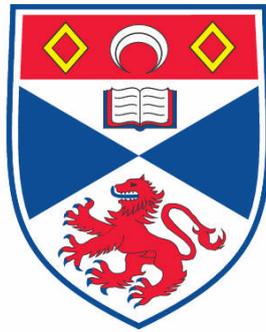


**CONTROL AND MEASUREMENT OF ULTRAFAST PULSES FOR  
PUMP/PROBE-BASED METROLOGY**

**Matthew R. Harper**

**A Thesis Submitted for the Degree of PhD  
at the  
University of St. Andrews**



**2007**

**Full metadata for this item is available in the St Andrews  
Digital Research Repository  
at:**

**<https://research-repository.st-andrews.ac.uk/>**

**Please use this identifier to cite or link to this item:**

**<http://hdl.handle.net/10023/430>**

**This item is protected by original copyright**

**This item is licensed under a  
[Creative Commons License](#)**

Control and measurement of  
ultrafast pulses for  
pump/probe-based metrology

MATTHEW R. HARPER

PHD.

JULY 2007

I, Matthew R. Harper, hereby certify that this thesis, which is approximately 45,000 words in length, has been written by me, that it is the record of work carried out by me and that it has not been submitted in any previous application for a higher degree.

Date:

Signed:

I was admitted as a research student in August 2000 and as a candidate for the degree of PhD. in August 2000; the higher study for which this is a record was carried out at NPL between 2000 and 2007.

Date:

Signed:

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

Date:

Signed:

In submitting this thesis to the University of St Andrews I understand that I am giving permission for it to be made available for use in accordance with the regulations of the University Library for the time being in force, subject to any copyright vested in the work not being affected thereby. I also understand that the title and abstract will be published, and that a copy of the work may be made and supplied to any bona fide library or research worker, that my thesis will be electronically accessible for personal or research use, and that the library has the right to migrate my thesis into new electronic forms as required to ensure continued access to the thesis. I have obtained any third-party copyright permissions that may be required in order to allow such access and migration.

Date:

Signed:

For Sally.  
And for Joe too, now.

## Abstract

In this thesis the control of ultrafast ( $10^{-13}$  s) optical pulses used for metrological applications has been investigated. Two different measurement set-ups have been considered, both based around the ‘pump-probe’ technique, where an optical pulse is divided into two parts, one to ‘pump’ or excite a physical system of interest, the other to ‘probe’ or measure the outcome. In both cases the measurement uses electro-optic sampling (EOS), where an electric field is measured by detecting changes in the optical probe pulse polarisation after interaction with the field. In the first study, a method for wavelength metrology in the terahertz (THz) region has been demonstrated by producing an optical pulse shaper and genetic algorithm to control pump pulses and so indirectly influence the THz spectra they generate. In the second study an OPO (optical parametric oscillator) has been developed to provide ultrafast optical pulses for the generation of  $< 100$  fs electrical pulses for metrology using quantum interference control (QUIC). QUIC electrical signals have been demonstrated successfully by charge accumulation measurements and the QUIC electrical pulse temporally measured using EOS, though the low signal levels due to power restrictions mean the QUIC electrical pulse is unsuitable for metrology at this time. Finally, a portable optical pulse measurement device based around frequency-resolved optical gating (FROG) has been designed, built and tested. This has been shown to be capable of retrieving amplitude and phase information in both the temporal and spectral domains for optical pulses as short as 20 fs duration. The ability to characterise shaped pulses also has been demonstrated successfully, with the requirements for full automation identified.

# Acknowledgements

Progress in science happens in collaboration with many others whose contributions can range from keen insights, to hands-on help, to being just a handy listener. As such I am eternally grateful to my colleagues at the National Physical Laboratory (NPL) for all their help, advice and input to this work over the last few years. To list them all would almost be another chapter in itself, but I must mark out a few for special mention. Foremost is Andrew Smith, who has been a friend, advisor and manager over the course of both this work and my other research. Thanks to Richard Dudley for his help and advice on the terahertz measurements and special mentions to Helen Margolis and Stephen Lea, especially to Stephen for his help and humour with the pulse shaper and FROG. I also extend my appreciation to NPL as an organisation for funding this PhD during the course of my employment with them.

