Characterisation and Optimization of Ultrashort Laser Pulses

by

James B. Macpherson

A thesis
presented to the University of Waterloo
in fulfilment of the
thesis requirement for the degree of
Master of Science
in
Physics

Waterloo, Ontario, Canada, 2003

© James B. Macpherson 2003

Author's Declaration for Electronic Submission of a Thesis

I hereby decalare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Abstract

The ultrafast optical regime is defined, as it applies to laser pulses, along with a brief introduction to pulse generation and characterisation technologies. A more extensive description of our particular amplified pulse generation and SPIDER characterisation systems follows. Data verifying the correct operation of the characterisation system is presented and interpreted. Our laser system is then characterised in two different configurations. In each case, the data describing the system is presented and analyzed. Conclusions are made regarding the performance of both the characterisation and laser systems, along with suggested improvements for each.

Acknowledgements

I would like to acknowlege the assistance of my supervisor, Dr.D.Strickland, as well as very helpful explainations of the SPIDER technique given to me by Dr.I.A.Walmsley and Dr.C.Dorrer.

Contents

1 Ultrafast Laser Pulse Characterization **1**

1.1 Historical Methods 1

1.2 Modern Methods 2

1.3 Target Laser System 3

2 Theoretical Background **5**

2.1 Coherence 5

2.2 Analytical Pulse Description 5

2.2.1 Single Variable Approximation 5

2.2.2 Temporal Representation 6

2.2.3 Spectral Representation 7

2.2.4 Function Notation 9

2.3 Dispersion 9

2.3.1 Material Dispersion 10

2.3.2 Path Length Dispersion 11

2.4 Spectral Phase Measurement 13

2.4.1 Spectral Interference 13

2.4.2 Fourier Transform Spectral Interferometry 13

2.4.3 Self-Referencing Spectral Interferometry 16

2.4.4 The SPIDER Method 17

3 The SPIDER System **20**

3.1 Optical Apparatus 20

3.2 Spectrometer 22

3.3 Data Processing 24

4 Measured Data **26**

4.1 SPIDER Verification 26

4.1.1 Interferometer Imbalance and Fringe Sampling Rate 26

4.1.2 Spectral Shift Adjustment 27

4.1.3 Shear Calibration 27

4.1.4 Shear Approximation Error 28

4.2 Pulse Variation 30

4.3 Target System Characterization 31

4.4 Enhanced Compression Experiment 33

5 Data Interpretation and Analysis **36**
5.1 SPIDER Verification 36
5.1.1 Interferometer Imbalance 36
5.1.2 Fringe Sampling Rate 37
5.1.3 Spectral Shift Adjustment and Shear Calibration 38
5.1.4 Shear Approximation Error 38
5.2 Pulse Variation 40
5.3 Target System Characterization 40
5.4 Enhanced Compression Experiment 41

6 Conclusions and Recommendations **45**
6.1 SPIDER 45
6.2 Target Laser System 45

A The Fourier Transform **48**
A.1 Definition 48
A.2 Shift Theorem 48
A.3 Sampling Theorem 49

B SPIDER Software Description **51**
B.1 Spectrometer Readout 51
B.2 Calibration 51
B.3 Determining Spectral Phase 52
B.4 Temporal Reconstruction 52

C Optimization Code for the Enhanced Compressor **54**

List of Tables

1 Phase Fluctuation Statistics 30
2 Test Points for the Enhanced Compressor 33
3 Statistical Fluctuation Error 40
4 Compressor Parameter Effect on Phase Coefficients 42

List of Figures

1 Target Laser System 4
2 Pulse Phasor Representation 8
3 Dispersive Optical Systems 12
4 Interference Spectrum and Transform 14
5 Spectral SPIDER Optical Schematic 19
6 Temporal SPIDER Optical Schematic 19
7 The Michelson Interferometer 20
8 The Sum Frequency Generator 21
9 The Spectrometer 23
10 Interferometer Dispersion 27
11 Spectral Shifting Test 29
12 Shear Calibration Comparison 29
13 Shear Approximation Test 29
14 Phase Fluctuation Plot 30
15 Target System Spectral Phases 31
16 Derived Quantities for the Target System 32
17 Spectral Phase Coefficient Plots for the Enhanced Compressor . . . 35
18 Enhanced Compressor Optimized Pulse Data 44

1 Ultrafast Laser Pulse Characterization

Pulsed lasers form a large and important part of modern laser technology. Rather than producing a continuous electromagnetic wave indefinitely, pulsed lasers produce short bursts of radiation which typically last for less than a second. Laser pulses have a great variety of applications in science and engineering, as well as being of direct scientific interest themselves. Their utility and novelty is due to the very small time scales on which they act and to the high instantaneous power which may be produced using a relatively small amount of energy. Today pulses routinely reach femtosecond ($10^{-15}s$) durations and produce irradiances of $10^{20}\ W/cm^2$.

The earliest lasers, including the original ruby laser [8], produced pulsed outputs having durations in the millisecond or microsecond range. This pulsed operation was by limitation rather than choice however, as the conditions for laser output could simply not be maintained for longer periods. The first major advance in pulsed laser technology was the deliberate pulsed operation of lasers by "Q-Switching". This technology allowed lasing conditions to be imposed or removed quite suddenly, producing a high power pulsed output. Modern Q-switching produces pulse durations in the nanosecond range.

Pulsed operation by "mode-locking" [12, 1] followed shortly after the development of Q-switching. This method produces a short pulse by superposing many synchronised longitudinal laser cavity modes. Mode-locking allowed the generation of pulses as short as picoseconds ($10^{-12}\ s$), and more recently, femtoseconds. The reduction of pulse duration below the nanosecond range was a very significant milestone for laser pulse technology. It was at this point that pulse durations began to fall below the temporal resolution limit of electronic detectors, forcing the development of new diagnostic techniques. The generation, measurement and application of pulses exceeding this limit became known as the field of "ultrafast" or "ultrashort" pulse optics.

1.1 Historical Methods

Of the earliest ultrafast diagnostic methods, non-linear correlations and streak camera techniques were most common.

Non-linear correlations [14] are a family of pulse measurements which are mathematically similar to convolution. Two or more pulses are superposed and sent to a detector which is sensitive to their temporal overlap. Changes in the degree of overlap with the relative timing of the superposed pulses may be used to estimate

their longitudinal extent and determine a rough pulse duration.

The simplest form of correlation is the intensity autocorrelation. Two replicas of the unknown pulse are combined with a varying delay τ. The detection system then produces a non-linear power measurement which is time averaged over a period much greater than the pulse length. The ("background free") autocorrelation A may be expressed as a function of the unknown pulse intensity and the relative delay:

$$A\{\tau\} = \int_{-\infty}^{\infty} I_u\{t\} I_u\{t-\tau\} dt$$

The width of the measured signal vs. delay implies a pulse width if a particular pulse shape may be assumed. This required assumption is obviously a significant limitation. Further, the method lacks any sensitivity to the phase of the field, returning only intensity characteristics. Other forms of correlation may produce more definite and complete pulse characterization, but these measurements tend to become impractical as the pulse reconstruction becomes more complete.

Like correlations, streak camera measurements [15] convert pulse timing information into spatial variation. In a streak camera, the unknown pulse is sent to a cathode target, where it frees electrons. The electrons are accelerated by an applied electric field and then drift to a detector or detector array. Ramping or other variation of the applied field alters the acceleration, so that electrons produced by different time segments of the pulse are separated spatially. Their spatial distribution may be detected and forms a measurement of the pulse intensity in time.

The resolution of streak cameras tends to be limited to the picosecond range. There are obvious limits to the rate at which the accelerating potential can change, and resolution is degraded by the distribution of initial electron velocities as well as the mutual repulsion of the drifting electrons. As a further limitation, the cameras lack sensitivity to the phase of the pulse field.

1.2 Modern Methods

State of the art ultrafast diagnostics apply a similar principle to the early methods, in that temporal pulse variation is converted to some other form which is measurable. The two most important types of transduction for this purpose are spatial and spectral conversion. In some cases, a combination of temporal, spatial and spectral measurements are used to characterize unknown pulses, but the methods which are proving to be the most powerful do not require time dependent data. These "slow" measurements are not effected by the limited response time

of detection electronics.

The two most common modern ultrafast diagnostics are Spectral Phase Interferometry for Direct Electric Field Reconstruction (SPIDER) and Frequency Resolved Optical Gating (FROG). The former requires only spectral measurement, while the latter requires spectral and spatial measurement. Each of these techniques characterizes a pulse by measuring its relative spectral phase.

SPIDER [5] is a variant of a measurement technique known as Fourier Transform Spectral Interferometry (FTSI) [7]. SPIDER may determine pulse characteristics from two power spectrum measurements. The first measurement is simply the power spectrum of the unknown pulse. The second power spectrum is that of two unknown pulse replicas which are modified and recombined at a fixed relative delay. SPIDER is particularly suited to rapid measurement, as it requires a relatively small input data set and is not computationally intensive. SPIDER systems have been assembled which operate at a rate of 40 Hz [16].

In the FROG method [6], the unknown pulse is mixed with a gate pulse, such that only the overlapping pulse segments are routed to a spectrometer. The measurement of power spectrum vs. relative pulse timing produces a two dimensional data set. An iterative data process may be applied which uses the unknown pulse power spectrum and the gated pulse segment spectra to determine the unknown pulse spectral phase. The main advantage of FROG is a very simple optical apparatus. The two-dimensional data set also contains redundancies which allow calibration to be verified by a self-consistency check. The major disadvantage of FROG is the relatively large amount of data collection and processing required.

These two methods allow the most complete temporal pulse characterization available to date. The instantaneous intensity of the pulse is determined definitely. The phase of the field oscillations is determined within a constant, which is equivalent to determining the instantaneous frequency.

1.3 Target Laser System

As reported in this document, a SPIDER system was built and used to characterize the output of a particular pulse laser system, which is described here.

The target system (figure 1) consists of a mode-locked laser oscillator and an accompanying laser amplifier. The oscillator produces an infrared pulse using an optically pumped Titanium:Sapphire crystal as the gain medium. A four-mirror cavity with an intra-cavity prism pair produces approximately 60 thousand resonant longitudinal modes within the very wide gain bandwidth of the crystal. These modes radiate simultaneously and synchronise naturally, producing pulses of

roughly 20 nm bandwidth centered at approximately 800 nm. Output pulse energy is around 5 nJ. Output pulse duration is on the order of tens of femtoseconds and the pulses repeat at approximately 100 MHz.

The chirped-pulse amplifier [9] consists of three stages. The initial stretcher spreads the oscillator pulses in time to a duration in the hundreds of picoseconds. After stretching, single pulses are selected from the oscillator pulse trains and amplified. These seed pulses enter the regenerative cavity via a Faraday isolator and make several round trips through a second Ti:Sapphire gain medium. The gain medium is optically pumped by a synchronised pulse from a Q-switched Neodymium:Yttrium-Lithium-Fluoride laser. After a set number of passes, the amplified pulse has its polarization altered by a Pockel's cell and is thus routed back in the direction from which it came. The Faraday isolator uses a direction dependent polarization change to direct the amplified pulse to the compressor, rather than allowing it to retrace the path of the seed pulse. The grating compressor approximately reverses the effect of the stretcher. Amplified pulses have energies in the milliJoule range and durations in the tens to hundreds of femtoseconds. The amplified pulses are produced at a rate of roughly 1 kHz.

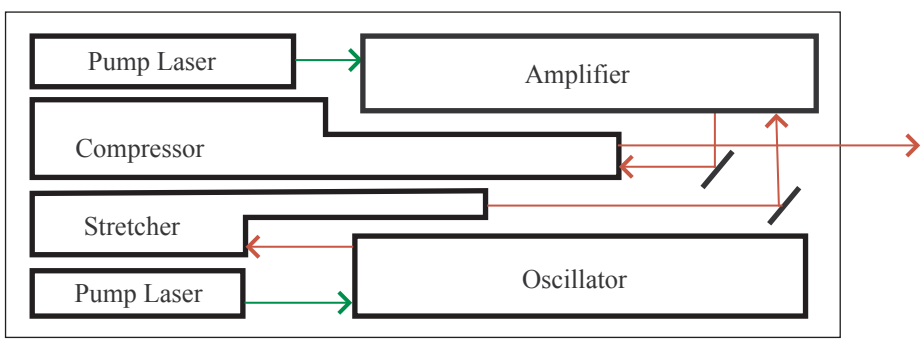

Figure 1: The laser system under study.

2 Theoretical Background

2.1 Coherence

A laser pulse is an electromagnetic wave packet. It is a propagating fluctuation of the electric and magnetic fields covering a finite volume. The property which makes a laser pulse differ from other forms of light is coherence.

Coherence describes the regularity of a field. It may be explained in terms of the spherical wavefronts given off by a point source and the net field generated by many such sources. The waves emitted at any point will never be perfectly regular in spacing, and will tend to show occasional jumps in phase due to processes such as collisions on the atomic scale. Because of this irregularity in wave phase, replicas of the wave will only superpose with a consistent phase relation over a limited relative delay. This statistical measure is known as the coherence time. Replicas delayed by more than the coherence time will form a noisy superposition in which some parts of the wave pair interfere destructively while other parts interfere constructively. In the coherence time, a field disturbance will propagate a corresponding distance known as the coherence length.

For a spatially distributed source, the different emitting points may show varying degrees of synchronisation. Well synchronised sources will generate wavefronts with a consistent phase, while poorly synchronised sources will form noisy wavefronts. Spatial coherence describes the degree of consistency of phase in the transverse dimension, and is important when superposing waves from different parts of the source.

A laser pulse is highly coherent in both its transverse and longitudinal directions, having a zero amplitude field preceeding and following a set of clearly defined wavefronts. In forming such a regular wave-packet pattern, it differs from the noisy traveling disturbances of other sources such as incandescent lamps.

