use of $C_{Pearson}$ to tracking interference maxima can be carried out by again analysing the Fabry Perot etalon. Using the same raw data as Figure 5.3, $C_{Pearson}$ was used to generate the phase transfer function and group delay shown in Figure 5.7.

An ambiguity arises due to the reciprocal nature of $C_{Pearson}$(i.e. because $C_{Pearson}(x_i...x_n, y_i...y_n) = C_{Pearson}(y_i...y_n, x_i...x_n)$). The magnitude of phase difference between two traces can be quantified, but no indication of which trace is ahead in phase is given.

In practise three projections were taken per pair of traces, the first quantifying the phase difference. In the second of these the last data point of the reference trace and the first data point of the filtered trace were dropped be-

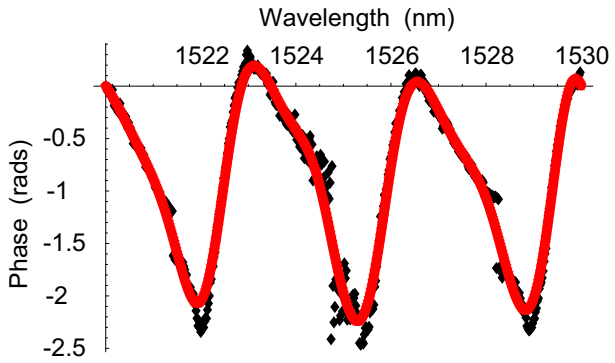

Fig. 5.7: Phase transfer function of Fabry Perot etalon extracted from raw data (shown in **??**) using $C_{Pearson}$ (equation 5.2) with fit function (red line).

fore taking the projection. For the third projection the opposite was done. Comparing these latter two $C_{Pearson}$ values indicates in which direction the traces would have to move relative to each other to increase or decrease phase difference. This in turn gives the sign of the phase difference. It should be noted that phase traces derived using this method jump/wrap around at 2π intervals. These jumps were removed to give phase traces, such as the one shown in Figure 5.8, which are far less noisy than those taken from the same data by following the position of interference maxima.

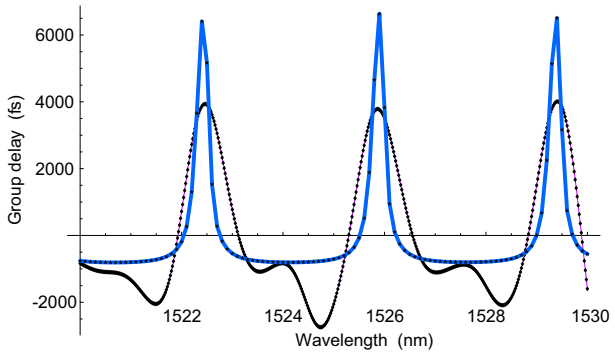

Fig. 5.8: Modelled group delay (blue line) and experimental group delay (black line, differentiated fit function from Figure 5.7.

The modelled mirror separation was chosen to be 335μm, to match the position of resonances in both traces. Two distinct differences between the

predicted and experimental graphs are obvious; the experimental peaks are shorter and wider than the model peaks (as with the peak tracking method above) and artifacts are seen between peaks, where none were expected. The former discrepancy is likely due to the etalon not being "ideal" as modelled. The etalon was aligned manually with several micrometers and so the mirrors were unlikely to be exactly parallel. Also the mirrors were old and possibly no longer as reflective as expected. Either, or both, of these factors would have the effect of lowering and broadening the group delay peaks as seen above (reducing theoretical mirror reflectivity to 71% gave a near perfect match), and so it seems highly plausible that the experimental group delay peaks are realistic. Choosing a fit with less smoothing (i.e. decreasing the Gaussian weighting function of the fit) increases peak height and narrows peak width, but only to a slight extent. Unfortunately this also increases the magnitude of the artifact seen between peaks. Were the artifacts to be removed a less smoothed fit might give a more realistic group delay graph. The artifacts result from a systematic error, which is evident in the phase transfer function in Figure 5.7. Jumps are seen around 1525nm and 1528.5nm in this graph. These spectral positions correspond to the reference trace and filtered trace being very close in phase and also the sample transmission being low. The low transmission results in highly noisy traces being given by the signal detector, although the reference detector traces remains unaffected. Taking the projection of the two, when they are close to being in phase, will give a value of $C_{Pearson}$ closer to 0 than would be given for noiseless traces. This effect occurs most strongly when it would be expected that the correlation should be close to 1 or -1. A value of 1 (or -1) can only result from the correlation of two directly proportional data sets (or where one data set is directly proportional to the negative of the other). Theoretically the spectrally filtered and the unfiltered trace should only differ in phase (relative amplitudes are ignored in the correlation). In reality both traces are subject to noise, and this is most noticeable in the spectrally filtered trace when the device is only weakly transmitting. This is clearly the case when the monochromator has selected a wavelength between transmission peaks in a Fabry Perot etalon or in the bandgap of a photonic crystal waveguide for example. The magnitude of the jumps is proportional to the noise present

in the signal trace. The trends in the phase transfer function are still clearly visible, but this effect is undesirable. To remedy this the average of a number of traces can be taken (reducing noise proportionally), with a corresponding increase in acquisition time. Another solution would be to alter OPL offset (the difference in OPL in the two arms of the interferometer) by a fraction of a wavelength and take another pair of plots. The jumps would then occur at different spectral positions and by combining the two phase transfer functions, the jumps could be removed.

Initially traces were correlated with their previous spectral slice in the same plot (i.e. correlations were taken between each spectral slice and its neighbour within the signal plot and similarly for the reference plot). The spectrally dependent phase profiles from each plot were then subtracted to yield a phase transfer function. This treatment does not work in practise, due to the problem discussed above; with noisy traces the correlation values will be lower than they should be. Consequently in the situation where the reference traces are close to noiseless and the spectrally filtered traces are noisy, comparing the two phase functions will result in an erroneous phase transfer function. The steps in the reference (i.e. spectrally unfiltered) trace will almost always be smaller than the steps in the filtered trace, leading to a component of the phase transfer function being dependent on the transmissivity of the device. For this reason correlations were carried out between reference traces and filtered traces rather than between adjacent spectral slices of the same trace.

To further test this method, another measurement was taken from the same etalon, but the OPL offset was increased. Increasing the offset should add a slope of constant gradient to the phase transfer function, which equates to a constant shift in the group delay. The magnitude and shape of features in the group delay graph should be identical to those found previously, although due to changes to the OPO alignment the pulse bandwidth (and hence spectral window of the plots) had shifted and broadened since the previous measurement.

In this measurement (Figures 5.9, 5.10, 5.11) although jumps in phase are present, they are so small that the fitting function was unaffected. It can clearly be seen that the measured phase profile is an accurate representation of that

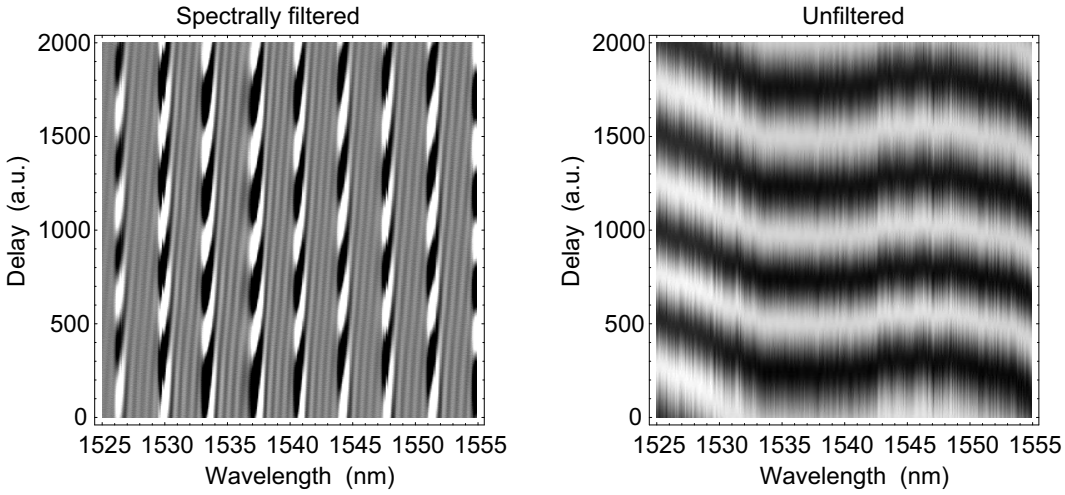

Fig. 5.9: Filtered and unfiltered plots of etalon with large OPL offset. Note the correspondingly steep gradient of maxima between transmission peaks.