From St Andrews I would naturally like to thank my supervisor Wilson Sibbett for his help, advice and all-round general approachability and in-depth knowledge over the seven years of this PhD, and I can only apologise that he ended up with so much of the thesis to look through in so little time. I also extend my appreciation to Tom Brown for his interest and encouragement, and very special thanks to Pablo Loza-Alvarez for his great help at the beginning of this work and for fuelling my enthusiasm for OPOs and FROGs.

I would also like to acknowledge and thank Derryck Reid (Heriot-Watt) and Adam Shaver (Reading) for their help and advice on the FROG and genetic algorithms, respectively.

And now, on to the personal side, which is the most difficult if only because

the right words are so hard to find. I thank my parents for the under-rated gift of life, for their endless love and support and for whatever they did right in my upbringing. I would like to thank my young son Joseph, for being a considerate enough baby to start sleeping through the night from two months old so that I actually had a chance of finishing this, and for being just so much fun when he isn't asleep. And most of all I would like to thank Sally, who was my partner when I started this and my wife and mother to my son when I finished. I'd like to thank her for all the love, support, cajoling, encouragement, help, advice, humour and belief in me through the course of this PhD, and for all the wonderful times we have together outside of it. I really would not have been able to do this without her.

And to all those friends and family who ever asked me “How’s the PhD going?”, I apologise for my shortness when answering, and can say now, “Here it is.”

# Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
1.1	Describing ultrafast optical pulses . . . . .	4
1.2	Nonlinear effects . . . . .	7
1.2.1	Linear polarisation . . . . .	8
1.2.2	Quadratic polarisation . . . . .	8
1.2.3	Cubic polarisation . . . . .	10
1.3	Applications of ultrafast pulses . . . . .	11
1.4	Thesis structure . . . . .	12
<b>2</b>	<b>Sources of ultrafast pulses</b>	<b>14</b>
2.1	Modelocked lasers . . . . .	14
2.1.1	Active modelocking . . . . .	17
2.1.2	Passive modelocking . . . . .	18
2.1.3	Kerr-lens modelocking . . . . .	19
2.2	Optical Parametric Oscillators . . . . .	21
2.2.1	Phase matching . . . . .	23
2.3	The NPL OPO . . . . .	25
2.3.1	SHG generation . . . . .	28
2.3.2	PPKTP OPO . . . . .	29
2.4	Optical pulse shaping . . . . .	31
2.4.1	Theory of LCM pulse shaping . . . . .	33
2.4.2	The NPL pulse shaper . . . . .	37
2.5	Genetic Algorithms . . . . .	41
2.5.1	Population . . . . .	43

2.5.2	Breeding . . . . .	44
2.5.3	Mutation . . . . .	45
2.5.4	Insertion . . . . .	45
2.6	Algorithm coding . . . . .	46
2.6.1	Testing of the algorithm . . . . .	47
2.6.2	Minimising the evaluations . . . . .	51
2.7	Practical application of the pulse shaper . . . . .	55
2.7.1	Experimental outline . . . . .	55
2.7.2	Experimental setup . . . . .	61
2.8	Conclusions . . . . .	68
<b>3</b>	<b>Quantum Interference Control</b>	<b>72</b>
3.1	Basic theory . . . . .	72
3.1.1	QUIC in semiconductor . . . . .	74
3.2	QUIC and ultrafast electrical metrology . . . . .	77
3.3	Principle of the work . . . . .	80
3.3.1	Device design . . . . .	81
3.3.2	Experimental Layout . . . . .	82
3.4	Charge Accumulation Measurements . . . . .	86
3.4.1	Variation with 750 nm polarisation . . . . .	86
3.4.2	Variation with 750 nm power . . . . .	89
3.4.3	Variation with pump power . . . . .	91
3.4.4	Variation with relative delay . . . . .	93
3.4.5	Discussion of charge accumulation results . . . . .	95
3.5	Time-resolved Measurements . . . . .	97
3.5.1	Electro-optic sampling . . . . .	98
3.5.2	The NPL EOS system . . . . .	99
3.5.3	Alterations to the NPL EOS system . . . . .	102
3.5.4	Experimental set-up . . . . .	105
3.5.5	EOS Measurement with 1500 nm modulator . . . . .	106
3.5.6	EOS Measurement with 750 nm modulator . . . . .	108
3.6	Conclusions . . . . .	111