2.2 Analytical Pulse Description

2.2.1 Single Variable Approximation

The laser pulse is conventionally described in terms of its electric field. A truly general description of this electric field requires many variables. Four dependences (three spatial and one temporal) are required to describe the magnitude of each of two field polarizations. Fortunately, the transverse spatial variation of the pulse is not of interest in many cases. Often in the remaining cases, it can be assumed identical or analogous to the transverse variation of a continuous laser beam.

Further, the two polarizations may usually be considered independent, so that the problem of analytical representation is simplified to describing a linearly polarized electric field magnitude along the propagation axis. This magnitude varies only as a function of the longitudinal spatial co-ordinate and time.

For the propagation in air of pulses whose dimensions are large with respect to wavelength, one may usually neglect absorption of the radiation, transverse spatial spreading of the beam (diffraction), and longitudinal spreading of the pulse (dispersion). Under these conditions, the pulse maintains its shape and on-axis magnitude as it propagates, so that the pulse description is similar in each of the two remaining variables. The variation of the pulse with time at any given point is then identical to the variation of the pulse with distance at some corresponding time, so that a pulse description in terms of either variable contains all information of interest. Conventionally, pulses are described as functions of time, with their longitudinal spatial behaviour implied by the velocity of light. The time frame chosen is usually local, following the pulse as it propagates so that the description changes only if the physical characteristics of the pulse change. Note that a typical plot of the pulse field vs. time is like an instantaneous snapshot of the pulse in the process of travelling from *right* to *left*.

2.2.2 Temporal Representation

The simplified (one dimensional) laser pulse is represented analytically by a time dependent complex function whose real part is proportional to the electric field magnitude at any time t:

$$E_t\{t\}e^{i\phi_t\{t\}}$$

This field description is possible because of the high temporal coherence of the laser pulse. A low coherence field would require a more complicated function capable of describing abrupt and random changes in field strength.

The modulus of the temporal function E_t (considered real) defines the pulse amplitude or amplitude envelope, while the (real) argument ϕ_t is referred to as the temporal phase. The amplitude is proportional to the square root of the instantaneous intensity of the pulse $I_t\{t\}$, which is the average power determined over one period of oscillation. The amplitude function limits the electric field, restricting the pulse to a finite time span and setting the local magnitude of field oscillations.

The temporal phase changes the real/imaginary component decomposition, allowing the field to oscillate between the limits defined by the amplitude envelope.

The rate of change of phase determines the instantaneous frequency of the pulse. Note that the sign of the argument is irrelevant due to the even symmetry of the real component; the choice of the above representation over its complex conjugate is just a matter of convention. Assume here that the phase angle increases with time.

Temporal phase is often expressed in terms of a power series expansion in which the centre point t_c is zero:

$$\phi_t\{t\} \;=\; \phi_{t0} + \phi_{t1}(t - t_c) + \phi_{t2}(t - t_c)^2 + \ldots$$

This series may diverge for very large or small values of the time, since the phase is only of significance where the amplitude function is non-zero.

The constant part of the temporal phase, which is independent of time, is referred to as the absolute temporal phase. It determines where the peaks of the oscillation are positioned with respect to the envelope. Quadratic and higher order contributions to temporal phase result in changes of instantaneous frequency over the duration of the pulse. A pulse of continuously increasing or decreasing instantaneous frequency is said to be "chirped".

The analytical pulse description may be understood intuitively in terms of phasor diagrams. The analytical function is a phasor of time dependent magnitude and angle (figure 2). The rate of change of the angle determines the instantaneous frequency of the pulse, while the real axis projection represents the electric field strength at any instant. The phasor has zero magnitude outside the time span corresponding to the pulse.

2.2.3 Spectral Representation

Like any function, the time dependent pulse description may be converted to various other forms. The Fourier transform (appendix A) may be used to convert the time-dependent pulse description into a function of frequency. The inverse mapping is also possible.

The pulse representation in frequency has a similar form to the temporal function. It is a complex function of a single variable:

$$E_\omega\{\omega\}e^{i\phi_\omega\{\omega\}} \;\equiv\; \mathcal{F}\big[E_t\{t\}e^{i\phi_t\{t\}}\big]$$

The single variable ω is the angular frequency, which is greater than conventional frequency by a scaling factor of 2π. The (real) modulus E_ω and (real) argument ϕ_ω of the spectral analytical function are known as the spectral amplitude and spectral phase respectively.

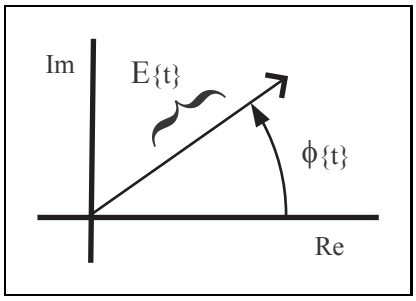

(a) Time Domain Phasor

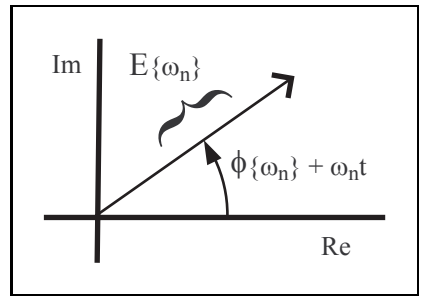

(b) Frequency Component Phasor

Figure 2: Phasor representations of a laser pulse. The electric field magnitude is represented by the real part of the time dependent pulse phasor (a). In the frequency representation, the time dependent phasor is decomposed into phasors of constant angular frequency (b).

The value of the spectral analytical function at a given frequency can be considered a projection onto a single basis vector in frequency space. This component corresponds to a sinusoidal wave in time. The wave will have a frequency as given, an amplitude equal to the spectral amplitude and a position on the time axis determined by the spectral phase. Contributions from all frequencies form the time domain pulse as a superposition of sinusoids. The square of the spectral amplitude is the spectral intensity $I_\omega\{\omega\}$, which expresses the power present in each frequency component. As with the temporal case, this simple functional description is possible because of the high temporal coherence of the laser pulse. A low coherence field would not maintain a consistent superposition of frequency components, and would therefore have a fluctuating spectral phase.

In terms of the phasor description, the spectral analytical function describes a set of phasors which each rotate in time at a unique, constant frequency, and which each produce a field component proportional to their real projection. The spectral amplitude determines the magnitude of each phasor, while the spectral phase describes the various phasor angles at time zero. The temporal analytical phasor corresponds to the net phasor obtained by summing the entire set of spectral phasors. Figure 2 illustrates a component phasor.

As with temporal phase, the spectral phase may be considered a series expansion of the single variable:

$$\phi_\omega\{\omega\} \;=\; \phi_{\omega 0} + \phi_{\omega 1}(\omega - \omega_c) + \phi_{\omega 2}(\omega - \omega_c)^2 + \ldots$$

The centre frequency ω_c is usually chosen so as to minimize the magnitude of the coefficients and optimize convergence of the series. Common choices include the frequency having the peak intensity, or the intensity-weighted average frequency

[2]. Once again this expansion need only be valid over the neighbourhood where the amplitude function is non-zero, so that it may diverge for extreme frequencies.

The constant spectral phase is that part of the spectral phase function independent of frequency. It may be identified with the absolute temporal phase, since adding a constant angle value to all frequency component phasors is the same as adding a constant angle value to their net phasor. The constant spectral phase thus determines the position of the field oscillations with respect to the pulse amplitude envelope.

Variations of the spectral phase which are linear with respect to frequency are of no significance to the pulse characteristics. By the Fourier shift theorem, pulses differing only in linear spectral phase occur at different times, but are otherwise physically identical:

$$E_\omega\{\omega\}e^{i(\phi_\omega\{\omega\}+\omega\tau)} \;=\; e^{i\omega\tau}\big(E_\omega\{\omega\}e^{i\phi_\omega\{\omega\}}\big) \;=\; \mathcal{F}\big[E_t\{t+\tau\}e^{i\phi_t\{t+\tau\}}\big]$$

The time shift τ of the pulse is simply the proportionality constant of the added linear spectral phase.

In the visual terms of the phasor diagram, a linear spectral phase shift is equivalent to advancing or retarding each component phasor by an amount proportional to its frequency. This does not alter the superpositions that occur, but does change the time at which they are observed.

2.2.4 Function Notation

Two conventions are often used to simplify analytical pulse description. First, the symbols representing a function are used to imply the type of function. Thus E, I and ϕ refer to amplitude, intensity and phase respectively in both temporal and spectral cases. Secondly, temporal and spectral functions are distinguished solely by the symbol used for their function argument. For example, $\phi\{t\}$ may be identified with the *temporal* phase $\phi_t\{t\}$, because the argument t is commonly used to denote time. This convention simplifies analytical expressions and reserves subscripts for other purposes.

2.3 Dispersion

To a good approximation temporal laser pulse characteristics remain constant as the pulse travels through air or vacuum. Significant modification may occur when the pulse travels through materials of higher density, or through certain systems of optical elements. These systems are best described in terms of the spectral pulse representation.

In traversing an optical system, a pulse frequency component will oscillate through some number of cycles determined by the optical path length, which is a function of both the geometrical path length s and the wavelength in the medium. If the optical path length varies with frequency, the system is said to be dispersive, and the superposition of components will be altered. The angular phase ϕ acquired by any component may be expressed in several forms using the angular wavenumber k, material wavelength λ, vacuum wavelength λ_0, refractive index n and the transit time t:

$$\phi \;=\; sk \;=\; \frac{2\pi s}{\lambda} \;=\; \frac{2\pi sn}{\lambda_0} = \;\omega t$$

Dispersion is often divided into two categories. Path length dispersion occurs when the geometrical path length varies with frequency. Material dispersion occurs when the refractive index varies with frequency (in which case the remaining medium-dependent quantities $k\{\omega\}$, $\lambda\{\omega\}$ and $t\{\omega\}$ above will show their frequency dependence accordingly). The sign convention used to describe dispersion is such that lower frequency components suffer the least time delay for a positively dispersive system. Positive dispersion will increase instantaneous pulse frequency vs. time.

2.3.1 Material Dispersion

For propagation in a dense medium, significant positive material dispersion is almost always present. This dispersion may be expressed as a function of any of the medium-dependent quantities listed above, but it is most common to express it in terms of the angular wavenumber:

$$\phi\{\omega\} \;=\; sk\{\omega\}$$

This form may be differentiated with respect to frequency to relate spectral phase to the "group velocity" v_g of the pulse:

$$\frac{d\phi}{d\omega} \;=\; s\frac{dk}{d\omega} \;=\; \frac{s}{v_g}$$

An additional differentiation relates spectral phase to "group velocity dispersion", which is the derivative of group velocity :

$$\frac{d^2\phi}{d\omega^2} \;=\; s\frac{d^2k}{d\omega^2} \;=\; \frac{s}{dv_g/d\omega}$$

The group velocity describes the speed at which the pulse propagates in an optical medium. This is not necessarily equal to the phase velocity v_p, which is

the speed of any given component:

$$v_p = \frac{\omega}{k} = \frac{s}{t}$$

In a medium of constant refractive index, all phase velocities and the group velocity are equal, and the pulse does not experience any physical change.

If the index varies linearly with frequency, different cycles of the sinusoids will come into alignment to form the pulse amplitude envelope. The shifting results in a change of absolute phase which is proportional to the propagation distance in the medium. The envelope (and thus the pulse itself) will propagate at the group velocity without changing shape, and will appear to move with respect to the net field oscillation. When the pulse exits the medium and returns to vacuum or air, the group and phase velocities will become equal, and the absolute phase of the pulse will remain fixed.

Any non-zero value of the group velocity dispersion will introduce non-linear spectral phase, altering both the pulse envelope and phase.

2.3.2 Path Length Dispersion

Path length dispersive systems modify pulses by physically separating their component frequencies and recombining them after they have traversed a frequency dependent geometrical path. Two such systems are of interest here, the grating pair and the prism pair.

A reflective optical grating is a surface whose reflective phase and or amplitude varies periodically in one dimension. The most common type found in ultrafast optics is a mechanically ruled gold surface, where the grooves are spaced on the order of a micron ($10^{-6}m$). When excited by part of an incoming wavefront, each groove will act as a line source, re-emitting the incident wave in all directions. In addition to the reflected beam observed for a regular mirror, the grating will produce beams in any direction for which which adjacent groove outputs are an integer number of cycles out of phase. The number of cycles describes the order of the reflection. For any non-zero order, the exact reflected direction is wavelength dependent. The various frequency components of the incident beam therefore leave the grating at different angles, fanning out in the plane perpendicular to the grooves.

A grating pair arranged with their faces and grooves oriented parallel to one another may be used to induce a controlled amount of dispersion. First order reflections are applied as shown in figure 3. The incident pulse is separated into frequency components after the first incidence. The second grating then reflects

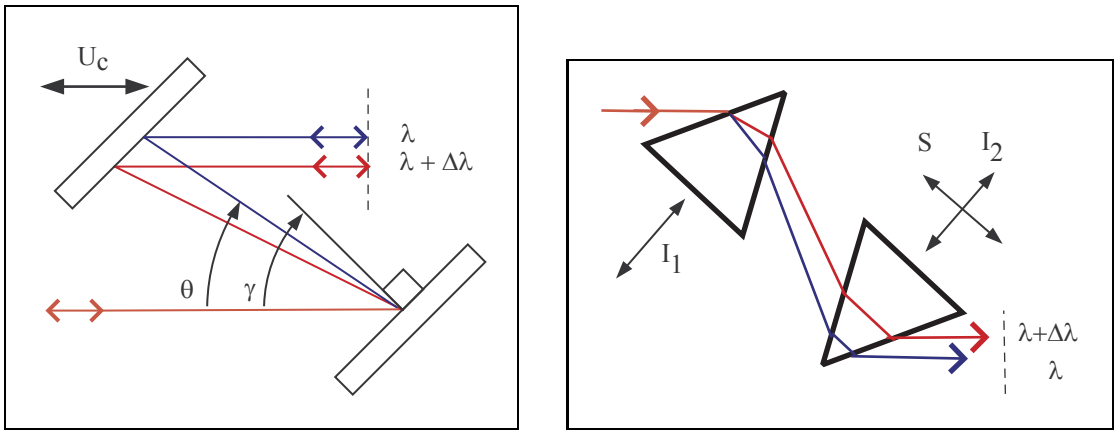

(a) Grating Pair

(b) Prism Pair

Figure 3: Two dispersive optical systems. The grating pair (a) produces path length dispersion only, while the prism pair (b) introduces both path length and material dispersion.