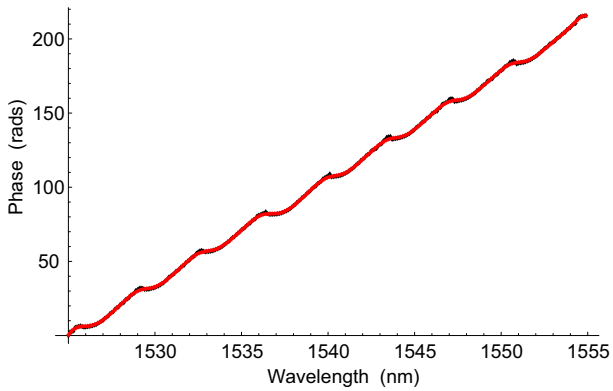

Fig. 5.10: Experimental phase transfer function of Fabry Perot etalon (black points) with fit function (red line) taken from data shown in Figure 5.9 using $C_{Pearson}$.

expected from such an etalon. The group delay graph now shows no artifacts between transmission peaks. It is likely that this was due to the increased OPO output power upon realignment, and a correspondingly less noisy plot from the signal detector. Unfortunately no comparison could be made to the peak tracking method here, because the phase gradient is so large.

It has been established that tracking peaks is a simple, accurate and yet

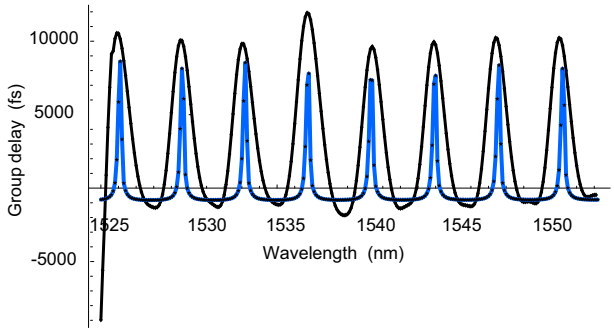

Fig. 5.11: Modelled group delay (blue line) and differentiated fit function from Figure 5.10.

limited method. Using projections is robust and requires no knowledge of the plots generated (indeed phase could be calculated as soon as each trace is taken, giving a close to real-time measurement), but can in some circumstances produce spurious artifacts.

To attempt to quantify noise inherent to the system, a final comparison between using projections and tracking peak positions was carried out by taking plots with no sample present. Ideally the phase transfer functions and group delay graphs generated should be straight lines and horizontal lines respectively. Any deviation from this can be regarded as a measurement of noise. Factors such as transmissivity of sample, input pulse energy and acquisition (i.e. signal integration) time per trace will all affect the magnitude of error due to noise, and so the traces were taken under similar conditions to the Fabry Perot traces. To compensate for the lack of a low transmission sample, the OPO output power was reduced to approximately half its previous value and the oscilloscope averaged over 8 traces instead of the 16 traces used for the other measurements in this chapter.

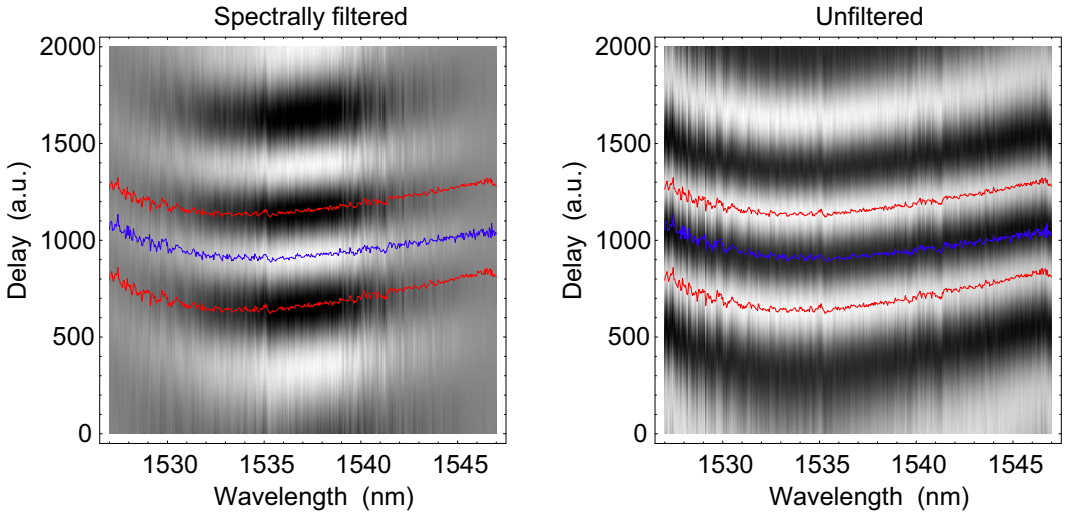

Fig. 5.12: Filtered and unfiltered plots with no sample present.

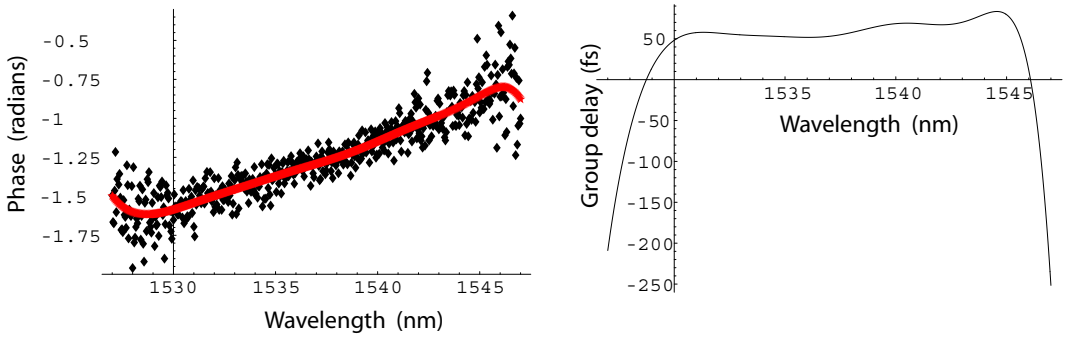

Fig. 5.13: Resultant phase transfer function and group delay, measured by tracking maxima, with no sample present (raw data shown in Figure 5.12).

Although at either edge of the spectral window the phase and group delay values aren't realistic, which is to be expected (considering the pulse intensity has fallen to near zero), the traces are consistent with each other. They both show the group delay to be of the order of 50fs and have little deviation from this in the spectral region where the pulse intensity is significant. Using $C_{Pearson}$ gave τ=54.5 $\pm$2fs, and the peak tracking method gave τ=54.6 $\pm$5fs. Further measurements taken in the absence of a sample gave similar values, indicating that these values could be regarded as an upper limit of noise-induced

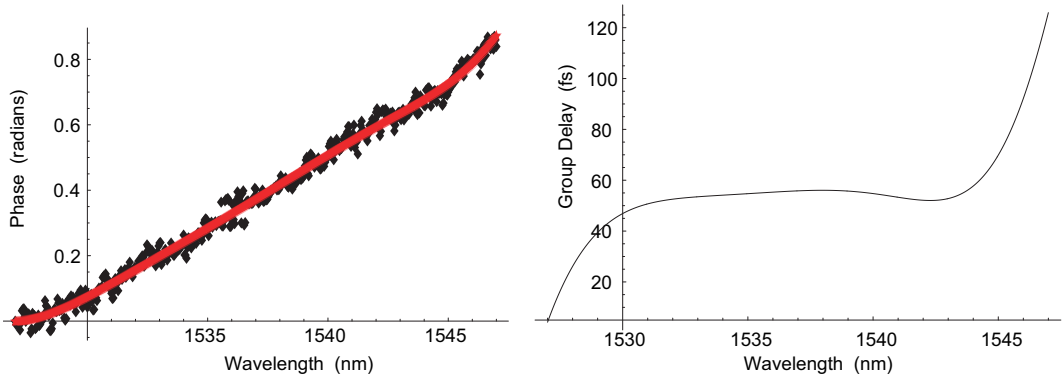

Fig. 5.14: Resultant phase transfer function (left hand side) and group delay (right hand side), again from data shown in Figure 5.12, measured using $C_{Pearson}$, with no sample present.

error. These values give an indication of the accuracy of the measurements in this chapter, and could be improved by decreasing the effect of noise through further averaging. Sprik and Kop[7] averaged over 256 traces in their measurements, compared to the 16 traces used for the measurements in this chapter. These values suggest very strongly that measurements from this experimental setup can be regarded as exceptionally reliable, considering that a variation of 10fs in any of the previously shown group delay plots would alter the plots negligibly.