<b>4</b>	<b>Measurement of Ultrafast Pulses</b>	<b>113</b>
4.1	Autocorrelation . . . . .	114
4.2	Frequency Resolved Optical Gating . . . . .	119
4.3	The NPL FROG system . . . . .	122
4.3.1	Beamsplitters . . . . .	123
4.3.2	The Delay Line . . . . .	131
4.3.3	Crystal thickness . . . . .	137
4.3.4	Spectrometer . . . . .	138
4.3.5	Retrieval Algorithm . . . . .	143
4.4	Test FROG measurements . . . . .	144
4.4.1	Chirped pulses . . . . .	144
4.4.2	Sub-12 fs pulses . . . . .	145
4.4.3	Shaped pulses . . . . .	147
4.4.4	Uncertainties in FROG . . . . .	152
4.5	Conclusions . . . . .	157
<b>5</b>	<b>Conclusions and future work</b>	<b>159</b>
5.1	QUIC for electrical pulses . . . . .	159
5.1.1	Generation and EOS of QUIC signal . . . . .	159
5.1.2	Further work on QUIC . . . . .	160
5.2	Shaping of ultrafast pulses . . . . .	161
5.2.1	Pulse shaper and genetic algorithm . . . . .	161
5.2.2	Further work . . . . .	162
5.3	Compact pulse measurement system . . . . .	163
5.3.1	The NPL FROG system . . . . .	163
5.3.2	Further work . . . . .	164
5.4	Polarisation shaping project . . . . .	165
	<b>References</b>	<b>166</b>
<b>A</b>	<b>LabVIEW Code</b>	<b>170</b>
A.1	The FROG program . . . . .	172
A.2	The GA program . . . . .	177



# List of Figures

2.1	Modelled non-modelocked output . . . . .	15
2.2	Modelled modelocked output . . . . .	16
2.3	Synchronously-pumped OPO operation . . . . .	23
2.4	Destructive interference in absence of domain inversion . . . . .	25
2.5	Quasi-phasematching through domain inversion . . . . .	26
2.6	First OPO design . . . . .	27
2.7	Photo of PPLN OPO . . . . .	28
2.8	Second OPO design . . . . .	30
2.9	Third OPO design . . . . .	32
2.10	Photo of PPKTP OPO . . . . .	33
2.11	OPO spectrum . . . . .	34
2.12	OPO autocorrelation . . . . .	34
2.13	Schematic of a $4f$ pulse shaper . . . . .	36
2.14	Photo of pulse shaper . . . . .	39
2.15	Spectral effect of max/min pixel settings . . . . .	40
2.16	Spectral minimisation at 845 nm . . . . .	41
2.17	Spectral minimisation at multiple wavelengths . . . . .	42
2.18	Genetic algorithm flowchart . . . . .	44
2.19	Genetic algorithm evolution (linear scale) . . . . .	48
2.20	Genetic algorithm evolution (logarithmic scale) . . . . .	48
2.21	Variation of generations-taken with child proportion . . . . .	52
2.22	Variation of evaluations-made with child proportion . . . . .	52
2.23	Variation of generations-taken with population size . . . . .	53
2.24	Variation of evaluations-made with population size . . . . .	54
2.25	NPL terahertz system . . . . .	57