---

each component at an angle complimentary to its initial reflection. The components leave the second grating traveling parallel to one another, though slightly offset. The path length through the system is frequency dependent, so that the exit pulse will have experienced dispersion. If the lateral spatial offset of the components is large with respect to the beam diameter, it may be corrected by reflecting the beam back through the grating pair in a different vertical plane. This second pass will double the net dispersion. The grating pair dispersion has been determined analytically [17] and may be calculated given the line spacing d, the distance between grating face planes G and the angle of incidence γ:

$$\frac{dt}{d\nu} = \frac{-2cG}{d^2\left(1-((c/\nu d)-sin\{\gamma\})^2\right)^{\frac{3}{2}}}$$

Here ν is frequency, c is the speed of light, and the double pass geometry is accounted for by the factor 2.

Prisms isolate pulse frequency components in much the same way as a grating. The variation of refractive index with frequency in the prism material causes slightly different exit angles for each component of an incident pulse. A prism pair (figure 3) may form a system similar to the grating pair, in which an initial prism sends frequency components in different directions and a second prism realigns them [10]. Again, a double pass may be used for perfect recollimation. In the case of the prism pair, both material and path length dispersion are present.

2.4 Spectral Phase Measurement

2.4.1 Spectral Interference

When two waves are mixed they interfere. The net resulting wave depends on both the amplitude and relative phase of the component waves. This is a familiar effect in the temporal (or spatial) domain.

Interference may also be observed in the frequency domain. The power spectrum observed for two combined pulses is not simply the sum of the individual pulse spectra. Instead, the observed spectral intensity will vary as the square of the net spectral amplitude, which is phase dependent. Spectral intensity maxima will be observed where the individual pulse frequency components have reinforcing phase. Minima will be observed at frequencies where the individual pulse components have canceling phase.

Analytically, superposing two coherent pulses (in the linear optical case) is equivalent to adding their phasors. In the spectral basis, each pulse is viewed as the sum of a set of constant magnitude, constant frequency phasors, and component phasors of like frequency may be added to form the spectral representation of the net field.

The general expression for the two pulse spectrum $D\{\omega\}$ has an analytical form which is common in optics and other wave phenomena,

$$D\{\omega\} \;=\; E_1\{\omega\}^2 + E_2\{\omega\}^2 + E_1\{\omega\}E_2\{\omega\}\cos\Big[\pm\,\phi_1\{\omega\} \mp \phi_2\{\omega\}\Big]$$

The first two terms are simply the power spectra of the individual pulses, while the third term describes interference effects. Note that the interference depends only on the absolute value of the phase difference, so that the sign of the phase difference (or equivalently, the labeling of the pulses) is irrelevant.

2.4.2 Fourier Transform Spectral Interferometry

Under appropriate conditions, the spectral phase difference between two pulses may be recovered from their combined spectrum by the method of Fourier Transform Spectral Interferometry (FTSI) [7]. The procedure is based on frequency filtering, a common signal processing technique.

If two pulses are well separated by some time delay τ, their spectral phase difference will be approximately linear in frequency. The interference term of the combined spectrum will oscillate through several periods over the bandwidth of the pulse, so that the spectrum forms a fringe pattern (figure 4) analogous to the spatial intensity pattern observed in Young's 2 slit experiment [13]. Several pulse

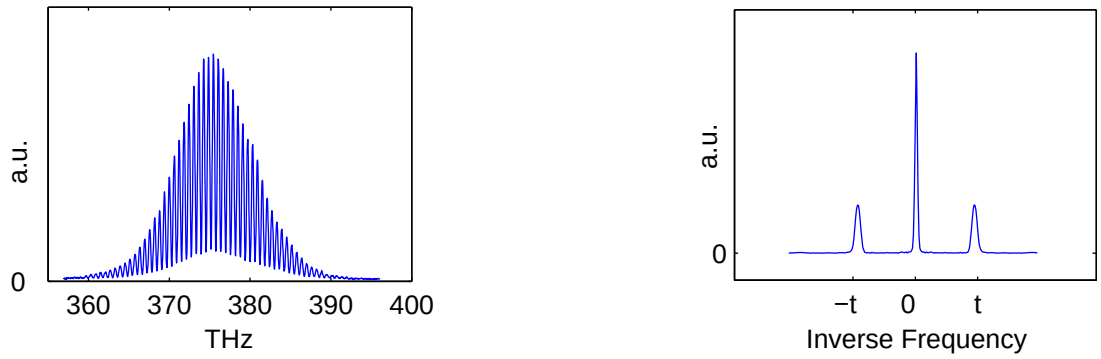

(a) Pulse Pair Spectrum

(b) Transform of Spectrum

Figure 4: A measured pulse pair spectrum (a) exhibits fringes caused by spectral interference. The nominal fringe spacing is the inverse of the pulse time separation t. A Fourier transform maps the spectrum to three distinct regions of the transform abscissa (b). All spectral phase information is contained in each of the two lobes on either side of the central maximum.

---

pairs may be collected to form an averaged interference spectrum, so long as the pair spacing is consistent, and the separation between consecutive pairs is large. The returned phase properties are understood to be an average.

The oscillation of the spectrum contains the desired phase information. To isolate the phase dependent part of the data, the spectrum is Fourier transformed. This transformation should not be confused with another common process, namely inverse transforming the spectral pulse representation; note that it is a *forward* Fourier transform which is applied to the spectral *intensity*. The transformed function (figure 4) is a representation of the spectrum in a reciprocal space, which has units of time. The usual roles of time and frequency are thus reversed.

Components of the spectrum are mapped to regions of the time axis equal to the inverse of their spectral width. Broad features which vary slowly with spectrometer frequency are mapped to a central ("D.C.") region around the zero time point. Narrow spectral features are mapped to more extreme times, both positive and negative. If the spectral intensity changes due to phase difference are rapid relative to the intensity changes of the individual pulse spectra, then the oscillation term will map to a unique region of the time axis. Specifically, it will form a distinct pair of ("A.C.") sidebands. A filter may be applied to remove the central lobe, isolating the interference data.

The transformed spectrum representation is a complex function of time. Because the measured spectral intensity is a real quantity, the positive and negative sides of the transform are complex conjugates. Taken as a pair, the side lobes

may be inverse transformed to form a real data set which is the spectral phase oscillation. If only a single lobe is inverse transformed, a complex magnitude and phase representation of the oscillation term results. The real part of this representation is equal to half the observed value of the oscillation term. The complex angle is equal to the individual pulse phase difference. By convention, assume that the positive time sideband is retained, so that the angle of the complex spectrum representation increases with spectrometer frequency.

The linear component of the recovered spectral phase difference arises from the time separation of the pulses, as described in a common reference frame. It is often preferable to describe the pulses by similar functions, such that they would occur simultaneously if viewed in a common time frame. This choice of functional description is practical for comparing the pulses and results in identical linear spectral phase components. This description may be applied to the phase difference measurement by writing the linear component of the phase difference explicitly, outside the pulse functions. One assumes that the functional descriptions have been time shifted to a common frame by adding a linear spectral phase to one of the pulses. The sign chosen for the linear term depends on previously assumed conventions, specifically pulse labeling, pulse order in time, the rotation direction of the complex pulse and spectrum representations and the sign of the Fourier transform. Assuming that the time delay τ is measured with respect to the frame of pulse 1, and applying all previous sign conventions, the signed phase difference may be written:

$$\phi_1\{\omega\} - \phi_2\{\omega\} \quad \rightarrow \quad \phi_1\{\omega\} - \phi_2\{\omega\} + \omega\tau$$

The linear term of the measured phase difference $\omega\tau$ may be removed numerically or by a calibration measurement. The remaining data is the non-linear spectral phase difference Φ between the two pulses:

$$\Phi\{\omega\} \;\equiv\; \phi_1\{\omega\} - \phi_2\{\omega\}$$

The traditional application of FTSI is measurement of the phase transfer function of an optical element. The transfer function describes the non-linear spectral phase change experienced by a pulse due to dispersion. If the two pulses used to generate the spectrum are replicas created in an interferometer, then the non-linear phase difference between them will be due to dispersion differences in the arms of the interferometer. A phase difference measurement may therefore be used to determine the spectral phase transfer function of the beamsplitter substrate, or the beamsplitter substrate plus or minus the transfer function of any optical

element placed in an arm of the interferometer. A calibration measurement of the substrate dispersion may be used to algebraically isolate the dispersion due to the optic under test.

In addition to measuring the *change* in non-linear spectral phase imparted by an optic, FTSI may be applied to the problem of measuring the pulse spectral phase itself. In this case, the unknown pulse under test (labeled u) is mixed with a reference pulse (labeled r) whose spectral phase is known. The reference phase is simply combined algebraically with the measured phase difference to return the unknown phase:

$$\begin{aligned} \Phi\{\omega\} &= \phi_r\{\omega\} - \phi_u\{\omega\} \\ \phi_u\{\omega\} &= \phi_r\{\omega\} - \Phi\{\omega\} \end{aligned}$$

Unfortunately, this characterization method cannot be used to make definite pulse spectral phase measurements because the reference pulse phase must either be measured by some other means or assumed. In the former case spectral interferometry is dependent on other methods, while in the latter case the measurement is not definite.

2.4.3 Self-Referencing Spectral Interferometry

Self-referencing spectral interferometry allows limited but definite measurement of the spectral phase of an unknown pulse waveform (labeled u) without the use of a second, characterized pulse. The method combines two replicas of the unknown pulse (labeled 1 and 2). One of the pulse replicas (say 2) is frequency modulated, so that its spectrum is shifted by the modulating frequency Ω:

$$E_2\{\omega\}e^{i\phi_2\{\omega\}} = E_u\{\omega - \Omega\}e^{i\phi_u\{\omega-\Omega\}}$$

The applied frequency shift is chosen to be a relatively small fraction of the total pulse bandwidth. The pulse replicas now have spectral functions which are displaced in frequency but otherwise identical, and they are said to be spectrally sheared. As usual, FTSI may be applied to the combined pulse spectrum to return the non-linear spectral phase difference between the two pulse replicas vs. spectrometer frequency. Relating the two replicas back to the unknown pulse however, shows that the measured phase difference is actually a comparison of the unknown pulse spectral phase value at adjacent frequencies:

$$D\{\omega\} \Rightarrow \phi_1\{\omega\} - \phi_2\{\omega\} = \phi_u\{\omega\} - \phi_u\{\omega - \Omega\}$$

The phase measurement compares the unknown phase at each frequency to its value at a neighbouring frequency differing by the spectral shear. When combined with the spectral shear value, this measured phase difference may be related to the first derivative of the unknown phase:

$$\frac{\phi_u\{\omega\} - \phi_u\{\omega - \Omega\}}{\Omega} \simeq \frac{d\phi_u\{\omega\}}{d\omega}$$

At this point, a constant may be added to the derivative expression which is equivalent to a linear term in the unknown phase. By the Fourier shift theorem, this constant simply determines the time frame chosen to describe the pulse. If we assume that the temporal pulse is described in a local time frame such that it peaks around the zero point, then the unknown spectral phase function must have zero slope at the intensity weighted centre frequency. The constant may be set accordingly. The choice of constant does not effect the physical characteristics determined for the unknown pulse.

A continuous first derivative could be integrated to determine the non-linear spectral phase. The discrete measurement may be summed instead. The partial sum of the discrete derivative at any point in the series is equivalent to the non-linear phase at that point, aside from some constant value. This second unknown constant is analogous to a constant of integration and must be set arbitrarily. This leaves the absolute phase of the reconstructed pulse unknown. This limitation applies to the FROG technique as well.

When combined with a trivial spectral amplitude measurement and inverse transformed, the measured spectral phase function determines the instantaneous intensity and frequency of the unknown pulse. Only the absolute value of the phase remains undetermined. Note that the discrete nature of the measured spectral function does not limit the accuracy of the pulse reconstruction. By the Whittaker-Shannon sampling theorem (appendix A), the time domain pulse may be *exactly* described by a phase function with a sampling interval exceeding the inverse of the pulse duration. It is the spectral *width* which limits the quality of the pulse reconstruction by setting the highest pulse frequency component reproduced.

2.4.4 The SPIDER Method

Self-referencing spectral interferometry as described in the previous section cannot currently be realized for pulses of optical frequency. This is strictly a technology limitation, in that no frequency modulator is known which can generate the required spectral shear. Spectral Phase Interferometry for Direct Electric Field Reconstruction (SPIDER) [5] adds one additional step to pulse modification and

data processing to allow self-referencing spectral interferometric measurements of laser pulses. In SPIDER, non-linear frequency mixing is used to produce a spectral shear.

Certain combinations of propagation medium and high field strength result in "non-linear" optical conditions. In the non-linear regime, one observes effects not predicted under the usual approximations of traditional optics. One such effect is frequency mixing, in which two coincident waves are combined to form a third wave whose frequency is an integer combination of the incident wave frequencies. If one of the incident waves is a pulse and the other a continuous (single frequency) wave, then the output mixed wave is simply a frequency-shifted version of the original pulse. Both amplitude and phase functions are shifted in frequency without a distortion in shape.

Unfortunately, non-linear frequency mixing cannot be used to directly generate the required spectral shear, because it is not currently possible to mix frequency pairs as dissimilar as an optical frequency ($\approx 500\ THz$) and the required shear value ($\approx 1\ THz$). Instead, SPIDER frequency modulates *both* pulse replicas, shifting them by unique optical frequencies:

$$\begin{aligned} \phi_1\{\omega\} &= \phi_u\{\omega - \Omega_1\} \\ \phi_2\{\omega\} &= \phi_u\{\omega - \Omega_2\} \end{aligned}$$

The slight difference in frequency shifts results in a spectral shear Ω between the shifted pulses:

$$\Omega \equiv \Omega_2 - \Omega_1$$

This shifting process is illustrated schematically in figure 5.