5.3 Results and discussion

5.3.1 Cavity waveguides

The cavity devices, as shown in Figure 5.15, consist of a (in this case four sided) cavity and a thin slot which separates the cavity from the ridge waveguide. The slot acts as a beamsplitter; incident light will either be reflected immediately or enter the cavity and after one or more round trips exit. In either case the cavities have been designed to emit all incident light back into the waveguide in the opposite direction from which it entered. In theory these devices should behave very similarly to a Gires Tournois interferometer (GTI) (a specific case of Fabry Perot etalon, where the second mirror is 100% reflective). A perfect (and therefore impossible to fabricate) device would have a spectrally uniform 100% transmission. The cavity waveguide samples behave similarly

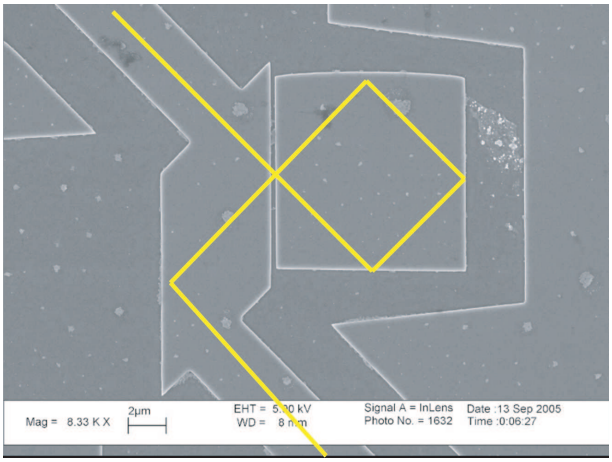

Fig. 5.15: SEM image of Mazilu cavity with overlaid beam path (yellow line)

to the Fabry Perot etalon. In both cases phase changes of π occur at each resonance. This should be expected and can be thought of as follows: Consider a cavity with uniform spectral transmission and a fixed optical path length per round trip. A GTI is an uncomplicated example of such a cavity. For a given frequency of light, a whole number of half wavelengths will fit exactly in to the cavity, and it follows that the light exiting the cavity will be in phase with the light entering the cavity. Lengthening (or shortening) the wavelength of the

input light causes a phase difference between the incident and exiting light. As the wavelength is scrolled this phase difference will rise to a maximum of π and fall back to zero periodically. The Fabry Perot does not illustrate this well because it exhibits a spectrally dependent transmission (the periodicity of the transmission changes and the phase changes are equal). The cavity devices were designed to have a uniform spectral transmission. For each pulse that enters a cavity, the phase of the light leaving the cavity corresponds to the superposition of a series of temporally separated pulses. Taking the GTI as an example cavity, the light reflected from it will correspond to:

$$S = r_1 + t_1^2 r_2^2 r_1 + t_1^2 r_2^2 r_1 + ... = t_1^2 r_2^n r_1^{|n-1|} = t_1^2 r_1^{|n-1|} \tag{5.3}$$

(noting that $r_2 = 1$ for a GTI). Here S is the normalised sum of all light reflected from the interferometer, r_1 is the reflectivity of the first mirror, r_2 is the reflectivity of the second mirror and $t_1(= 1 - r_1)$ is the transmissivity of the first mirror. Note that only magnitudes are represented here, no phase terms have been included.

The equation above is only true for a GTI and not for other cavity types, but every cavity type can be described by a similar series of terms. Each successive term will be delayed with respect to its predecessor by the OPL of one cavity round trip. This is significant when the OPL is not an integer or half integer number of wavelengths. At these wavelengths, each term will differ in phase from its adjacent terms. Consequently the relative strengths of each of the terms determines how wide (or narrow) a wavelength range a π phase shift will cover. This can be thought of as follows; each successive term is weaker than the previous one, but has accumulated a larger phase difference with respect to the initially reflected light (the first, and therefore strongest, term in the above series). If each term is far weaker than the previous one, there will be little contribution to the superposition of phases from later terms. This means that as wavelength is moved away from the condition of $OPL_{cavity} =$ integer/2, the resultant superposition of phases will change slowly. Conversely if each term in the above series is close in magnitude to the previous term, a slight change in wavelength will give rise to a large overall phase change.

The relative strengths of terms in the series is determined by the reflectivity of the mirror (or slot in the case of the Mazilu cavity structures). Essentially, a highly reflective slot will cause light in the cavity to be retained longer and in doing so accumulate an increased phase difference with respect to the light reflected from the slot. A weakly reflective slot will do the opposite and so a π phase change will occur over a narrower spectral range for a highly reflective slot than would occur with a weakly reflective slot. The reflectivity of a slot depends on its thickness, and so altering the slot width allows the spectral width of each resonance to be engineered.

To engineer the spectral position and separation of resonances the cavity size can be altered. Consider the condition for a resonance:

$$\frac{OPL_{cavity}}{2} = m\lambda_1 = (m+1)\lambda_2 \tag{5.4}$$

where OPL_{cavity} is the round trip OPL of the cavity, m is an integer, and λ_2 and λ_2 are the central wavelengths of adjacent resonances. For m $\gg$ 1 the difference between λ_1 and λ_2 will be far smaller than their difference when m $>$ 1. For a large cavity therefore, the spectral separation between resonances will be small and vice versa for a small cavity. To give a desired separation between resonances, one can simply choose the corresponding cavity length. Altering the cavity length by less than one wavelength changes the phase relation between light reflected from the cavity entrance (the first term in the above series) and light exiting the cavity after one or more round trips. This effectively moves each resonance spectrally.

With respect to engineering a phase transfer function, each cavity will contribute a comb of resonances, and for each comb the resonances can, as discussed above, have an engineered spectral width and position. It can be seen that with sufficient cavities cascaded together a waveguide structure with any desired phase transfer function should be able to be constructed. This assumes such waveguide devices are designed and fabricated to sufficiently narrow tolerances, which is why an accurate measure of their dispersion is desirable.