2.26	EOS measurement of terahertz pulse . . . . .	57
2.27	Spectrum of terahertz pulse obtained by FFT . . . . .	58
2.28	Physical filtering of terahertz spectrum at 0.7 THz . . . . .	58
2.29	Integrated THz and pulse-shaper layout . . . . .	60
2.30	Spectra obtained with and without pulse shaper, absolute . . .	61
2.31	Spectra obtained with and without pulse shaper, normalised .	62
2.32	Example spectrum obtained using EOS settings for genetic algorithm . . . . .	63
2.33	Pulse-shaped filtering of 0.75 THz . . . . .	65
2.34	Pulse-shaped filtering of 0.7 THz . . . . .	66
2.35	Pulse-shaped filtering of 1.5 THz . . . . .	67
2.36	Pulse-shaped filtering of 1.5 THz (rescaled) . . . . .	68
2.37	Pulse-shaped maximisation of 0.7 THz . . . . .	69
2.38	Evolution progress for 1.5 THz filtering result . . . . .	71
3.1	Band diagrams for photoconductive switch . . . . .	79
3.2	Band diagrams for QUIC . . . . .	80
3.3	Semiconductor device for QUIC . . . . .	83
3.4	QUIC charge accumulation layout . . . . .	85
3.5	QUIC signal dependence on 750 nm polarisation . . . . .	88
3.6	Signal dependence on 750 nm polarisation in absence of 1500 nm	89
3.7	QUIC signal dependence on 750 nm power . . . . .	91
3.8	QUIC signal dependence on OPO power . . . . .	92
3.9	QUIC signal dependence on relative delay between 750 nm and 1500 nm . . . . .	94
3.10	NPL EOS system . . . . .	100
3.11	NPL EOS system altered for QUIC measurements . . . . .	103
3.12	Temporal shift of 750 nm pulses as observed on oscilloscope .	107
3.13	EOS measurements of QUIC signal . . . . .	109
3.14	Average EOS measurement of QUIC signal . . . . .	110
4.1	Intensity autocorrelation layout . . . . .	114
4.2	Interferometric autocorrelation layout . . . . .	116
4.3	Interferometric autocorrelation of a 'Mira' Ti:S laser . . . . .	118

4.4	Interferometric autocorrelation of a laser with chirp . . . . .	119
4.5	SHG-FROG layout . . . . .	120
4.6	Flowchart of FROG retrieval algorithm . . . . .	122
4.7	Schematic of the NPL FROG system . . . . .	124
4.8	Correct arrangement of FROG beamsplitters . . . . .	125
4.9	Beamsplitter/periscope arrangement . . . . .	127
4.10	FROG signal variation with arm powers . . . . .	128
4.11	FROG signal with/without periscopes . . . . .	129
4.12	Excess pulses from FROG . . . . .	130
4.13	Spectral measurement before and after traversal of FROG . . .	131
4.14	ESA-C front panel voltage vs commanded voltage . . . . .	132
4.15	Newport graph of AD-100 displacement vs voltage . . . . .	134
4.16	Interferometer arrangement for FROG calibration . . . . .	135
4.17	Interferometer measurement of AD-100 displacement vs voltage	136
4.18	Interference effect in post-FROG spectral measurement . . . .	137
4.19	Theory vs experiment for FROG signal from Mira . . . . .	139
4.20	FROG signal from Mira for 0.3 mm and 10.1 mm SHG crystals	140
4.21	He-Ne calibration of AD150/monochromator . . . . .	141
4.22	Photo of FROG . . . . .	142
4.23	Pulse retrieval (modelled) . . . . .	143
4.24	Spectrogram measurement of chirped pulses . . . . .	144
4.25	Retrieval of chirped pulses . . . . .	145
4.26	Spectrum of Femtolaser . . . . .	146
4.27	Spectrogram of Femtolaser . . . . .	147
4.28	Retrieval of Femtolaser pulses . . . . .	148
4.29	Pulse shaping of laser spectra . . . . .	149
4.30	Spectrograms of shaped pulses . . . . .	150
4.31	Retrieval of shaped pulses from spectrograms . . . . .	151
4.32	Spectrogram signal beyond range of PZT . . . . .	152
4.33	Effect of limited delays on shaped-pulse retrieval . . . . .	153
4.34	Sequential measurements of shaped-pulse spectrogram . . . . .	154
4.35	Composite spectrogram of shaped pulse . . . . .	155
4.36	Spectrogram marginals . . . . .	156