The phase shifts may be determined and used to relate the measured spectral phase difference to the unknown pulse phase.

$$\begin{aligned} D\{\omega\} \quad \Rightarrow \quad \phi_u\{\omega - \Omega_1\} - \phi_u\{\omega - \Omega_2\} &= \phi_u\{(\omega - \Omega_1)\} - \phi_u\{(\omega - \Omega_1) - \Omega\} \\ &= \phi_u\{\omega'\} - \phi_u\{\omega' - \Omega\} \end{aligned}$$

Concatenation of the measured values then proceeds exactly as in the previous section, allowing the measurement of instantaneous intensity and frequency of the unknown pulse. Note that several spectral phase data sets may be reconstructed from a given interference spectrum by varying the starting point used for concatenation.

As a matter of convenience, the two modulation frequency components are usually derived directly from the unknown pulse. The optical setup is illustrated in

figure 6. A third replica of the pulse is sent through a grating system and stretched in time. The segments of the stretched pulse which arrive at the sum frequency generator coincident with the interferometer pulses are of different frequency and approximately monochromatic.

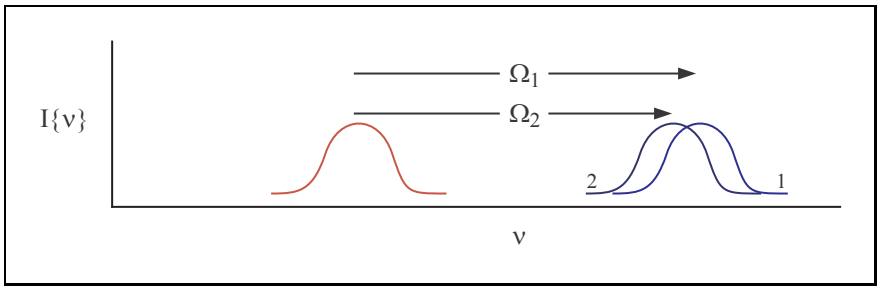

Figure 5: Frequency shifts used in the SPIDER method. The particular case illustrated here is the more common frequency upconversion.

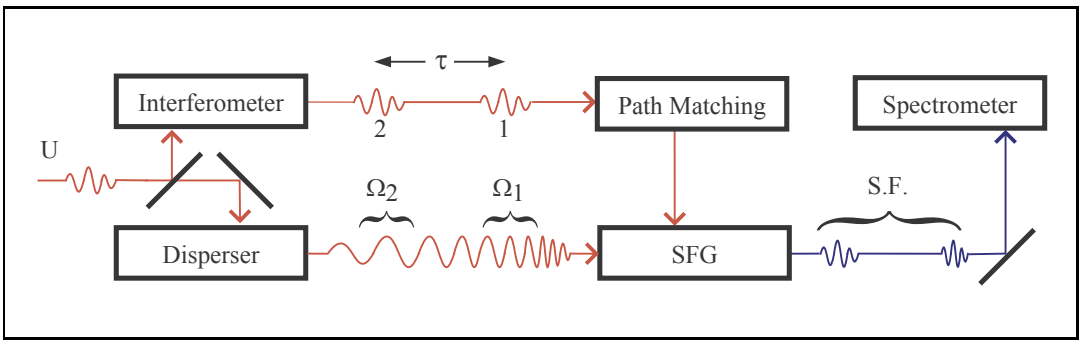

Figure 6: Temporal schematic of the frequency shifting process commonly used in SPIDER.

3 The SPIDER System

3.1 Optical Apparatus

The optical part of our SPIDER apparatus is built on a standard $24'' \times 12''$ optical board housed in a plastic box. Functionally, it may be divided into four smaller systems: a Michelson interferometer, a dispersive grating system, a path length adjuster and a sum frequency generator. This is similar to the schematic shown in figure 6. Any remaining optical components are trivial items such as mirrors or beamsplitters.

The horizontally polarized input pulse is divided into two replicas at an initial glass beamsplitter. The transmitted pulse replica is directed to the grating system while the reflected replica is sent into the interferometer. This choice of routing is deliberate. It divides the input beam power in a ratio of 60% transmitted to 40% reflected, which partially balances the different losses of the stretcher and interferometer. Dispersion errors due to the beamsplitter substrate are also avoided by sending the reflected replica to the interferometer.

In the stretcher system, the input pulse is dispersed by two gold surface gratings ruled at $d^{-1} = 1200\ lines/mm$ each. The arrangement is as shown in figure 3. After the initial pass, a periscope displaces the beam vertically and sends it back through the grating pair at a slightly greater height. The dispersion of the stretcher for frequencies within the pulse bandwidth is approximately 2×10^6, as calculated using the measured incidence angle $\gamma = 37°$ and the distance between grating faces $G = 10\ cm$.

The pulse sent into the interferometer is divided and recombined at a single coated glass beamsplitter as shown in figure 7. One arm of the interferometer is adjustable via a micrometer and translation stage, while the other is fixed. The adjustable arm is set such that its total path length delay is *greater* than that of the fixed arm by roughly $2\ ps$, with a more accurate value determined by spectral measurement. The interferometer is the part of the apparatus which is most

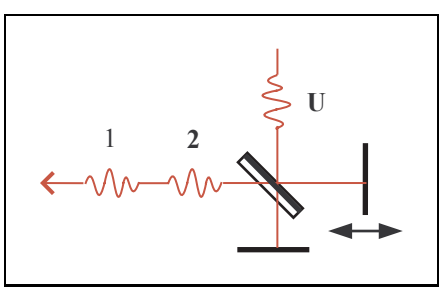

Figure 7: The Michelson interferometer used to replicate the unknown pulse.

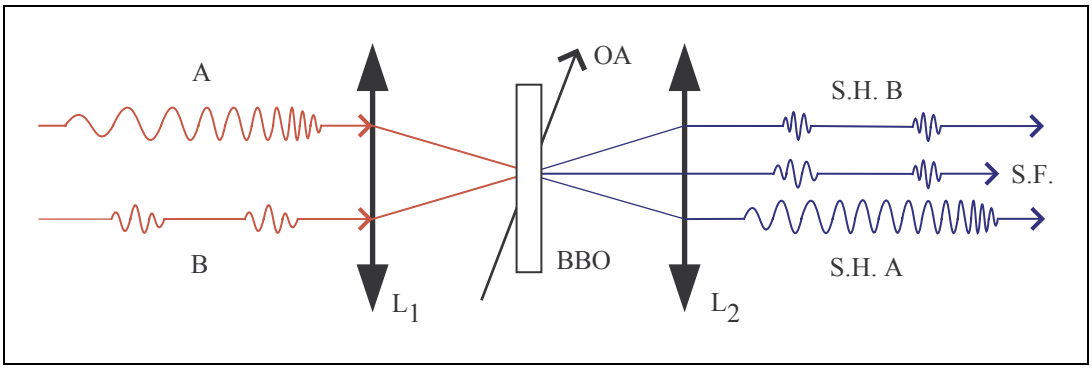

Figure 8: The type I non-collinear sum frequency generator. Second harmonic (SH) and sum frequency (SF) signals are indicated for inputs A and B. The optic axis (OA) of the BBO crystal falls in the plane of the interacting beams as shown.

---

sensitive to vibration, so for improved stability several points of the beamsplitter are mounted adhesively to a thick aluminum ring, which is then bolted to a solid block of aluminum. The adjustable arm is built by placing a small translation stage at the top of another solid aluminum block. The substrate of the beamsplitter is approximately 1 mm thick and oriented to face the fixed arm. Interference spectra generated with the unamplified pulses direct from the oscillator were observed to be very consistent in fringe spacing, indicating that the interferometer is quite stable with respect to mechanical vibration.

The interferometer output is sent through a variable path length and then aligned below the stretcher output. The path length is adjusted so that the pulse pair from the interferometer is syncronised with the much longer pulse from the grating stretcher. Entering the sum frequency generator, the beams are parallel but vertically offset, with a separation of about 1 cm.

The sum frequency generator is based on a birefringent, non-linear crystal. Non-linear optical interaction in the crystal generates three new outputs in the 400 nm (blue) region of the spectrum by combining pairs of infrared input photons [3]. Two of the outputs are due to second harmonic generation, where both combined photons come from the same input beam. This forms a frequency-doubled replica of each input. The third output is the sum frequency combination of the two input signals, generated by combining one photon from each input beam. The direction of each output is determined by conservation of photon momentum. The specific arrangement used in this case is referred to as a non-collinear type I geometry and is shown in figure 8.

Entering the sum-frequency generator, the parallel input beams are focused by

a common lens and intersect over a region centered at the 20 cm focal length. Focusing increases the field intensities to create non-linear optical conditions. A 100 μm thick beta barium borate (BBO) crystal window is placed in the region of overlap. Each molecule of the crystal may separately emit any of the three possible outputs, and these various point sources must be syncronised or "phase-matched" over the bulk of the crystal to produce a significant net output. This is accomplished by orienting the crystal so that its "optic axis" falls within the plane formed by the intersecting input beams. The input beams are then of "ordinary" polarization while the vertically polarized outputs are of "extraordinary" polarization. Rotation of the crystal in the plane of the input beams alters the refractive index experienced by the output beam and may be used to partially synchronize the molecular outputs. When properly oriented, the directional, polarization and wavelength dependences of the crystal's refractive index partially cancel. A second lens following the crystal is used to recollimate the diverging output beams and direct them to a periscope. In the periscope the pulse polarization is rotated to the horizontal before the pulses are sent to the spectrometer. An aperture preceeding the lens may be used to single out one of the three crystal outputs.

The sum frequency beam is the frequency shifted pulse replica pair used for SPIDER measurement. The stretched pulse is chirped to a length of about 20 ps, so that its instantaneous frequency varies continuously from one end of the pulse to the other. There should be little change in instantaneous frequency over the ($\approx$ 100 fs) duration of a single pulse, but a significant change over the 2 ps time interval between pulses. Each interferometer pulse replica arrives at the crystal coincident with a different, approximately monochromatic frequency of the stretched pulse, generating a spectral shear of approximately 1.7 THz, with the first pulse being shifted by about 374 THz.

3.2 Spectrometer

Pulse spectra are measured by a spectrometer of the Czerny-Turner type, specifically Oriel model MS127i, shown in figure 9. The pulse to be measured is directed onto a 10 μm entrance slit where it diffracts, acting as a line source that re-emits the incident wave along a cylindrical wavefront. This arrangement reduces sensitivity to input beam width, position and direction. Curved mirrors collimate the cylindrical wave and direct it to a grating which spreads the various frequency components. The diverging component frequencies are redirected by additional curved mirrors and sent to an array of detectors. The frequency band and resolution of the spectrometer may be controlled by the orientation and line density of

the interchangeable grating. A 1200 $line/mm$ ruled gold grating is used to measure the infrared fundamental pulse spectrum, while a 2400 $line/mm$ holographic grating is used to measure the blue frequency-shifted spectra.

The detector array is an Oriel LineSpec linear charge-coupled device (CCD). The single line of detectors is positioned horizontally, so that each element receives light from a different range of frequency components. An electrical charge accumulates at each element in rough proportion to the incident light intensity, so that the relative charges form a power spectrum measurement of the input. During the readout cycle, each pixel charge is converted to a voltage sequentially, forming an analog voltage signal. This analog output is digitized at the computer by a GageScope oscilloscope board. Array readout syncronisation and pixel readout syncronisation signals aid in timing the data conversion.

The combination of grating and detector produces a wavelength resolution of approximately $0.07nm$ for the infrared setting and $0.03nm$ for the blue setting. Wavelength calibration is accomplished using gaseous electrical discharge lamps. The infrared range is calibrated using the Argon (Ar) emission spectrum, while the blue range is calibrated using Helium (He) and Hydrogen (H) emission spectra. Several lines in the immediate neighbourhood of the pulse spectrum are assigned a pixel number value according to their intensity weighted centre point. Standard tables [18] are used to assign a wavelength value to each line as well. The pixel/wavelength data pairs are then fitted to a fourth order polynomial which expresses wavelength as a function of pixel number, allowing the interpolation of wavelength values for all pixels. To reduce error, the polynomial fit is centered at a pixel number near the region of interest, and calibration lines far from the region of interest are not included.

Background calibration of the spectrometer is also necessary. Charge sources such as stray light and thermally generated carriers produce dark current at each

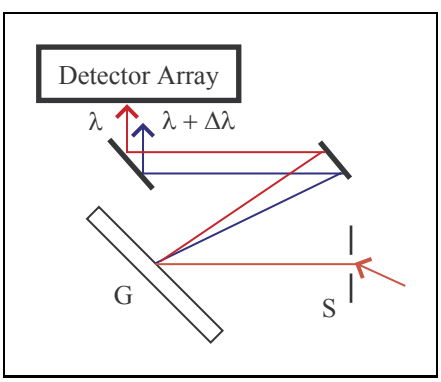

Figure 9: The Czerny-Turner spectrometer. The entrance slit (S), grating (G) and detector array are indicated. Some reflective and focusing optics are omitted for clarity.

pixel independent of the input light. An averaged value of the background level is determined for each pixel and removed from the initially read signal.

3.3 Data Processing

SPIDER data processing is performed on an IBM compatible personal computer (PC) running LabView programming software under the Windows 2000 operating system. The processing is approximately instantaneous or "real-time", finishing well within one second of pulse generation.

Measured interference spectra are Fourier transformed directly using a built in LabView algorithm which accepts a one dimensional set of data and assumes a regular sampling rate. This is a simpler alternative to converting the measured wavelength spectrum to an evenly sampled function of frequency. Additional phase appears when the sampled spectrum is viewed as a function of pixel number, but the Fourier transform filtering process may still be applied because the pixel to frequency relation is approximately linear over the pulse bandwidth. The additional phase disappears when the pixel values are converted to frequency. Execution speed is improved by the use of fast Fourier transform (FFT) algorithms. The frequency filter applied in the transformed space is a "hyper-Gaussian" which is a Gaussian curve raised to a power. Exponents of about 4 work well, providing a wide, flat and apodized filter. To further increase execution speed, the filter function is evaluated only during an initialization loop and stored for use in future iterations. This improves execution time by reducing the calculation of exponential functions.