To compare the predicted behaviour of the Mazilu cavity waveguides with

the experimental results a simple model was used to show the expected phase transfer function, group delay and GDD. A single cavity was modelled using the equation:

$$T = \frac{r - e^{ink_0L}}{1 - e^{ink_0L}} = e^{i\theta} \tag{5.5}$$

Where the phase, θ and the transmitted light intensity, I are given by,

$$Arg(T) = \theta$$
$$Abs(T)^2 = I$$

Here T represents the transmitted electric field, r the slot reflectivity, n refractive index, k_0 the free space wavevector ($k_0 = 2\pi/\lambda_0$, with λ_0 representing free space wavelength) and L the cavity length. Plotting the phase transfer function calculated from this in the same wavelength range as one of the dips seen in the CW transmission spectrum gave the plot shown in Figure 5.16. Figures 5.17 and 5.18 are the first and second derivatives respectively. The

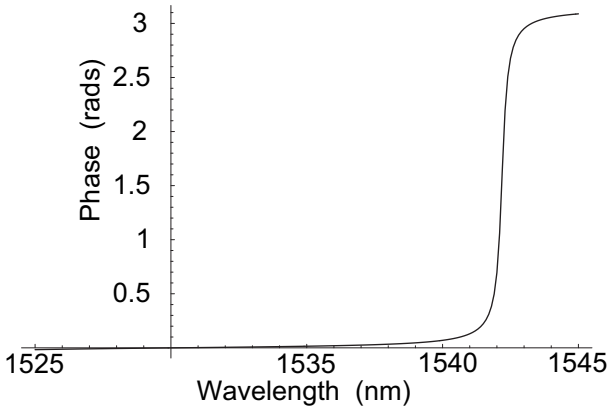

Fig. 5.16: Plot of theoretical phase transfer function, yielded by taking $Arg(T)$, with T given by Equation 5.5.

values used in equation 5.5 were r=0.96, n=3.29 (the effective index due to the waveguide rather than the material index of 3.182) and L=30 μm. The values for n and L were given by the composition and design (confirmed by SEM images) respectively. The reflectivity value of 0.96 corresponded with

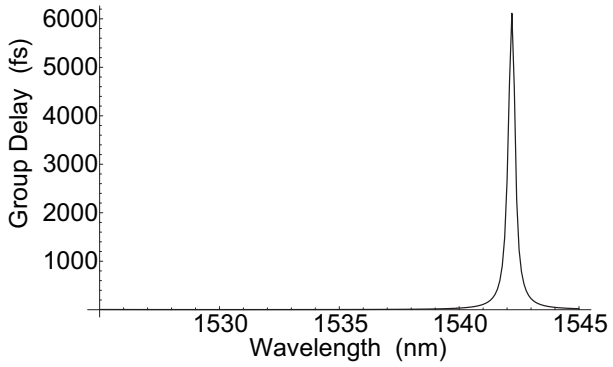

Fig. 5.17: Plot of theoretical group delay, given by differentiating phase transfer function shown in Figure 5.16.

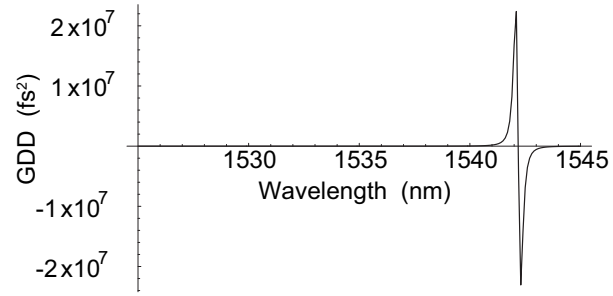

Fig. 5.18: Plot of theoretical GDD, given by differentiating group delay shown in Figure 5.17.

a slot width of ∼130nm. Unfortunately fabrication problems led to a larger than expected slot thickness, in some cases almost twice the desired value. Additionally there was a finite aspect ratio, meaning that the thinner the desired slot size, the shallower the etch depth around the slot. These fabrication problems were advantageous in that CW transmission measurements, which should have been spectrally uniform, showed decreased transmission around resonances. This provided a map of the spectral position of resonances. Taking plots over the spectral ranges indicated by the CW measurements (such as Figure 5.19), gave phase transfer functions, group delay and GDD traces as shown in Figures 5.20, 5.21 and 5.22 respectively. Please note: in the following graphs the cavity length, L was modified to 30.273μm, to spectrally match

the theoretical and experimental data. The experimental data indicated a far lower slot reflectivity than the 96% expected, and a value of 70% was used to give the theoretical traces.

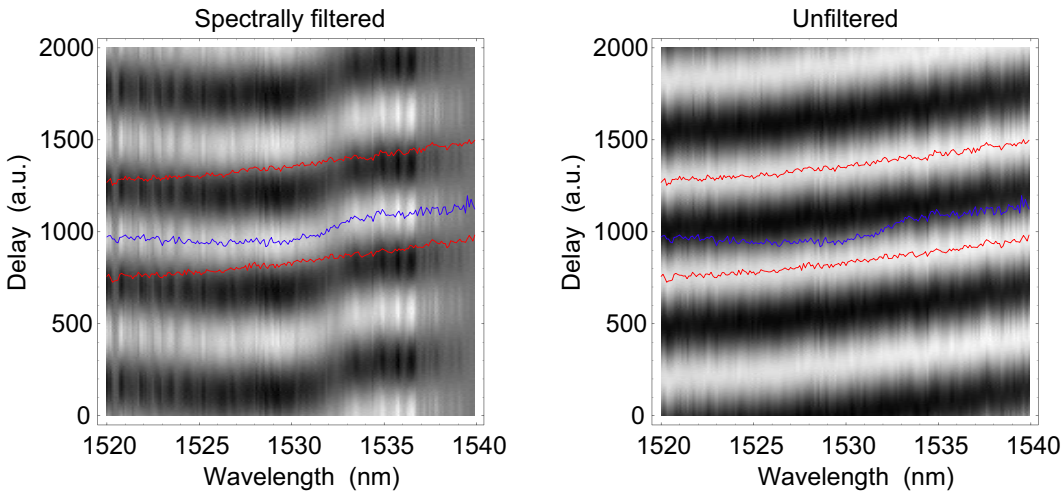

Fig. 5.19: Raw data plots from single cavity waveguide structure, corresponding to model given in Equation 5.5.

Figures 5.21 and 5.22 suggest that the experimental data and the chosen method of data treatment produce a good match to the theoretically predicted behaviour, assuming that altering the reflectivity value in the model is valid. Unfortunately the effect of the fabrication problems cannot be quantified and it is possible that the slot reflectivity is above 90%. Losses in the cavity and

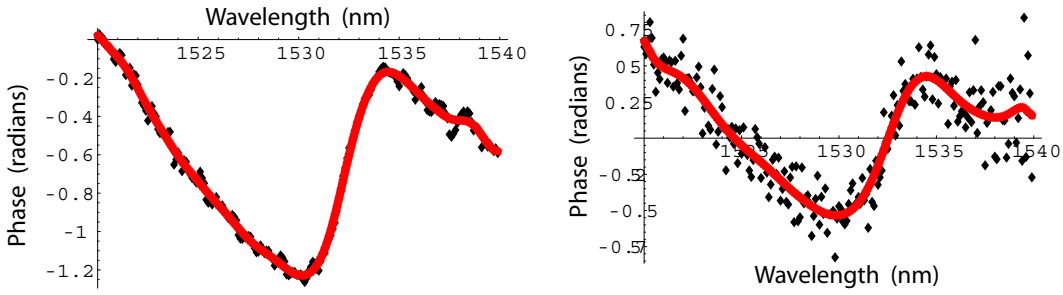

Fig. 5.20: Phase transfer functions extracted from Figure 5.19 using $C_{Pearson}$ (left hand side) and peak tracking (right hand side) methods.

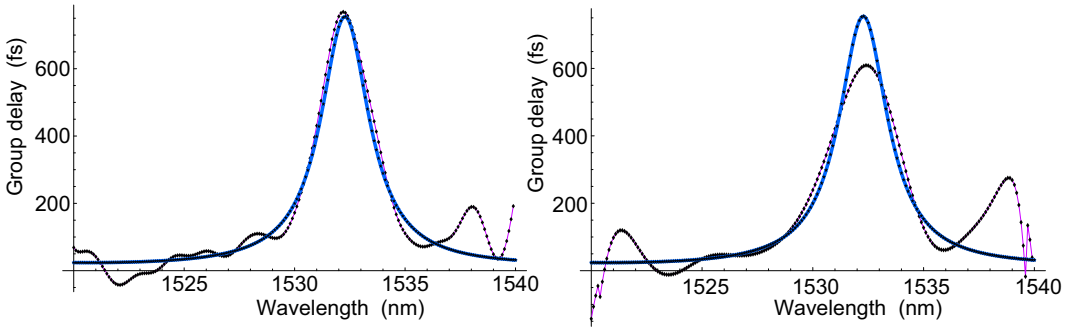

Fig. 5.21: Experimental delay (black lines) extracted from fits in Figure 5.20 using $C_{Pearson}$ (left hand side) and peak tracking (right hand side) and modelled group delay with R=70% (blue lines).