4.37 Retrieval of shaped pulse using composite spectrogram . . . .	157
A.1 Main FROG program . . . . .	173
A.2 FROG program hierarchy . . . . .	173
A.3 FROG user interface . . . . .	174
A.4 FROG user interface code . . . . .	174
A.5 HR4000 driver . . . . .	175
A.6 Spectrogram acquisition routine . . . . .	175
A.7 Delay line driver . . . . .	176
A.8 FROG algorithm . . . . .	176
A.9 GA user interface . . . . .	178
A.10 Main code for GA . . . . .	179
A.11 GA hierarchy . . . . .	180
A.12 CRi LCM driver . . . . .	180
A.13 0.7 THz suppression evaluation routine . . . . .	181
A.14 GA Create . . . . .	181
A.15 GA Replicate . . . . .	182
A.16 GA Ranking . . . . .	182
A.17 GA Select . . . . .	182
A.18 GA Recombine . . . . .	183
A.19 GA Mutate . . . . .	183
A.20 GA Reinsert . . . . .	184

# Chapter 1

## Introduction

To measure something, to *perceive* something, is always a matter of scale. The method of measurement needs to work at the same, or lesser, scale as that which is to be observed. The response time of human visual perception is in the region of an eighth of a second and everyday phenomena such as the motion of a ball can be observed comfortably. Processes on faster scales, such as the beating of an bird's wing, happen too fast to be seen and understood first-hand. So the human need for understanding drives people to develop mechanisms capable of seeing where they cannot. The development of the strobe flash, where a burst of light illuminates a subject for a fraction of a second and that instant caught on camera film, enabled the observation of events on the millisecond and even microsecond scale. Wing motion could be seen in flight, liquid drops caught mid-splash, the moment of impact of a bullet on plate-glass recorded. The usefulness of a short pulse of light to explore previously inaccessible timescales was established beyond all doubt, and when pulses became available on the picosecond and femtosecond scale, even the motions of molecules were opened up to scrutiny.

The National Physical Laboratory (NPL) based at Teddington, UK, has a particular interest in measurement techniques. As the national standards laboratory for the United Kingdom since 1901, NPL maintains the country's primary measurement standards and to this day sets new standards for measurement accuracy that are internationally recognised. One of the many

areas of research applies femtosecond optical pulses to the production of picosecond electrical pulses for the metrology of high-speed electrical instruments such as oscilloscopes (which, in terms of features, frequency response, and processing power, are a far cry from the cathode-ray devices of years back). It was the interest in bringing the electrical pulses into the femtosecond range, and in finding new applications for the lasers that produce the optical pulses, together with the expertise at St Andrews University and the close links between the two organisations, that led to the research covered in this thesis.

The picosecond and femtosecond scales are generally referred to as *ultrafast* regimes. Measurements made using ultrafast pulses can be considered broadly similar to those taken using strobe-flash photography, in that a short pulse of light is generated, interacts with the subject for an interval roughly equal to its duration, and is then detected in such a way that any changes in the pulse brought about by the interaction can be used to infer something about the state of the subject at that interval. This process of generation, interaction, detection and inference can be considered to apply to any measurement achieved using light (including continuous-wave, although this thesis will only be concerned with pulses). The common thread that links all four parts is, of course, the pulse itself. The generation will determine the nature of the pulse, and the nature of the pulse will determine how it interacts with the subject under study. The detection method will have to be sensitive to the kinds of changes the pulse is expected to undergo and a correct inference will require knowledge of the pulse both before and after interaction as well as an understanding of how any interaction will affect it. Yet, while in the strobe-flash example we can make broad assumptions about the pulse of light produced, we require far more detailed knowledge in the extremes of the ultrafast regime where even the very air the pulse passes through can affect its shape and so the final result.