The observed fringe pattern phase is shifted down in frequency by subtracting the shift value applied to the highest frequency pulse replica. This is determined by comparing the fundamental and shifted spectra.

The delay dependent component of the fringe pattern phase is determined by a separate measurement of the interferometer output second harmonic. The pulses are essentially identical in this case so that the phase returned from the interference spectrum is the desired linear phase:

$$\begin{aligned} D\{\omega\} \quad \Rightarrow \quad \phi_1\{\omega\} - \phi_2\{\omega\} &= 2\phi_u\{\omega/2\} - 2\phi_u\{\omega/2\} + \omega\tau \\ &= \omega\tau \end{aligned}$$

The directly measured data may be used, or the time delay may be determined from a linear fit and used to generate the background phase.

The derivative phase at the desired centre frequency is set to zero by subtracting a common constant value from the entire data set.

The spectral shear may be determined by calculating the dispersion of the grating pair at the fixed pulse frequency and combining this with the pulse time-separation as measured during calibration. Alternatively, the shear may be estimated by comparing the intensity weighted centre frequencies of the individual shifted pulses. A third shear estimating method which compares theoretical and observed fringe pattern modulation amplitude did not produce consistent results.

Integration of the measured phase derivative proceeds by concatenating values separated by the spectral shear. A linear interpolation between nearest neighbouring sample points is used to evaluate the phase derivative at the required points. To increase execution speed, the point pair used for each interpolation and the relative weighting of each point is calculated only during an initialization loop. It is then stored for future loop iterations. One or several spectral phase data sets may be constructed by shifting the starting point of the concatenation. All data sets have their absolute phase set to zero at a common centre frequency, which allows their phases to be averaged in the frequency domain. The measured spectral phase values may optionally be fit with a polynomial for analysis or smoothing.

An optional transformation to the time domain may be performed by use of an inverse Fourier transform. Spectral amplitude and phase data sets are constructed at sampling rates similar to the spacing of the measured spectral amplitude data. Values are interpolated over the measured spectral ranges and set to zero elsewhere. Extra zero values added to the extreme frequency ends of the spectrum allow the use of an inverse FFT algorithm and increase the sampling rate for the plotted temporal pulse function, though they do not improve the actual resolution of the data.

4 Measured Data

Data captured with the SPIDER system is presented in this section. The data is analyzed in section 5. By convention, spectral phase curves are presented in radian units as a function of proper (as opposed to angular) frequency. Phase fitting coefficients are defined by the relation

$$\phi\{\nu\} \;=\; \phi_0 + \phi_1(\nu - \nu_c) + \phi_2(\nu - \nu_c)^2 + \ldots$$

where the phase is in radian units, ν is the proper frequency, and ν_c is the centre frequency. The SPIDER software used femtosecond time units and nanometer length units for convenience, and as a result the phase expansions assume frequencies in PetaHertz. Conversion to other units or coefficient conventions is straightforward.

4.1 SPIDER Verification

The following data was collected to verify the correct operation of the SPIDER system.

4.1.1 Interferometer Imbalance and Fringe Sampling Rate

Dispersion differences between the arms of the interferometer are a source of error. In traditional FTSI it is usual to measure this phase difference and correct the observed data, or to balance the arms using a glass blank. The blank is identical to the beamsplitter substrate and is placed in the arm opposite the substrate to make the glass path lengths identical. For SPIDER, dispersion throughout the *entire* optical apparatus is a source of error. In our SPIDER apparatus little glass was used outside the unbalanced interferometer, and a glass blank was not available to match the beamsplitter substrate. It is therefore expected that the interferometer imbalance is the greatest source of dispersion error. The imbalance of the interferometer is measured here and used to estimate that error.

The interferometer was used to create fringe spectra of different pulse time separation. Seed pulses from the laser oscillator were used as input, and the output was sent directly into the spectrometer. The use of the fundamental frequency was convenient and resulted in more accurate data, as the seed pulse amplitude function was very smooth and unlikely to cause noise in the observed oscillation. Several multiple-shot fringe spectra were recorded, including figure 4, pg. 14. The single pass substrate dispersion determined at three different time separations is shown in figure 10.

The imbalance data may also be used to check for error due to spectrometer resolution. Fringe definition requires two data points at minimum, and improves with additional sampling. The sample spacing of the spectrometer is fixed for either scale, so that the fringe width determines the number of samples per fringe. The samples here have different fringe width, and therefore different fringe resolution.

4.1.2 Spectral Shift Adjustment

The spectral phase returned from a SPIDER measurement should be independent of the particular frequency shifts applied to the individual pulse replicas, as long as the values fall within reasonable limits and are accurately determined by calibration. The phase data sets presented in figure 11 are used to verify correct operation and calibration. Two data sets differ in the total amount of spectral shift. Two data sets have different spectral shear. The parameter changes were accomplished by varying the interferometer delay and the grating/interferometer path matching delay. The SPIDER was recalibrated at each new setting.

4.1.3 Shear Calibration

The spectral shear may be determined by comparison of upconverted pulse spectra, or by dividing the pulse time separation by the SPIDER stretcher dispersion. For the first method, alternating arms of the interferometer are blocked off, and the shear is taken as the difference in intensity weighted centre frequency of the upconverted pulses. For the second method, the stretcher dispersion is calculated from measured stretcher geometry [17] and the time separation is determined from the linear component of the second-harmonic interference spectrum. Figure 12 shows phase curves constructed by the different methods.

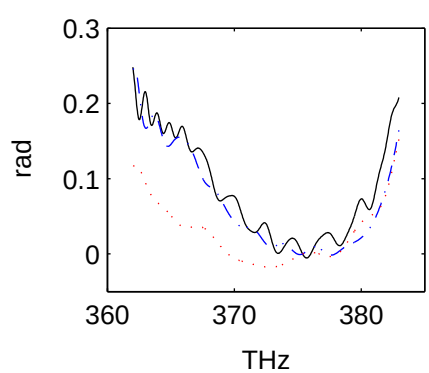

Figure 10: Single pass dispersion of the Michelson interferometer beamsplitter substrate, measured at the fundamental. The pulse time separations used were 3.75 ps (solid line), 1.98 ps (dash-dot line), and 0.64 ps (dotted line).

4.1.4 Shear Approximation Error

Each method of spectral shear measurement described above leaves the SPIDER susceptible to a small error which varies with the chirp of the unknown pulse.

When the shear is determined using the dispersion of the SPIDER stretcher, the unknown pulse is assumed to be of minimum duration for its bandwidth ("transform limited"). If the unknown pulse enters the SPIDER stretcher with an initial chirp (of either sign), this assumption is only approximate. When shift and shear are determined by comparison of upconverted spectra, values will be inexact whenever the chirp of the unknown pulse changes from the value used for the calibration. The data presented here was recorded in order to estimate the significance of these errors.

The chirp of the unknown pulse was adjusted by changing the grating spacing in the laser system compressor. The spectral phase was reconstructed over a range of compressor settings, using both fixed calibration values determined at the compressor centre and values measured separately for each compressor position. The values were measured by comparing upconverted spectra. Figure 13 shows the phase coefficient values as a function of compressor setting for each case.

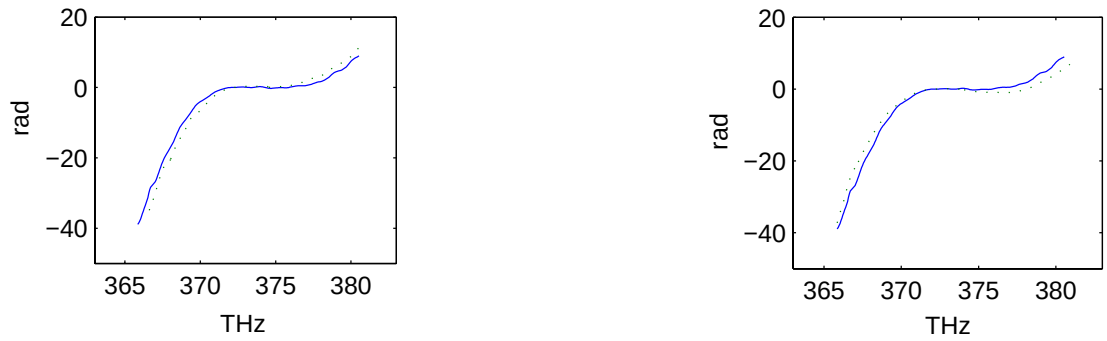

(a) Shift Test

(b) Shear Test

Figure 11: Spectral shifting test. Two phase curves were measured using identical shear values and different first pulse shifts (a) of 373.9 THz (solid line) and 371.1 THz (dotted line). Two phase curves were measured using identical first pulse shifts and different shear values (b) of 1.5 THz (solid line) and 0.9 THz (dotted line).

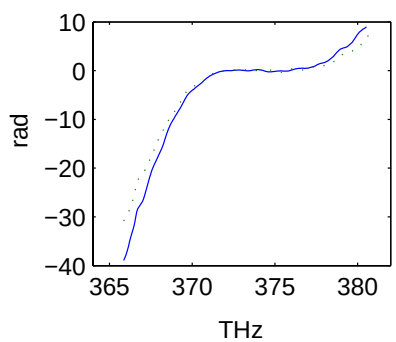

Figure 12: Phase traces comparing different reconstruction methods. The reconstructions used shear values determined by SPIDER timing and grating geometry (solid line) and power spectrum comparison (dashed line).

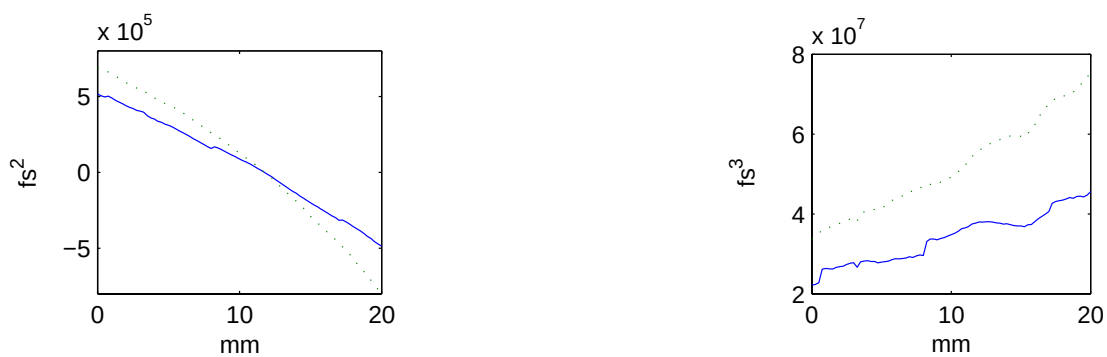

(a) Quadratic Phase

(b) Cubic Phase

Figure 13: Quadratic (a) and cubic (b) spectral phase coefficient values vs. laser compressor setting for the cases of corrected (solid line) and fixed value (dashed line) calibration.

4.2 Pulse Variation

This data describes the variation in the reconstructed spectral phase. Amplified pulse systems tend to produce somewhat inconsistent pulses from one shot to the next, so that the measured variation is expected to be due more to the physical performance of the laser than the precision limits of the SPIDER system.

One thousand phase measurements were made, each corresponding to a single pulse. This required the use of an optical chopper with an appropriate duty cycle, as the minimum exposure time of the CCD was about six times longer than the pulse repeat period. To identify short term drift in system behaviour, the spectra were recorded at evenly spaced intervals over a period of one hour. The entire laser system and detector were given more than an hour to reach a steady operating state before recording began.

A third order polynomial was fit to each returned phase. The phase coefficients and their statistical parameters are presented in figure 14 and table 1.

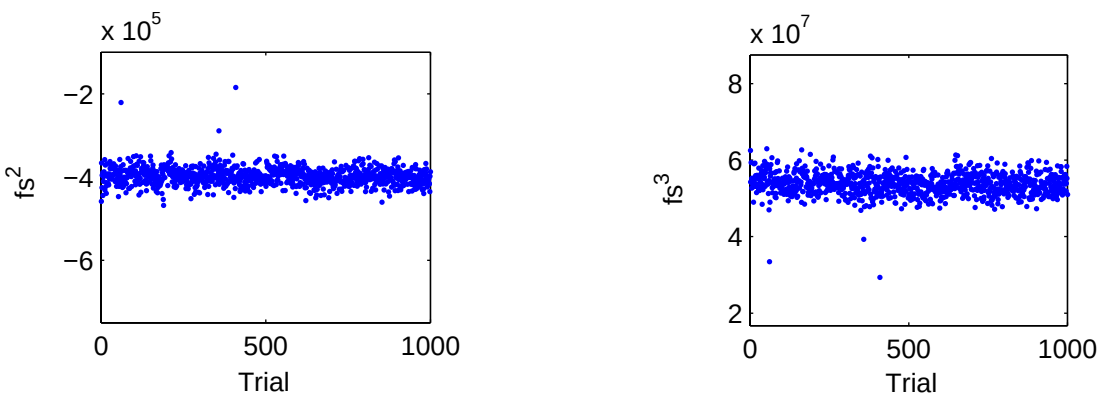

(a) Quadratic Coefficient Values

(b) Cubic Coefficient Values

Figure 14: Single-shot quadratic (a) and cubic (b) spectral phase coefficients, as sampled over a period of one hour.

Order	Mean	Standard Deviation
2	$-3.98 \times 10^5\ fs^2$	$2.2 \times 10^4\ fs^2$
3	$5.37 \times 10^7\ fs^3$	$2.9 \times 10^6\ fs^3$

Table 1: Fluctuation statistics for the target system. These correspond to the phase coefficients shown in figure 14.

4.3 Target System Characterization

In order to characterize the target laser system, the amplified pulse spectral phase was measured as a function of compressor setting. Interference spectra were recorded over the whole compressor range at 0.5 mm intervals. Each interference spectrum was averaged over approximately 2500 shots and used to determine a spectral phase trace. Several of these traces are shown in figure 15.