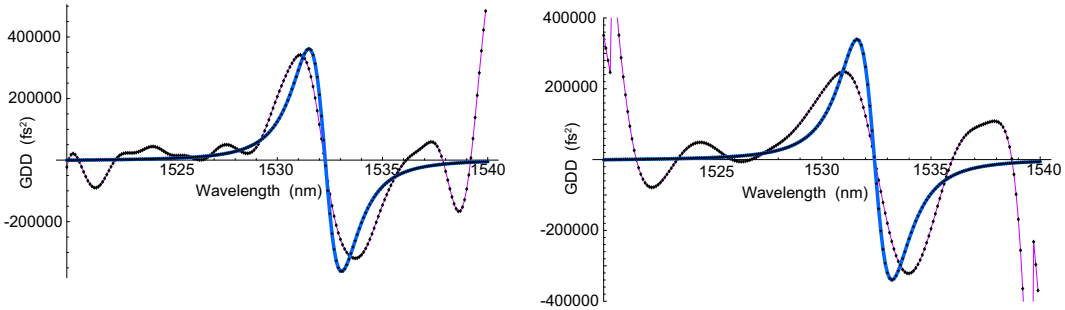

Fig. 5.22: Group delay dispersion yielded using $C_{Pearson}$ (left hand side) and maxima tracking (right hand side) yielded by differentiating respective delay graphs in Figure 5.20.

from the slot will effect the effective group delay seen and these have not been included in the simple model of the cavity. Losses per cavity round trip, either through absorption or out of plane scattering will reduce the contribution to phase from light remaining in the cavity for many round trips. This would alter phase transfer functions, group delay and GDD graphs in the same manner as lowering slot reflectivity. Seeing a correspondence between theory and experiment with a single cavity waveguide offered greater reassurance of the technique. Of greater interest however were the multiple cavity waveguides. No theoretical model had been made of these because the interaction between cavities was non-trivial and so their measurement was of particular interest.

The multiple cavity waveguides investigated showed multiple features with different dispersion characteristics. An example where two such features can be seen within one measurement is shown in Figure 5.23.

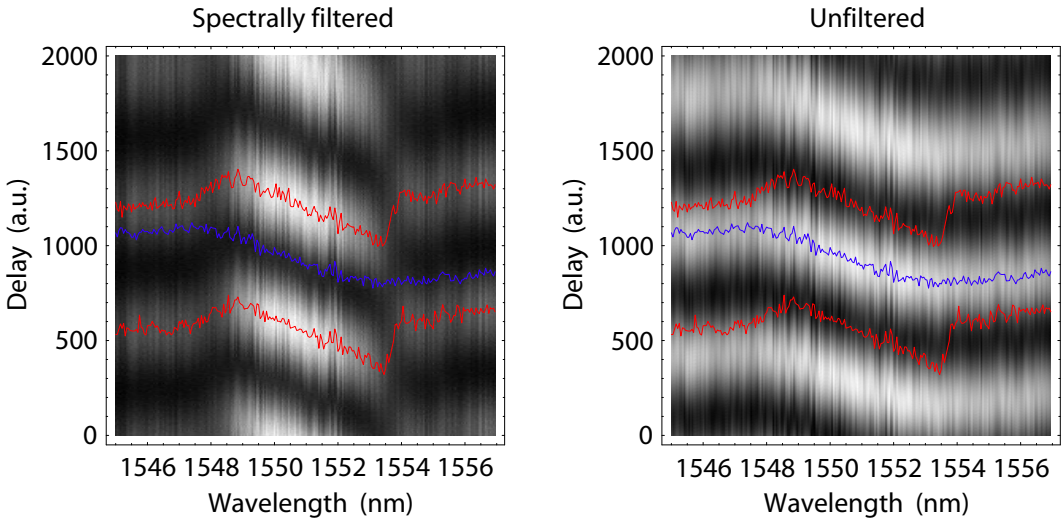

Fig. 5.23: Raw data of multiple cavity waveguide exhibiting two distinct dispersion features. The red lines trace adjacent maxima in the unfiltered plot and the blue line traces a maxima in the filtered plot.

Figures 5.24 and 5.25 illustrate that different resonances exhibit phase changes of variable magnitude. This is surprising considering the similarity

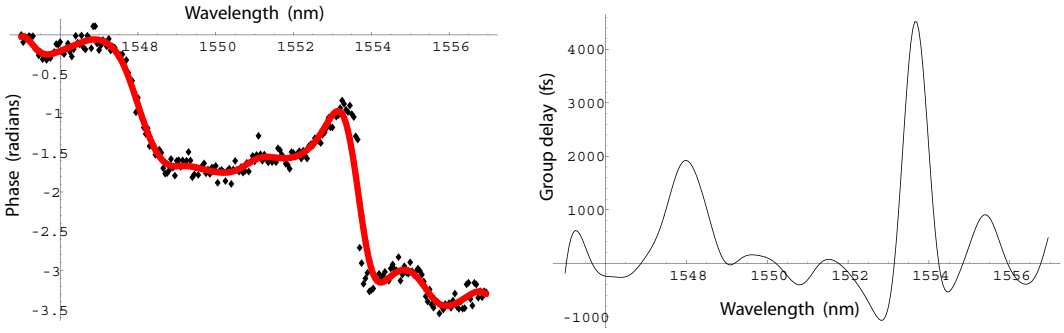

Fig. 5.24: Fitted phase transfer function and group delay from data shown in Figure 5.23, determined using $C_{Pearson}$. The fit (red line) to raw data (black points) in the left hand graph was differentiated to yield the group delay graph on the right.

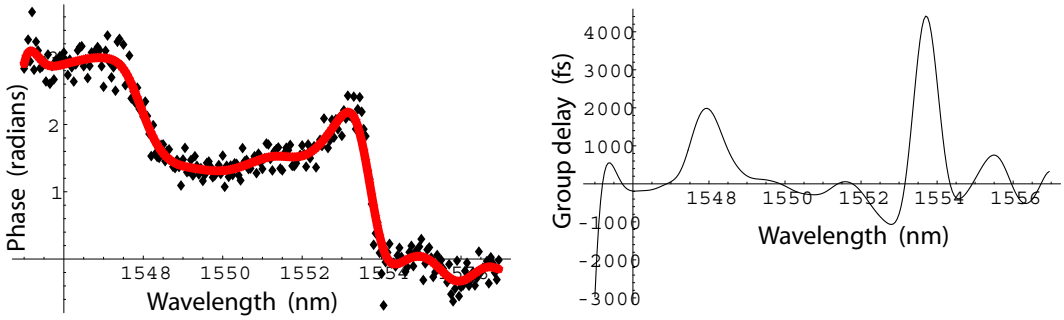

Fig. 5.25: Fitted phase transfer function and group delay from data shown in Figure 5.23, determined using peak tracking method.

of the waveguide cavities to GTIs; phase changes of π only may have been expected, but performing additional measurements yielded close to identical results. This perhaps indicates that the resultant dispersion of a multiple cavity waveguide is not a simple superposition of the dispersion features from each cavity. It may suggest that the equivalent of beating occurs between the cavities, resulting in complicated phase transfer functions. Whatever the cause(s) of this behaviour, devices of this type are interesting from the point of view of both fundamental physics and as potential commercial devices. Given the sensitivity and consistency of results offered by the phase measurement technique described in this thesis, further experimental investigation partnered with refined modelling could well lead to greater dispersion control being offered by these devices.

5.3.2 W3 AlGaAs photonic crystal waveguides

Figure 5.26 is an SEM image of the W3 waveguides whose potential as dispersion compensators was investigated. The waveguides were fabricated in the AlGaAs heterostructure described in chapter 4, designed for use around 1.55μm. The high aluminium content was chosen to avoid two photon absorption (TPA), which had been seen to severely limit transmission in waveguides fabricated in GaAs.