The analogy to photography differs in the case of *pump-probe* measurement. Pump-probe measurement is the name given to the technique, usually employing ultrafast optical pulses, where a system is first ‘pumped’ by a pulse to generate some process, be it an excitation to a higher state, the

creation of a photoelectron or the generation of some new electromagnetic frequency, and then this process is ‘probed’ by a second, time-delayed pulse which provides information about the change when measured itself. If this were photography then a strobe-flash would be creating the image as well as illuminating it. Pump-probe measurement is used in a wide range of applications in physics and chemistry. Typically the pump and probe pulses originate from the same source, being produced by dividing part of the output at a beamsplitter glass-air interface, for instance. This has the advantage that only one source is required and that the jitter (the time error between the arrival of the pump pulse and the arrival of the probe pulse at the points of interest) will be minimal. There may be disadvantages in this method, though, if different requirements in shape, wavelength, power or duration exist for each pulse.

In this thesis the application of ultrafast pulses to two forms of pump-probe experiments will be discussed, each with a different application to metrology. Key issues will be the detection and the shaping of the ultrafast pulses used. In the case of pulse shaping the nature of the pulse at the point of interaction can be controlled. This means, for example, that there can be compensation for other optical elements in the system to ensure that the profile of the pulse at the point of interaction will be optimised for the intended effect. Successful pulse shaping requires accurate pulse measurement, not just to check the ‘before and after’ nature of the pulse shaper but also to understand the effects of other parts of the system, so that these can be accounted for correctly. When pulses reach durations of a few femtoseconds, even a mirror cannot be assumed to reflect a pulse without introducing distortion.

There will also be discussion of the generation of ultrafast pulses. Although the nature of the various available sources will be touched on briefly, the technical detail is widely available elsewhere and only those aspects that are directly relevant to the thesis will be mentioned. Instead the concentration will be on the actual physics of the pulse itself, from the point of creation to the point where it is needed and how the pathway taken affects it. By examining the causes and effects of such aspects as frequency chirp a

greater understanding of what is needed from the pulse is obtained and the steps that ensure it is capable of doing what is wanted can be taken.

So far the shape of a pulse has been spoken of as if such knowledge were obtainable as a matter of course, yet such characterisation is no mean feat in the ultrafast regime. The advantages and disadvantages of some of the techniques commonly used to achieve characterisation will be briefly covered, and there will be detailed discussion of the method chosen for this work. There will also be consideration of how characterisation and shaping can be used together, each to improve the other.

Although ultrafast pulses have long been used to study and to measure physical processes, it is nowadays possible to actually *control* the processes as well, using the phase aspect of an ultrafast pulse to influence the very probability of a quantum-mechanical process. This *quantum interference control* is used to manipulate events at the chemical and atomic scale. One physical pathway can be selected in preference to another, increasing the efficiency of a process or producing effects not previously possible, such as generating an ultrafast electrical pulse almost an order of magnitude shorter than those currently achieved.

## 1.1 Describing ultrafast optical pulses

To see how a laser can produce pulses of light with ultrafast durations, consider that within a laser cavity there may be many excited longitudinal modes spread over a range of wavelengths determined by the laser medium and the resonator dimensions. Typically, these modes will be out of phase with each other and the resulting output will be incoherent and noisy. Occasionally, where the peaks of two or more modes coincide with each other and interfere constructively (and destructive interference occurs away from the peaks, as the modes step out of phase with each other due to their differing wavelengths), a peak may appear in the output. The greater the number of modes involved in this ‘constructive’ interference, the more intense and narrower the resulting peak, the flatter the amplitude away from the peak, and the more the output takes on the form of a pulse. Therefore, if it can be arranged for

a great number of laser modes to remain in phase this way there is now a mechanism for producing ultrafast laser pulses, and this is the basic principle behind *modelocking*. A range of modelocking methods exist, some active, some passive, and are discussed in brief in this thesis. An in-depth analysis of the various techniques is unnecessary but if the reader is interested the summary by French[1] and its associated references provide extensive information. As far as this thesis is concerned, it is considered that the main feature is that the laser employed is producing ultrafast pulses through modelocking and only those details that will affect the results will be discussed.