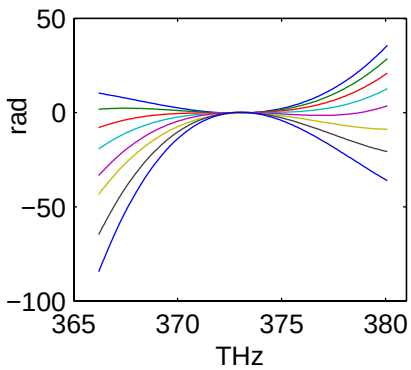

Figure 15: Phase curves recorded for the amplified laser pulse. Proceeding from the concave-up trace, the displayed phases were measured at grating micrometer positions of 1, 3, 7, 10, 13, 16, 19 and 22 mm.

In order to reduce the data, a third order polynomial was fit to each phase curve. The fit coefficients vs. compressor setting are shown in figure 16.

The time domain pulses predicted by the SPIDER measurement showed minimum duration around the middle of the compressor range. The optimized pulses showed a power distribution in which small "pre-pulses" preceeded the main lobe. Figure 16 shows a typical example of a time-reconstruction near the optimized point. The main lobe has a 135 fs full width at half maximum (FWHM).

The amplified power spectrum used for temporal reconstruction is shown in figure 16, along with the transform-limited temporal profile. This theoretical minimum duration pulse profile was determined numerically by assuming the power spectrum amplitudes were all in perfect phase. It has a 73 fs FWHM, which is narrower than the observed optimum by a factor of approximately 2.

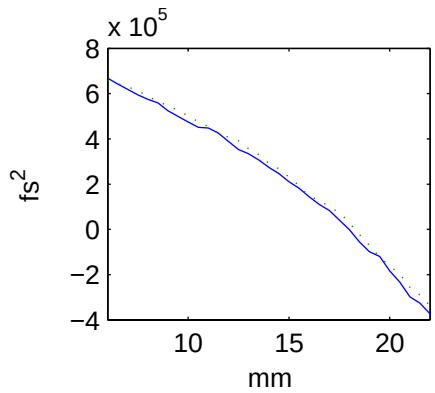

(a) Quadratic Phase

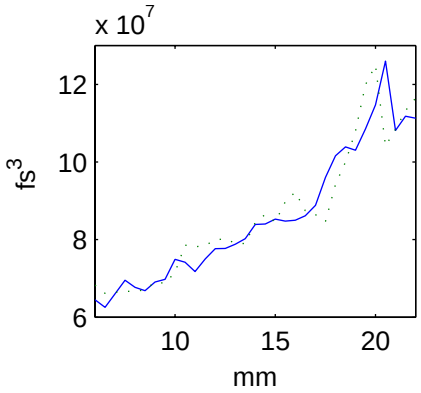

(b) Cubic Phase

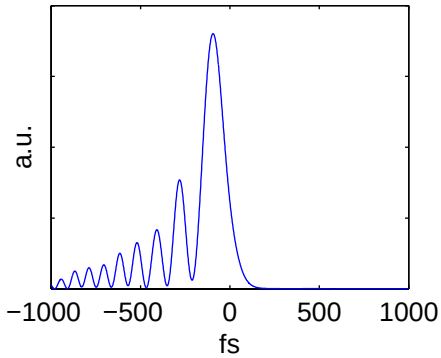

(c) Temporal Intensity

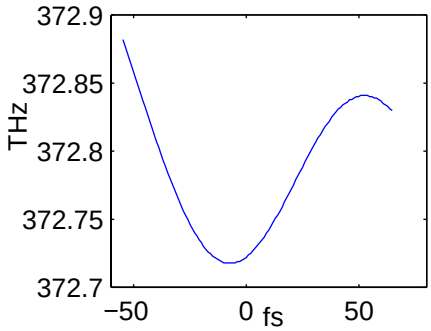

(d) Instantaneous Frequency

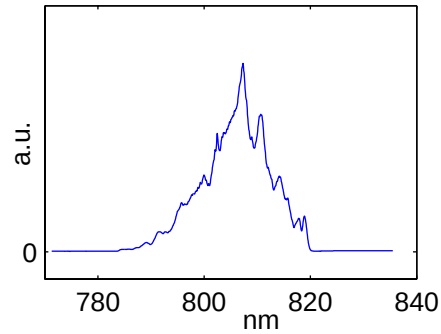

(e) Fundamental Spectrum

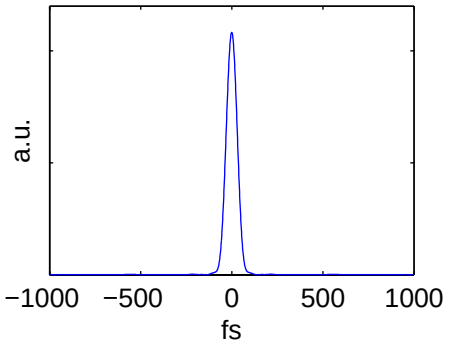

(f) Transform Limited Intensity

Figure 16: Derived quantities for the target laser system: Quadratic (a) and cubic (b) spectral phase coefficients vs. grating micrometer position. Temporal intensity (c) and instantaneous frequency (d) at the optimum compression. The amplified power spectrum used to derive time-domain data (e) and its transform-limited intensity profile (f).

4.4 Enhanced Compression Experiment

A dispersive prism pair was added to the laser system in an attempt to improve pulse compression. This was placed after the grating pair for experimental convenience. The prisms were set at a fixed, antiparallel orientation (as in figure 3, page 3) and arranged so that the position of the second prism could be varied in two dimensions. The prism separation was changed by moving the second prism along the mounting holes of the optical breadboard. The insertion depth of the second prism into the beam could be varied continuously using a translation stage.

With the added prism pair, the compressor had three degrees of freedom: grating separation, prism separation and prism insertion. The pulse spectral phase was mapped out over the full range of compressor adjustment, using the control parameter combinations summarized in table 2. 464 separate compressor states were sampled. The units used to describe the compressor settings were the grating micrometer reading μ_g, the mounting hole offset between the prisms L, and the prism insertion micrometer reading μ_p. At each setting an interference spectrum was averaged over 2500 shots and used to reconstruct the spectral phase.

μ_g	μ_p	L
$5-20$	$7-14$	2
$5-20$	$6-15$	12
$5-20$	$5-15$	25

Table 2: The control settings used to characterize the combined grating and prism compressor. The grating and prism micrometer ranges shown were covered inclusively in integer units.

A third order polynomial was fit to each measured phase. This was centered at the intensity weighted mean frequency 372.9 THz. Whenever the measured phase coefficients were plotted against a single one of the control variables (as in figure 17), the results were seen to be linear to a good approximation. Further, the slope of the traces was fairly consistent regardless of the constant values chosen for the remaining control variables. To reduce the data, the coefficients were fit to a 3 dimensional polynomial, the three dimensions being the compressor control parameters. A linear relation was chosen in each dimension, based on the behaviour seen in the single variable plots. The resulting empirical equations give the quadratic and cubic phase coefficients:

$$\begin{aligned}\phi_2 &= -9.4\times10^4\mu_g + -4.3\times10^3\mu_p + 4.5\times10^3 L + 7.5\times10^5\\ \phi_3 &= 4.3\times10^6\mu_g + 1.2\times10^6\mu_p + -2.7\times10^5 L + 3.0\times10^6\end{aligned}$$

A statistical parameter called the coefficient of determination describes the quality of the polynomial fit, and ranges from 0 for a complete lack of correlation

to 1 for a perfect fit. The coefficients of determination for the quadratic and cubic phase fits were 0.987 and 0.955 respectively. This validates the chosen form of the fitted empirical functions.

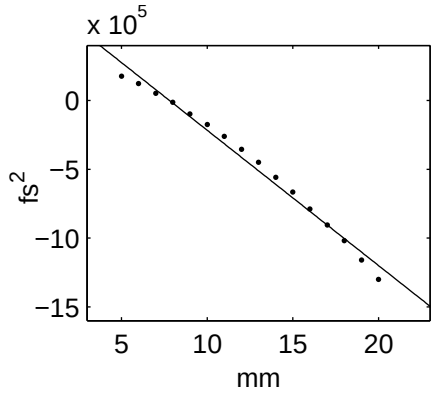

(a) Quadratic Grating Effect

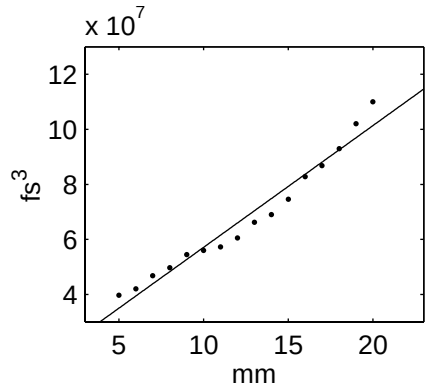

(b) Cubic Grating Effect

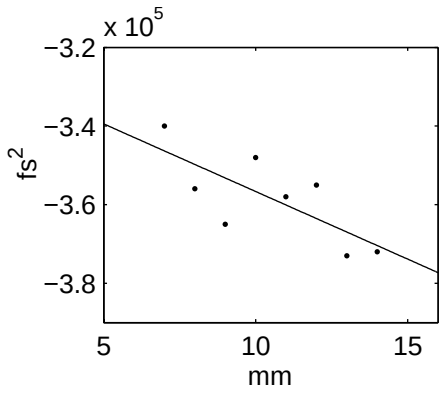

(c) Quadratic Prism Insertion Effect

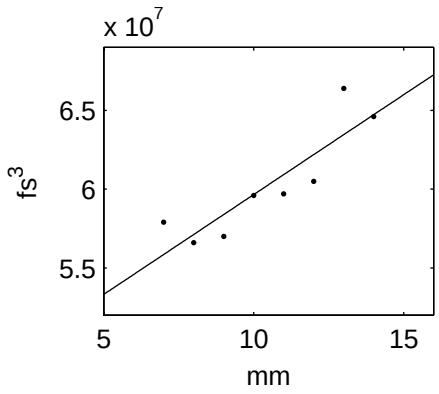

(d) Cubic Prism Insertion Effect

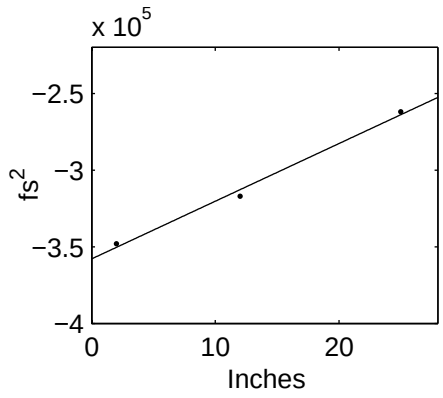

(e) Quadratic Prism Separation Effect

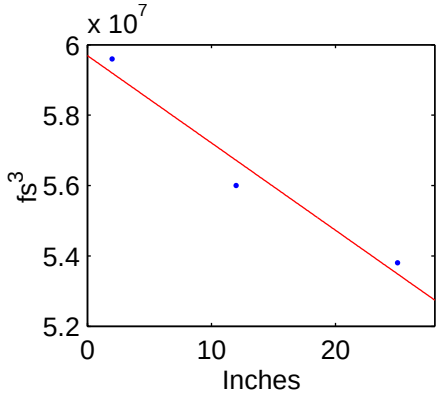

(f) Cubic Prism Separation Effect

Figure 17: Spectral phase coefficient plots for the enhanced compressor. Quadratic and cubic coefficients are plotted on the ordinates while a single compressor parameter is varied on the abscissa. The compressor parameters adjusted were grating spacing μ_g (a, b), prism insertion μ_p (c, d) and prism separation L (e, f).

5 Data Interpretation and Analysis

5.1 SPIDER Verification

Most of the verification data is sound, but one of the tests suggests possible quantitative problems with cubic phase measurement.

5.1.1 Interferometer Imbalance

Because the interferometer substrate phase (figure 10, page 27) is small with respect to typical pulse phases, it was initially assumed to be negligible. This assumption might be reconsidered however, as the pulses make multiple passes through the substrate, and the interference spectrum phase is integrated to determine the pulse phase.

The effect of the imbalance on the phase coefficients may be estimated, beginning with a polynomial fit to the measured substrate dispersion ϕ_s:

$$\begin{aligned}\phi_s\{\nu\} &= \phi_{s2}(\nu-\nu_c)^2 + \phi_{s3}(\nu-\nu_c)^3 + \phi_{s4}(\nu-\nu_c)^4 + \ldots \\ &= 3.2\times10^3\Delta\nu^2 - 8.0\times10^4\Delta\nu^3 - 1.6\times10^7\Delta\nu^4 + \ldots\end{aligned}$$

This is centered at the same point as the other measured phase expansions, so that all have a common definition of $\Delta\nu = \nu - 0.3729$.