Light was coupled in to and out of ridge waveguides which tapered down

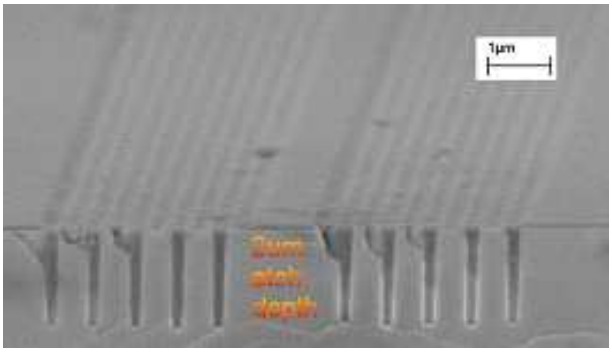

Fig. 5.26: SEM image of W3 waveguide

from their initial width (typically 5 μm) to the width of a section of photonic crystal waveguide in the centre of the sample. The photonic crystal section was a uniform 2D hexagonal array of holes with three rows of holes missing, forming a waveguide. Many such waveguides were fabricated on to one sample, each with a different combination of photonic crystal length and hole separation.

The W3 waveguides can not be modelled as simply as the cavity devices discussed in the previous section, and so to compare experimental and theoretical behaviour, existing data supplied by Tim Karle was used. The data, a graph of normalised frequency against normalised reciprocal wavelength, was calculated through a FDTD simulation[8]. In this simulation an initial electric field propagated in small discrete steps through a 2D model of the W3 waveguide, allowing the band diagram of the waveguide to be constructed.

Here the symbol a represents the lattice spacing (i.e. separation of adjacent holes) in the photonic crystal section and k_x represents the wavevector in the waveguide (i.e. $\frac{k_x}{2\pi} = \frac{n}{\lambda_0}$). The gradient of the above graph is therefore:

$$\frac{a}{\lambda_0} / \frac{k_x a}{2\pi} = \frac{2\pi}{\lambda_0 k_x} = \frac{1}{n} \tag{5.6}$$

The more widely familiar ω-k diagrams have gradients of

$$\frac{\omega}{k_x} = \frac{\omega\lambda_0}{n} = \frac{c}{n} = \frac{1}{v_g} \tag{5.7}$$

Thus the gradient of the graph above is directly proportional to the recip-

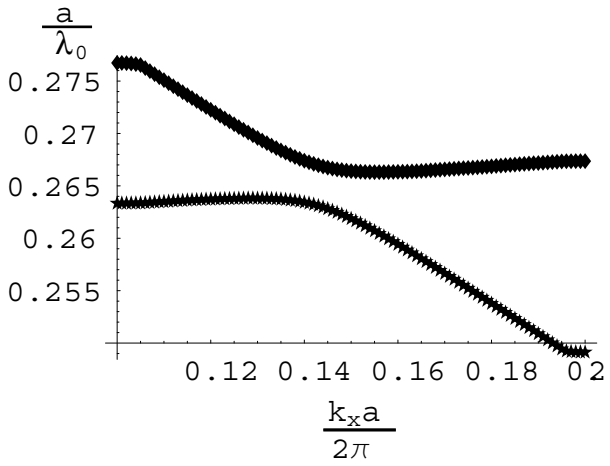

Fig. 5.27: Bandgap plot calculated from FDTD simulation, for W3 AlGaAs waveguides

rocal group velocity, with the speed of light as the proportionality constant.

To convert this to a more familiar group delay vs. wavelength plot the reciprocal of the derivative was taken and the x-axis was converted to wavelength (the value of a being known for every waveguide in the sample).

As can be seen from the equations, the length of a photonic crystal (of a given value of a), should be proportional to the scale of the group delay it imparts to light passing through it. Indeed, knowing the number of periods in each waveguide allows a correctly scaled graph of group delay vs. wavelength to be constructed. By multiplying the reciprocal group velocity by the length of photonic crystal (410nm x 297 periods = 121.77μm in the example below) the desired graph of group delay vs. wavelength was given.

The values of a and photonic crystal length determine the spectral positioning and magnitude of group delay respectively. Fitting the experimental data for a typical waveguide with the corresponding theoretical graph indicates what discrepancies in either value there may be. The curvatures of the two graphs should match, but due to time limitations the resolution of the theoretical graph is low. Despite this there is a clear similarity between experimental group delay measurements and theoretically predicted ones, as shown below. For ease of comparison the predicted group delay was shifted spectrally by

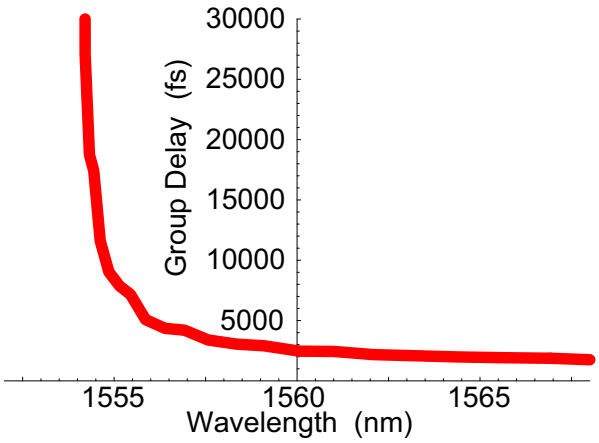

Fig. 5.28: Theoretical group delay calculated using data shown in Figure 5.27.

altering the value of a to 406.3nm.

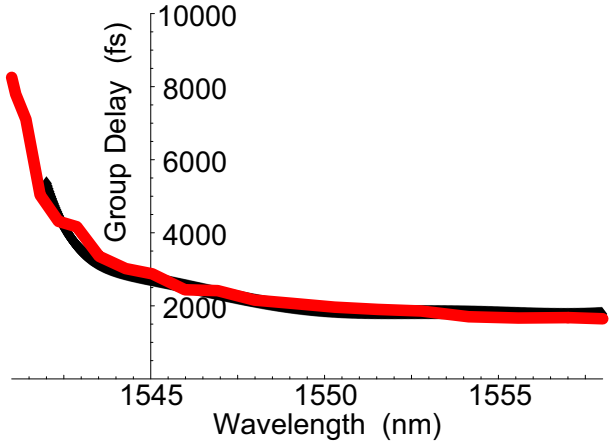

Fig. 5.29: Spectrally shifted theoretical group delay (red line) and experimental group delay (black line) for W3 AlGaAs waveguide.

Variances in fabrication quality limited the number of highly transmissive W3 waveguides, but several results similar to that shown in Figure 5.29 were taken, indicating a good correspondence between FDTD modelling and experiment.

An interesting feature (other than confirming the predicted dispersion near photonic bandgaps) exhibited by this sample was that of a "beating" effect seen

in several phase transfer functions, such as in Figure 5.30.

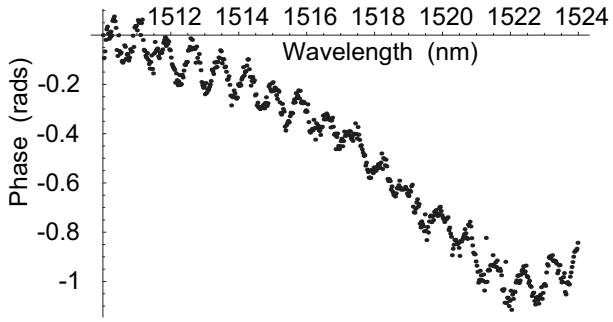

Fig. 5.30: Phase transfer function of W3 showing beating effect. The relatively low curvature of the graph corresponds to a region of low dispersion, far from the photonic crystal bandgap.

This effect was not seen in other samples with narrower waveguides and was initially attributed to the interaction of different modes supported by the waveguide. The W3 waveguides were too broad to only support a single mode, but higher order modes were not expected to be transmitted well or interact very strongly with the fundamental mode. Fourier transforming the group delay trace allowed the group delay difference between modes to be quantified as ∼10ps. Assuming the group delay occurs only in the photonic crystal section of the waveguide this implies that $n \sim 36$, which would normally imply very high losses (either through scattering or absorption). A more likely explanation is that there are reflections at one or both of the interfaces between ridge waveguides and photonic crystal waveguide. The beating equates to a distance of $\sim 900\mu$m, which is neatly the extra distance light reflected from an interface (or cleaved edge) would cover compared to light passing directly through the waveguide.

5.3.3 W1 Silicon on Insulator photonic crystal waveguides

The last of the sample types to be investigated was also an array of 2D photonic crystal waveguides. These were designed by Michael Settle and fabricated by the Intel corporation in silicon, designed for use around 1300nm.

With only one row of missing holes in the photonic crystal waveguide sections, these were far narrower (∼535nm wide) than the W3 waveguides. Due the wavelength requirements of the device a different OPO, capable of emitting pulses centred from 1280nm to 1350nm was used to investigate the W1 samples. The pulses from this OPO were short (∼250fs) which, in tandem with the narrow waveguide, indicated that nonlinearities may be excited which could in turn affect the dispersion experienced by transmitted pulses. However, the transmission (away from the bandgap) was high (approximately 20%), indicating little TPA.

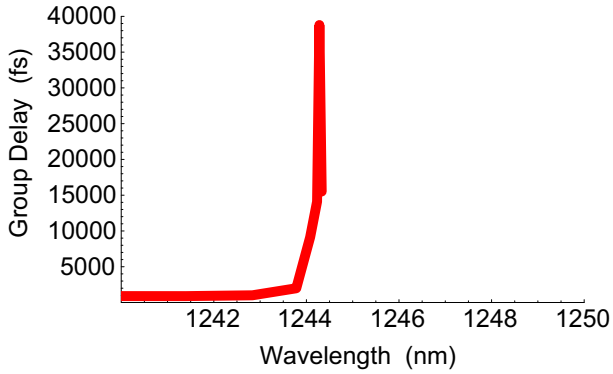

Fig. 5.31: Theoretical group delay for W1 Si waveguides.

Using a band diagram of the same type as for the W3 waveguide investigated previously, the graph shown in Figure 5.31 was plotted. As expected from two such similar structures, the modelled group delay for the W1 waveguides is similar to that of the W3 waveguides. The experimentally determined bandedge, confirmed by CW measurements, was not as predicted by the model.

Figure 5.35 shows a desirable shape, namely a central horizontal (or near horizontal) section with increasing GDD on one side and a decreasing GDD on the other. This characteristic could be exploited to compensate for dispersion imparted by other optical components. If a pulse had been subject to anomalous dispersion, meaning its short wavelength components led its long wavelength components temporally, then this could be (at least partly) corrected by passing through the W1 waveguide. Careful positioning of the pulse with respect to the bandedge would allow optimum, possibly even complete

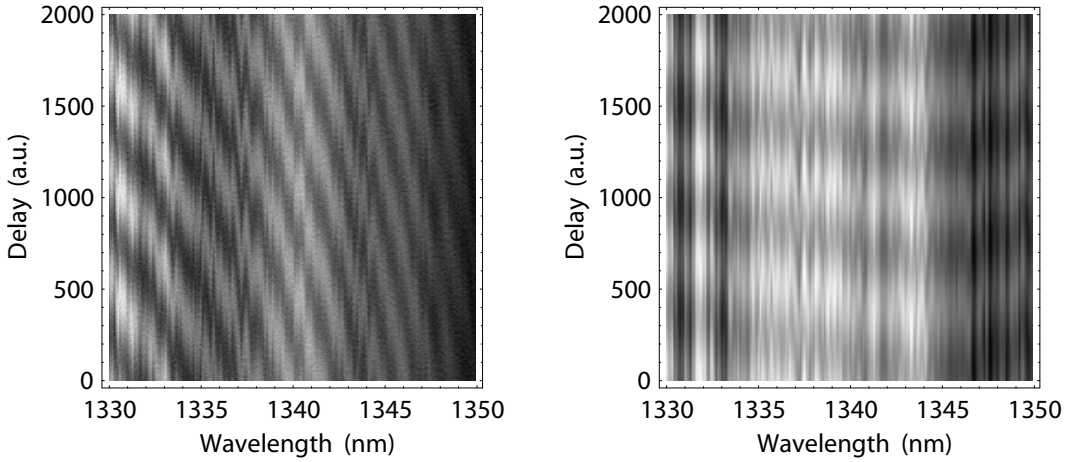

Fig. 5.32: Raw data plots taken from W1 photonic crystal waveguide

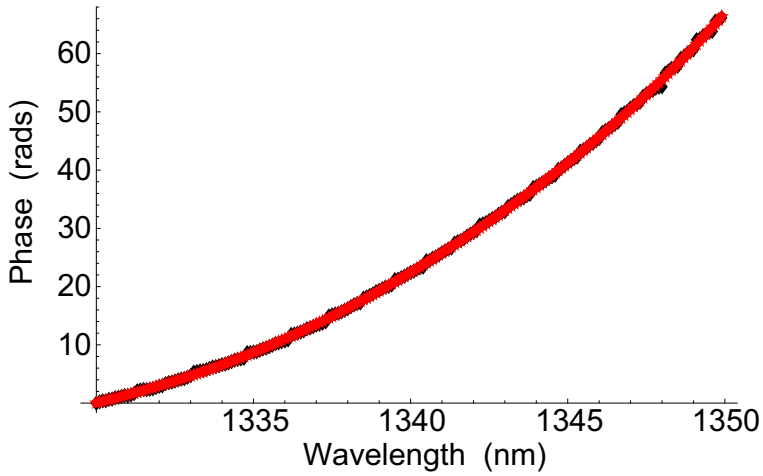

Fig. 5.33: Fitted phase transfer function found using $C_{Pearson}$. Again the red line is a fit to experimental data (black points).

dispersion compensation. Typically pulses experience normal dispersion and so a device exhibiting the spectrally reversed version of Figure 5.35 would be necessary to act as a dispersion compensator. This should occur at the other bandedge, but unfortunately this could not be experimentally verified due to the finite range of wavelengths the OPO could emit.

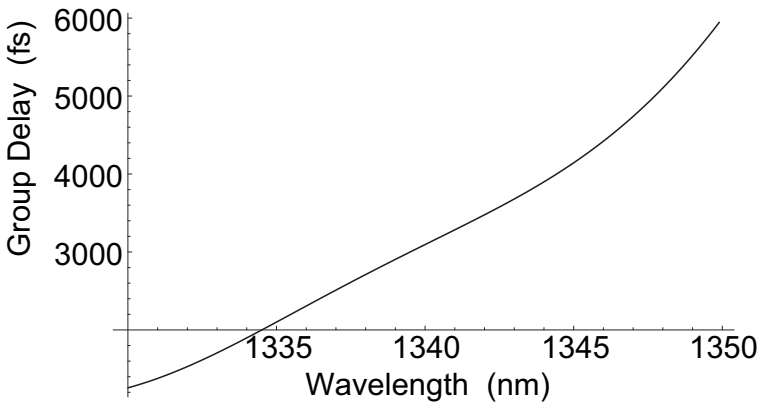

Fig. 5.34: Group delay given by differentiating fit in Figure 5.33.

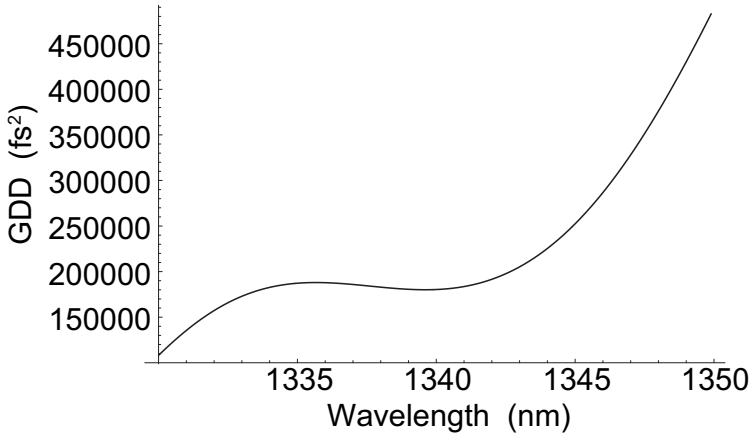

Fig. 5.35: GDD of W3 Si waveguide, given by differentiating group delay shown in Figure 5.34.