The extra phase effects the individual pulse replicas differently due to both their different spectral shifts and to their different number of passes through the substrate. Dispersion contributions also appear in the measured background phase. The first pulse makes three passes through the substrate in both the unknown pulse measurement and the background phase measurement. The second pulse makes a single pass through the substrate in each case. The phase difference returned from the interference spectrum becomes:

$$\begin{aligned}\phi_1\{\nu\} - \phi_2\{\nu\} &= \Big(\phi_u\{\nu\} + 3\phi_s\{\nu\}\Big) - \Big(\phi_u\{\nu-\Omega\} + \phi_s\{\nu-\Omega\}\Big) - \Big[3\phi_s\{\nu\} - \phi_s\{\nu\}\Big] \\ &= \phi_u\{\nu\} - \phi_u\{\nu-\Omega\} + \phi_s\{\nu\} - \phi_s\{\nu-\Omega\}\end{aligned}$$

Where the background phase contribution is shown in square brackets. The error may be expressed in terms of the dispersion polynomial. Because the shear Ω is neccesarily small with respect to the range of wavelengths, terms containing shear powers of two or more may be neglected in the expansions:

$$\begin{aligned}\phi_s\{\nu\} - \phi_s\{\nu-\Omega\} &= \Big(\phi_{s2}\Delta\nu^2 + \phi_{s3}\Delta\nu^3 + \ldots\Big) - \Big(\phi_{s2}(\Delta\nu-\Omega)^2 + \phi_{s3}(\Delta\nu-\Omega)^3 + \ldots\Big) \\ &\simeq 2\phi_{s2}\Omega\Delta\nu + 3\Omega\phi_{s3}\Delta\nu^2 + 4\Omega\phi_{s4}\Delta\nu^3 + \ldots\end{aligned}$$

The interference spectrum phase difference is divided by the shear to form the pulse phase. The error contribution to the final pulse reconstruction ϕ_E is therefore the above expression lowered one power in shear.

$$\begin{aligned}\phi_E\{\nu\} &= 2\phi_{s2}\Delta\nu + 3\phi_{s3}\Delta\nu^2 + 4\phi_{s4}\Delta\nu^3 + \ldots \\ &= 6.4\times10^3\Delta\nu - 2.4\times10^4\Delta\nu^2 - 6.4\times10^7\Delta\nu^3 + \ldots\end{aligned}$$

The quadratic and cubic terms are the error estimates for the measured spectral phase coefficients. The linear contribution is not significant to pulse shape, and may be ignored.

The estimated errors may be interpreted with reference to typical coefficient values observed in measurement (such as figure 16, page 32). Error in the quadratic coefficient is just less than 5% of the typical coefficient values, and is well within reasonable limits. The estimated error in the cubic coefficient however is of the same order of magnitude as the observed values. Cubic coefficient values may therefore be questionable due to dispersion in the interferometer.

5.1.2 Fringe Sampling Rate

The interferometer phase traces (figure 10) illustrate the dependence of the observed interference phase on sampling rate and fringe width. Where the fringe width is small, the number of samples per fringe becomes small, and a ripple appears on the measured phase. As the fringes become larger, the ripple reduces. For large fringe width, the measured phase begins to flatten.

The ripple was originally thought to be due to the sparse sampling of the fringes, caused specifically by the sampled points cycling into and out of phase with the fringe extrema. This does not seem to be the case however, as the ripple frequency does not match the beat frequency which would be found between the spectrometer sampling rate and the fringe width. The ripple frequency is different for each of the traces, and changes much like the fringe width, increasing with pulse time separation. This would suggest that the ripple might be fringe amplitude which is passing through the filtering process and appearing as phase, except that once again there is no numerical agreement. The fringe width and ripple differ by an order of magnitude. Regardless of the cause, the ripple was observed to fade when the fringes were made moderately larger. If very large fringes were avoided, the ripple could be reduced while maintaining the general shape of the phase trace.

The flattening of the phase curve at large fringe widths is probably due to the breakdown of the amplitude and phase approximation. When the fringes become large, they approach the feature size of the pulse power spectrum, and are not

properly isolated in the frequency domain. The phase effects are partially filtered out, resulting in a smaller returned phase value.

An intermediate value of 2 *ps* time separation was used for most of the remaining data collection. The intermediate value should avoid both the ripple and flattening problems encountered in phase measurement. This is particularly true because the spectrometer had a higher spectral sampling rate at the doubled pulse frequency, which either allows more points per fringe and/or smaller fringes.

5.1.3 Spectral Shift Adjustment and Shear Calibration

The different values of spectral shift and shear tested showed no significant difference in returned phase (figure 11). The shear was therefore set toward the lower end of the range of tested values, in order to improve the pulse phase sampling rate. The spectral shift was adjusted to place the interferometer pulse replicas near the centre of the stretched replica.

The two methods of shear calibration (figure 12) do not result in significantly different pulse phases. The method based on grating geometry calculation was favoured for remaining measurements. This method gives the most accurate results in the neighbourhood of the optimum pulse compression, where the assumption of a transform limited pulse becomes increasingly valid. The alternative method gives the most accurate results at whatever arbitrary pulse compression existed during calibration.

5.1.4 Shear Approximation Error

The measured and approximate shear reconstructions (figure 13) show expected trends. The measured phase reconstruction shows a linear dependence between the phase coefficients and the micrometer distance, as predicted by theory [11]. Further, the slope of the quadratic coefficient is negative, which is consistent with the known negative dispersion of a grating system and the sign conventions used.

The quadratic coefficient vs. micrometer setting shows curvature in the case of the approximated shear. This is an expected result. There will always be one compressor position for which the calibrated shear value is strictly correct. Error in the shear will be proportional to the offset of the compressor from this position, and will show sign dependence. This error results in stretching and compression of the apparent frequency scale and changes the phase curvature. The quadratic coefficient is therefore overestimated for reduced compressor length and underestimated for increased compressor length, as seen in the plot. The error

increases with offset, giving the approximated coefficient plot the observed curve.

The phase coefficients returned by the two calibration methods are reasonably similar. The validity of the approximate shear method is consistent with its common use [5], and may be explained by comparing the laser compressor and SPIDER stretcher. The residual chirp of the pulse over the entire compressor range tested remains small relative to the chirp introduced by the SPIDER. The relative error in the shear is therefore small, allowing the approximate shear method. This method is applied in taking the remaining data.

5.2 Pulse Variation

The set of measured single-shot phase coefficients (figure 14 and table 1) can be used to estimate statistical error in the multi-shot phase measurements. The statistical fluctuation error E in an averaged measurement is given by the formula [4]:

$$E = \frac{z\sigma}{\sqrt{n}} = \frac{2.58\sigma}{\sqrt{n}}$$

where n is the number of samples taken to form the average and σ is the standard deviation of the quantity being measured. The factor z describes the quality of the error estimate. Average values determined from a limited number of samples will vary, forming a statistical distribution about the true mean. The factor z is the half width of the statistical distribution falling within the determined error range. It is expressed in standard deviations. A value of $z = 2.58$ will result in error limits which are correct in 99% of all cases.

The above formula was used to estimate the statistical sampling error in the phase coefficients for various sampling sizes. Table 3 lists the estimated errors. Based on this calculation, sample sizes were set to 2500 shots when measuring phase coefficients. This should result in coefficient measurements where 99% of the measured values have statistical sampling errors below 1%.

Samples	Quadratic Error (fs^2)	Cubic Error (fs^3)
5	2.5×10^4	3.3×10^6
10	1.8×10^4	2.4×10^6
100	5.6×10^3	7.4×10^5
500	2.5×10^3	3.3×10^5
1000	1.8×10^3	2.4×10^5
2500	1.1×10^3	1.5×10^5
5000	7.9×10^2	1.1×10^5

Table 3: Estimated phase coefficient errors caused by statistical fluctuation.

5.3 Target System Characterization

SPIDER systems are usually verified against autocorrelation data. The autocorrelator available in this case was known to have bandwidth problems, giving results of limited resolution. The autocorrelation could only locate the optimum compressor setting within a few millimeters.

The SPIDER results are consistent with the autocorrelation, in that they predict an optimum compressor setting which is within the range specified by the

autocorrelator. The setting varies in day to day operation, but is usually around the mid-range of the compressor adjustment, near 12 mm.

The measured phase coefficients agree qualitatively with the values predicted by theory [11]. Specifically, both coefficient curves are linear in grating separation and of opposite sign. The sign of the coefficient slopes is consistent with the conventions adopted and the fact that the grating pair produces negative quadratic dispersion. The magnitude of the coefficient slopes do not agree with theory. This discrepancy may be due to the quality of the optical gratings. These have physical features which are only accurate to a quarter wavelength, and may not be a good approximation to the ideal gratings assumed in deriving the theoretical result.

The jagged features observed in the amplified pulse spectrum (figure 16) could be the source of a constant error in the measured phase, though this is unlikely. These spikes appear primarily in the amplifier stage, and could not be eliminated by tuning the laser system. They are of even spacing, as if caused by multiple reflections. Theoretically, they should produce two overlapped sets of constantly spaced spikes in the upconverted pulse. Even if these spikes were interpreted as interference effects however, they should not result in a contribution to the phase coefficients because their pattern would not change across the spectrum.

The phase curves and the observed and transform limited time profiles suggest that pulse compression is being limited by pre-pulsing due to third order phase. While there is a large measurement uncertainty in the actual value of the residual phase, its existence is reasonably certain since the false contribution due to interferometer dispersion should actually *reduce* the observed value, being opposite in sign. The estimated third order phase coefficient (with the assumed conventions) is on the order of 10^7. This residual third order phase may result from the fact that the target system stretcher and compressor are not matched such that their second and third order phases may cancel simultaneously. Dispersion in the amplifier may also contribute to this problem.

5.4 Enhanced Compression Experiment

The phase measurements for the enhanced compressor allow the effect of the control parameters to be compared. The empirical phase coefficient equations are used. The rate of change (slope) for each degree of freedom may be multiplied by the allowed range of parameter values. This gives an estimate of how much control each parameter exerts over each order of phase. Table 4 lists the estimated effects for the compressor controls.

For either order of phase, the grating is the dominant control parameter. The

Parameter	Order	Coefficient Range
Grating Separation μ_g	2	$2.4 \times 10^6\ fs^2$
Prism Insertion μ_p	2	$4.3 \times 10^4\ fs^2$
Prism Separation L	2	$1.2 \times 10^5\ fs^2$
Grating Separation μ_g	3	$1.1 \times 10^8\ fs^3$
Prism Insertion μ_p	3	$1.2 \times 10^7\ fs^3$
Prism Separation L	3	$7.0 \times 10^6\ fs^3$

Table 4: Estimated effect of each compressor parameter on the spectral phase coefficients.

coefficient change due to the grating is an order of magnitude greater than the change due to either of the prism parameters. This is in spite of the fact that the grating separation varies by only 2 cm compared to a 55 cm range for the prism separation. Because the prisms can have only a weak effect on the pulse phase relative to the gratings, the improvement in pulse compression will probably be marginal; when set to eliminate quadratic phase, the residual phase left by the gratings is large with respect to the range of third order phase spanned by the prisms. To determine this with more certainty requires an analysis which can find the optimum compressor setting given the observed data.

In the most general case, the parameter optimization would have to take into account the measured pulse amplitudes. The amplitude weighting is significant when comparing two phase functions. The pulse spectrum was largely symmetric however, and showed some change over the beam cross section. For these reasons, a phase-only analysis is preferable.

First note that the pulse has a narrow fractional bandwidth:

$$\frac{\Delta\lambda}{\lambda} = \frac{\Delta\nu}{\nu} \approx 10^{-2} \ll 1$$

The frequency differences $\Delta\nu$ appearing in the phase expansion will therefore be small over the pulse bandwidth, resulting in a rapidly converging series. We may assume that it is more important to reduce the quadratic coefficient than the cubic coefficient. Based on this, assume that the optimized pulse has the minimum quadratic coefficient available in the allowed range of compressor parameters. In this case, a zero value is possible. There will be a whole family of compressor settings which result in the minimum quadratic value, so next assume that the optimized setting is that member of the subset which has the minimum cubic phase coefficient.

The physical parameter limits and the zero quadratic phase requirement form constraints over which the cubic phase function must be minimized. The best approach would usually be a numerical evaluation over the range of allowed pa-

rameters, except that in this case the constraints and objective function may all be expressed as linear in the control parameters. Constrained linear optimization is analytically tractable, whereas constrained non-linear optimizations often are not.

The grating parameter is independent, so that its physical limits form simple constant inequalities:

$$0 \leq \mu_g \quad \text{and} \quad \mu_g \leq 25$$

Due to beam spreading and the geometry of the prism pair mounts, the working range limits of the 2 prism parameters are dependent. For example, at the lowest and middle prism separations $L = 2$ and $L = 12$, the prism insertion minimum limits were different, being $7 \leq \mu_p$ and $6 \leq \mu_p$ respectively. These limits may be expressed as linear inequations:

$$\begin{aligned} L &\geq \left(\frac{2-12}{7-6}\right)\mu_p + \left(L_i - \frac{2-12}{7-6}\mu_{pi}\right) \\ &\geq (-0.1)\,\mu_p + 12.6 \end{aligned}$$

where the equal limit of the inequation is the line formed by the two points. Six such inequations are required to describe the allowed prism parameter ranges.

The final constraint is the requirement that the quadratic coefficient be zero. The empirical coefficient equation is used:

$$0 = \phi_2 = -9.4 \times 10^4 \mu_g + -4.3 \times 10^3 \mu_p + 4.5 \times 10^3 L + 7.5 \times 10^5$$

The cubic coefficient equation may be optimized according to the constraints. The resulting optimum phase setting results in a residual cubic phase of $4 \times 10^7\ fs^3$ at $(\mu_g,\ \mu p,\ L) = (8.9,\ 5,\ 25)$. The optimization was performed using the "simplex" linear programming algorithm. This was implemented by the symbolic math program Maple, version 5.1. The Maple code used is included in appendix C.

As suggested by the empirical equation, the prism pair provides little correction to the grating phase. The optimized pulse data is shown in figure 18. The quadratic coefficient observed was $6.1 \times 10^4\ fs^2$, which is a small relative value. The cubic coefficient is $4.2 \times 10^3\ fs^3$, similar to the predicted optimum. A temporal reconstruction at the calculated optimum setting (figure 18) is virtually identical to the optimized pulse observed without the prism pair (figure 16). This has several pre-pulses and a 135 fs FWHM.

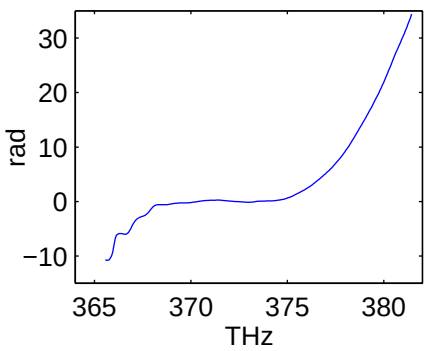

(a) Optimized Spectral Phase

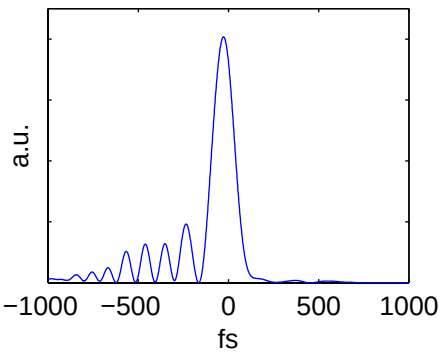

(b) Optimized Temporal Intensity

Figure 18: Optimized spectral phase (a) and temporal intensity (b) for the enhanced compressor.

6 Conclusions and Recommendations

6.1 SPIDER

The beamsplitter in the SPIDER interferometer is a potentially significant source of error. Reliability could be improved by replacing the beamsplitter with a less dispersive model. The use of a beamsplitter/compensating plate combination would probably result in some improvement, but will not completely remove dispersion errors, even when the interferometer arms are perfectly identical. A very low dispersion pellicle (membrane) beamsplitter was used initially, but was abandoned due to the poor quality of the split beams. If a pellicle must be used to eliminate dispersion, a type II non-linear crystal is recommended. This would allow collinear sum frequency generation, which would reduce the sensitivity to split beam quality. As an added benefit, this geometry is easier to align.

The calibration procedures for SPIDER are consistent.

6.2 Target Laser System

The pulse width of the laser system studied is limited by a large cubic dispersion which cannot be eliminated by the grating compressor. The output pulse is expected to show a series of pre-pulses. The best means of correction is probably to eliminate this phase at it's source by switching the target CPA system to a matched stretcher/compressor pair. Improved measurement reliability and pulse generation could also result from laser system improvements which would generate a smoother and more spatially uniform power spectrum. The amplifier would be the primary candidate for improvement.

The prism pair added to the compressor did not have strong enough dispersion effects to make a significant improvement in pulse duration. An increase in prism separation would be required to achieve the desired effect with this prism pair. This increase is not possible due to space restrictions. A practical alternative would be to use a double-pass geometry in the prism pair. This would double the dispersive effects.

References

[1] A.J.Demaria, C.M.Ferrar and G.E.Danielson Jr. Mode locking of a neodymium^{3+} doped glass laser. *Applied Physics Letters*, 8(1):22–4, January 1966.

[2] J.C.Diels and W.Rudolph. *Ultrashort Laser Pulse Phenomena: Fundamentals, Techniques and Applications on a Femtosecond Time Scale.* Academic Press, San Diego California, 1996.

[3] V.G.Dimitriev, G.G.Gurzadyan and D.N.Nikogosyan. *Handbook of Non-Linear Optical Crystals.* Springer-Verlag, Berlin, 1991.

[4] J.E.Freund. *Modern Elementary Statistics.* Prentice Hall, New Jersey, 3 edition, 1967.

[5] C.Iaconis and I.A.Walmsley. Self-referencing spectral interferometry for measureing ultrashort optical pulses. *IEEE Journal of Quantum Electronics*, 35(4):501–9, April 1999.

[6] D.J.Kane and R.Trebino. Characterisation of arbitrary femtosecond pulses using frequency resolved optical gating. *IEEE Journal of Quantum Electronics*, 29(2):571–9, April 1993.

[7] L.Lepetit, G.Cheriaux and M.Joffre. Linear techniques of phase measurement by femtosecond spectral interferometry for applications in spectroscopy. *Journal of the Optical Society of America B*, 12(12):2467–74, December 1995.

[8] T.H.Maiman. Stimulated optical radiation in ruby. *Nature*, 187:493–4, August 1960.

[9] P.Maine et al. Generation of ultrahigh peak power pulses by chirped pulse amplification. *IEEE Journal of Quantum Electronics*, 24(2):398–403, February 1988.

[10] O.E.Martinez, J.P.Gordon and R.L.Fork. Negative group-velocity dispersion using refraction. *Journal of the Optical Society of America A*, 1(10):1003–6, 1984.

[11] J.D.McMullen. Analysis of compression of frequency chirped optical pulses by a strongly dispersive grating pair. *Applied Optics*, 18(5):737–41, March 1979.

[12] H.W.Mocker. and R.J.Collins Mode competition and self-locking effects in a q-switched ruby laser. *Applied Physics Letters*, 7(10):270–3, November 1965.

[13] F.L. Pedrotti and L.S.Pedrotti. *Introduction to Optics*. Prentice-Hall, New Jersey, 2 edition, 1993.

[14] K.L.Sala, G.A.Kenney-Wallace and G.E.Hall. Cw autocorrelation measurements of picosecond laser pulses. *IEEE Journal of Quantum Electronics*, 16(9):990–6, September 1980.

[15] S.L.Shapiro, R.R.Cavanaugh and J.C.Stephenson. Streak-camera observations of the pulse emission from a synchronously pumped continuous-wave mode-locked dye laser. *Optics Letters*, 6(10):470–2, October 1981.

[16] T.M.Shuman, I.A.Walmsley et al. Real-time spider: ultrashort pulse chararcterisation at 20 hz. *Optics Express*, 5(6):134–43, September 1999.

[17] E.B.Treacy. Optical pulse compression with diffraction gratings. *IEEE Journal of Quantum Electronics*, 5(9):454–8, September 1969.

[18] R.C.Weast, editor. *CRC Handbook of Chemistry and Physics*. Chemical Rubber Company, Cleveland Ohio, 1976.

A The Fourier Transform

A.1 Definition

The Fourier transform is a common integral transform which maps a given function to a new representation in a reciprocal space. The Fourier transform $f_\omega\{\omega\}$ of a function $f_t\{t\}$ may be defined by the following:

$$f_\omega\{\omega\} \;=\; \mathcal{F}\big[f_t\{t\}\big] \;\equiv\; \frac{1}{\sqrt{2\pi}}\int_{-\infty}^{\infty} f_t\{t\}e^{-i\omega t}dt$$

Other definitions of the transform are sometimes used by convention. These involve changing the sign of the exponential argument or multiplying the integral by some normalizing constant.

The inverse Fourier transform:

$$\mathcal{F}^{-1}\big[f_\omega\{\omega\}\big] \;\equiv\; \frac{1}{\sqrt{2\pi}}\int_{-\infty}^{\infty} f_\omega\{\omega\}e^{+i\omega t}d\omega = f_t\{t\}$$

may be applied to a transformed function, returning it to the original space and basis. The form chosen for the inverse operation must be consistent with the conventions assumed for the original transform (as above).

The Fourier transform relates functions uniquely, so that the representation in either basis contains all information required to produce the remaining function. Two functions related by a Fourier transformation are referred to as a Fourier transform pair.

A.2 Shift Theorem

Modifying one member of a Fourier transform pair will necessarily result in changes to the remaining function. The shift theorem demonstrates that moving a function with respect to its coordinate system is equivalent to phase modulating its transform pair function. For example, when a temporal function f_t is delayed by some time τ,

$$f_t\{t\} \;\longrightarrow f'_t\{t\} \;\equiv\; f_t\{t-\tau\}$$

the related spectral function is phase modulated:

$$\begin{aligned}
f'_\omega\{\omega\} &\equiv \mathcal{F}\big[f'_t\{t\}\big] \\
&= \frac{1}{\sqrt{2\pi}}\int_{-\infty}^{\infty} f_t\{t-\tau\}e^{-i\omega t}dt \\
&= e^{-i\omega t}\frac{1}{\sqrt{2\pi}}\int_{-\infty}^{\infty} f_t\{t-\tau\}e^{-i\omega t-\tau}dt \\
&= e^{-i\omega t}\frac{1}{\sqrt{2\pi}}\int_{-\infty}^{\infty} f_t\{t-\tau\}e^{-i\omega(t-\tau)}d(t-\tau) \\
&= e^{-i\omega t}\frac{1}{\sqrt{2\pi}}\int_{-\infty}^{\infty} f_t\{t\}e^{-i\omega t}dt \\
&= e^{-i\omega t}f_\omega\{\omega\}
\end{aligned}$$

By a similar argument, shifting a function of frequency results in modulation of the transform pair temporal function:

$$\begin{aligned}
f'_t\{t\} &\equiv \mathcal{F}^{-1}\big[f_\omega\{\omega-\Omega\}\big] \\
&= e^{i\Omega t}f_t\{t\}
\end{aligned}$$

A.3 Sampling Theorem

In many practical control, analysis or measurement applications, the signal of interest is sampled periodically, resulting in a discrete set of data points. Fourier transformations may be applied to a discrete data set to generate a transform pair function, which is also discrete. Matching transform pairs have the same number of sample points.

The discrete sample and transformed function are an approximation to the continuous signal and its continuous transform. The accuracy of the approximation is determined by the number of samples taken. Specifically, the sampling rate of either function is inversely proportional to the domain over which its counterpart function is defined. For example, the rate at which a time signal is sampled determines the domain over which its frequency representation extends. Similarly, the total duration over which the time signal is sampled determines the resolution (spacing) of the frequency function samples.

The Whittaker-Shannon sampling theorem states that a continuous function which is non-zero over a finite span may be *exactly* reproduced over that span by any discrete counterpart function having a sampling rate equaling or exceeding the inverse width of the span. Increases in sampling rate beyond the limiting case

simply extend the reproduced function to include zero-value points on either side of the non-zero range.

B SPIDER Software Description

The following is an operational description of the LabView program used to implement SPIDER. LabView programs are referred to as "virtual instruments" or simply "VI"s. These may operate alone, or as subroutines for larger VIs. Larger programs are typically divided into several subroutines. Each VI consists of a graphical user interface called the "Front Panel" and a "Wiring Diagram" which describes the data processing schematically.

B.1 Spectrometer Readout

The SpecGraph VI reads the LineSpec spectrometer detector through the GageScope Oscilloscope. Three voltage signals are used. Channel A is the transduced CCD charge, channel B is a synchronisation trigger which indicates the start of each pixel and the Trigger is connected to the detector Sync to signal the start of each readout sequence. The pixel values are read into a one-dimensional array of 2048 elements.

A set of spectral calibration data points are entered on the SpecGraph front panel. These are used to interpolate wavelength values for each of the pixels. A fit polynomial is used for interpolation.

Several options are available when using SpecGraph:

A dark-current background may be subtracted from the pixel values. This is read from a file and then stored in memory.

Successive readouts may be averaged together to allow statistical noise reduction. This is particularly useful when recording weak calibration spectra.

The VI may be set to update its output spectrum only when the maximum pixel intensity exceeds a threshold value. This is useful when working with an asynchronous chopper.

Data may be saved to a tab-delimited ASCII text file. Data may also be repeated for the purposes of recording to file.

B.2 Calibration

When shift and shear values are calculated using pulse mean spectral amplitudes, the pulse spectra are recorded and saved to file with the SpecGraph VI. The ReadSpec VI then reads these files into memory and converts the spectrum abscissa to frequency.

Front panel graph cursors are used to set a spectral region of interest. The

intensity weighted average frequency is calculated within the region of interest. The calculated centres are used to determine shift and shear values manually. These are entered into the CalValue VI, which stores global variables.

The ReadSpec VI may also be used for wavelength calibration of the spectrometer. The spectrum abscissa may be expressed in pixel number and the weighted pixel centres of the spectral lines determined. Wavelength values may be added to the data set which is then exported to the SpecGraph VI.

When shift and shear values are calculated from the grating geometry, grating parameters are entered into the GratDisp VI, which determines the grating dispersion. The time separation is determined using the FTSI VI and entered into the CalValue VI.

The background phase is measured using the FTSI VI. This analyzes the second harmonic crystal output and determines the phase of the fringe spectrum. The derivative background phase is recorded to file. The CalValue VI must then be set to specify the location of the file. It may also specify a phase correction file which is applied to the final spectral phase. This feature was implemented but not used or tested extensively.

B.3 Determining Spectral Phase

The SpecPhse VI determines spectral phase. Calibration values are read from the CalValue and GratDisp VIs. The background spectral phase and the spectral intensity are read from the file specified in the CalValue VI.

Graph cursors are set to determine the region over which the spectral phase will be reconstructed. The fringe pattern amplitude is shifted down by the calibration shift frequency and displayed as an aid.

The fringe phase is determined using the built in LabView Fourier Transform VIs and a custom signal filtering VI called FreqFilt. Concatenation is performed. The SpecPhse front panel may be used to determine the number of phase data sets reconstructed and to fit a polynomial to the phase. The spectral phase or the polynomial fit to the spectral phase may be saved to file.

B.4 Temporal Reconstruction

When a time domain reconstruction is desired, the SpecPhse VI is called from a higher level VI called simply SPIDER. A switch on the SPIDER front panel allows the time domain calculation to be toggled on or off, as it effects execution speed. When calculating the time domain pulse, this VI passes the spectral phase to a

the Freq2Tim VI, which determines the pulse amplitude from a power spectrum file and performs an inverse Fourier transform using the given spectral phase. Front panel values may be set to cause Freq2Tim to assume zero spectral phase, in which case it calculates the transform limited time domain pulse for the given power spectrum.

C Optimization Code for the Enhanced Compressor

The following code was used to perform optimization for the analysis of the enhanced compressor. It was executed by the Maple symbolic math program.

Clear environment.

```
> restart;
```

Define the compressor parameter space. Note the prism parameter limits are dependent.

```
> paramlimits:= { 5<=Ug, Ug<=20, \
> 2<=L, L<=25, Up<=15, \
> L >= -0.0769*Up + 25.385, \
> L >= -0.1*Up + 12.6, \
> L >= 0.1*Up + 10.5 }:
```

Add C2 = 0 to form the full set of constraints.

```
> constraints := {-4280*Up+4450*L-94100*Ug+7.47e5=0} union
paramlimits;
```

Set the objective function C3.

```
> objective := 1.18e6*Up -2.67e5*L +4.29e6*Ug + 2.97e6;
```

Optimize. Check the sign to ensure a minimum absolute value.

```
> optpoint := simplex[minimize](objective, constraints union
{objective>=0});
> assign(optpoint);
> evalf(objective);
```