5.4 Conclusions

Raw data resulting from measurements taken with photonic crystal waveguides, cavity waveguides, a Fabry Perot etalon and no sample present were all analysed. The resultant phase transfer functions, group delays and GDDs given allowed a comparison of the relative strengths and applicability of the two data treatment methods discussed. Their similarity with modelled data and each other clearly demonstrates that this technique offers accuracy and sen-

sitivity, even when investigating challenging samples such as those with low transmission and spectrally narrow features. The samples discussed in this chapter, being previously untested and unrefined, gave an indication of the minimum quality of quantitative data this technique can yield. With greater integration time, more sensitive detectors or a fully calibrated piezo-transducer the low noise level and high sensitivity can only be bettered.

As far as is known this technique offers the first opportunity to very accurately quantify dispersion as a function of pulse duration and peak power for transmissive samples. It also offers to reveal the effect on dispersion of electrical bias (or optical pump power) and temperature in active devices. Using a CW source to investigate samples after using a pulsed source would easily allow the effect of nonlinearities to be established, which would provide critically important information allowing refinement of modelling and fabrication in photonic devices.

REFERENCE LIST

[1] V Binjrajka, D E Leaird C-C Chang, A W R Emanuel, and A M Weiner, "Pulse shaping of incoherent light by use of a liquid-crystal modulator array", *Opt. Lett.*, vol. 21, no. 21, pp. 1756–1758, 1996.

[2] O E Martinez R L Fork and J P Gordon, "Negative dispersion using pairs of prisms", *Opt. Lett.*, vol. 9, no. 5, pp. 150–152, 1984.

[3] M P Zinkin I I Tarhan and G H Watson, "Interferometric technique for the measurement of photonic band structure in colloidal crystals", *Opt. Lett.*, vol. 20, no. 14, pp. 1571–1573, 1995.

[4] M Beck and I A Walmsley, "Measurement of group delay with high temporal and spectral resolution", *Opt. Lett.*, vol. 15, pp. 492–494, 1990.

[5] Z Bor A P Kovacs, K Osvay and R Szipocs, "Group-delay measurement on laser mirrors by spectrally resolved white-light interferometry", *Opt. Lett.*, vol. 20, no. 7, pp. 788–790, 1995.

[6] I A Walmsley M Beck and J D Kafka, "Group Delay Measurements of Optical Components Near 800 nm", *IEEE J. Quantum Electron.*, vol. 27, no. 8, pp. 2074–2081, 1991.

[7] H J Kop and R Sprik, "Phase-sensitive interferometry with ultrashort optical pulses", *Rev. Sci. Instrum.*, vol. 66, no. 12, pp. 5459–5463, 1995.

[8] J D Joannopoulos S Fan, P R Villeneuve, "Large omnidirectional band gaps in metallodielectric photonic crystals", *Phys. Rev. B*, vol. 54, no. 16, pp. 11245–11251, 1996.

6. CONCLUSIONS

This thesis has introduced dispersion and its relevance to optical telecommunications. The fundamental dispersive properties of materials, nonlinearities and nanostructures have been discussed, highlighting the need for the development of compact, inexpensive, reliable dispersion compensators with dispersion profiles that can be arbitrarily engineered.

Current techniques employed to quantify dispersion have been critically reviewed, with emphasis placed on their applicability to measuring early generations of prospective dispersion compensators. Although these semiconductor nanostructure based prospective dispersion compensators offer great potential, as yet they can be awkward to measure. Limited coupling efficiency, low transmission and high nonlinear coefficients (particularly in the case of the AlGaAs / GaAs material system) all place demands on the technique used to quantify their dispersion. Sensitivity (i.e. the minimum power requirement necessary for a low noise measurement) was seen as a critical issue. Phase resolution and spectral resolution were also parameters that were vitally important, although to some extent they are mutually exclusive; in any given technique it is not possible (beyond a given point) to increase spectral/phase resolution without decreasing phase/spectral resolution.

Relying on nonlinear detection was found to reduce sensitivity to unacceptably low levels and because of this techniques based purely on nonlinearities as well as full pulse characterisation techniques (which require a linear and nonlinear measurement) were rejected. All linear dispersion measurement techniques rely on spectral interferometry (SI).

Techniques based on Fourier transforms are attractive in their experimental simplicity and possess a unique advantage; the phase transfer functions they produce are continuous functions. This is ideal when the group delay

or GDD is sought, no smoothing of the phase is needed before differentiation. Despite this advantage both time domain, and frequency domain, SI were rejected. In the time domain version, achieving high spectral resolution corresponded to scanning delay over a long range (>10cm), with interferometric accuracy (∼150nm). While technically possible, this is far from practical and thermal drift of the length of each interferometer arm would lower fringe contrast. In the frequency domain, the delay (i.e. optical path length difference between each arm of the interferometer) needed to be finite (determined by the required number of fringes and resolution of the spectral filter) and static. Thermal fluctuations invalidated data taken using this method. Even if these were compensated for (with an external reference, feedback loop and piezoelectric controlling delay), the finite delay was responsible not only for reducing maximum phase resolution, but reducing the strength of interference (and hence sensitivity) also.

The technique which offered the highest sensitivity was phase locked interferometry. Applying this complicated technique to waveguide based samples would be difficult and possibly result in invalid data. The use of an external reference to quantify drift in delay is valid assuming that the reference beam and characterising beam experience identical delays for any changes in optical path length within the interferometer. The ideal situation would intuitively seem to be propagating both beams nearly collinearly (i.e. parallel with a slight vertical/horizontal offset) through the interferometer. If the two beams were collinear then interference would occur inside the sample under investigation, which is clearly undesirable. If it was possible to propagate both beams through the sample, which is complicated by significantly different focal lengths and waveguiding conditions for well separated frequencies, the environments experienced by each beam are still likely to be different. The effect of free carriers and nonlinearities will invalidate changes in delay measured using an external reference beam. Passing the reference beam past the sample while the characterising beam passes through the sample will similarly produce invalid results; any changes in the sample (caused by temperature, free carrier population or nonlinearities) will not be taken in to account. For the purposes of measuring pulse dispersion, the use of an external reference cannot

be justified.

Every technique must make some compromise between sensitivity, speed of measurement, complexity, the range of samples that can be analysed, the range of sources used (incoherent, CW or pulsed), phase resolution, spectral resolution and noise. Accepting this, the technique outlined in this thesis is extremely competitive when compared to other spectral interferometry techniques, particularly when challenging samples such as the semiconductor nanostructures in this thesis are investigated.

In conclusion, a dispersion measurement technique which offers high sensitivity, requires no set features from the devices it is used to investigate (such as high reflectivity, high transmission or the presence / absence of waveguides) and yields reliable high resolution data was designed and tested. The flaws or limitations present in previously attempted techniques served to determine the essential features of the final robust technique.

6.1 Future work

Having established the validity and accuracy of the dispersion measurement technique in this thesis, many possible avenues of further work exist. The opportunity to examine the effect nonlinearities have on dispersion, particularly in complicated devices, such as semiconductor nanostructures exists. Measuring dispersion of a given structure with a broadband, incoherent source would allow the dispersion due to material and structure to be determined. By measuring dispersion as a function of peak power (either by varying average power or broadening/shortening pulse duration), greater understanding of the interplay between nonlinearities and nanostructures could be gained. An insight in to active devices (such as SOAs) could be granted by quantifying the dependence of dispersion on factors such as applied bias, temperature, pulse duration, peak power etc. Given sufficient time and a large enough range of samples, a complex model of dispersion as a function of pulse duration, power, free carrier population, temperature, and material could be built for commercially relevant structures.

An alternative could be to study the temporal dynamics of dispersion, by

inserting a beamsplitter in to the pulse trains before the interferometer. The pulses not passing through the interferometer could be used as a gate, for example optically pumping an SOA under investigation. By varying the delay between the pump and the dispersion characterising pulses, dispersion could be measured as a function of time.