One of the most important of such details to note is that the more modes that are locked in phase with one another (i.e. the broader the spectrum) the shorter the pulse can be. Given that pulses which we describe as being ultrafast in the *time* domain are formed from the combinations of various different modes from the *frequency* domain, there needs to be a method for describing such pulses in both domains. Beginning with the time domain, any electromagnetic wave consists of paired electric and magnetic field components. These components will always be orthogonal to each other and to the direction of travel of the wave, but their planar polarisations relative to some reference may change with space or time, as may their amplitudes. The components are also interdependent so it is usually sufficient to consider only one in mathematical treatments, and typically this will be the electric field.

A full expression for an electromagnetic wave would be particularly daunting, but fortunately a number of simplifications can normally be made. Since it will usually be linearly polarised light that is being used, just a single component of the electric field needs to be considered. Also, since the main interest is in the temporal features as opposed to the spatial (and often the assumption is that the pulse is spatially uniform over the area of interest), the field is considered only as a function of  $t$ . A pulse, then, can be described by an expression such as

$$E(t) = \sqrt{I(t)}e^{i(\omega t - \phi(t))} + c.c. \quad (1.1)$$

where  $I(t)$  is the intensity of the pulse as a function of time and  $\phi(t)$  is

the phase of the pulse as a function of time.  $\omega$  is the angular frequency of the carrier wave.

In the frequency domain, as might be expected, the description of the pulse is simply the Fourier transform of the electric field:

$$\bar{E}(\omega) = \int_{-\infty}^{\infty} E(t)e^{-i\omega t} dt \quad (1.2)$$

$\bar{E}(\omega)$  can be divided, analogously to  $E(t)$ , into ‘intensity’ and ‘phase’ parts:

$$\bar{E}(\omega) = \sqrt{S(\omega)}e^{-i\varphi(\omega)} \quad (1.3)$$

where  $S(\omega)$  is the spectrum and  $\varphi(\omega)$  is the spectral phase.

In essence, then, the pulse can be considered to have four important properties: its intensity, phase, spectrum and spectral phase. The intensity is generally the property of most interest, describing as it does the shape of the pulse in time. A typical parameter is the *full-width half-maximum* (FWHM), which is defined as the largest time interval between two points at half the maximum intensity and will be written throughout this thesis as  $\Delta\tau$ . In the frequency domain there is the analogous spectral FWHM,  $\Delta\nu$ . As noted, a broader spectrum  $\Delta\nu$  produces shorter pulses (smaller  $\Delta\tau$ ). It is common practice to benchmark the product of the two FWHMs to give what is known as the *time-bandwidth product* (TBP):

$$\Delta\tau\Delta\nu \geq k. \quad (1.4)$$

$k$  is not a unique value, but is dependent upon the pulse shape in question. Well-known values for  $k$  are 0.440 for a Gaussian pulse or 0.315 for a  $\text{sech}^2$  pulse. As a general rule, the more complicated the pulse shape, the greater the value of  $k$ . These values are the *minimum* values for those pulse shapes, and so these determine the shortest pulse of that shape that can physically be produced from the given spectrum. This situation is achieved in full-modelocking, when the spectral phase is flat implying that all available modes are in phase. In this case the pulse is referred to as being *transform-limited*.

The case where the spectral phase is not flat is referred to as *frequency chirp*. Frequency chirp is normally considered more in the context of a time-dependence of the instantaneous frequency of the pulse of the electric field, either increasing or decreasing with time. Various interactions can introduce frequency chirp to a pulse; a common cause is dispersion when the pulse travels through a material. A pulse containing frequency chirp has a longer duration than an unchirped pulse and for this reason chirp is often considered undesirable and steps taken to reduce or eliminate it. Yet there are cases where chirped pulses are positively needed, such as optical amplification. Whether frequency chirp is advantageous, disadvantageous or of no consequence depends entirely on the experimental assessment being undertaken.

## 1.2 Nonlinear effects

Much of the emphasis so far has been on the usefulness of the ultrafast aspect of the pulses, to allow the ‘illumination’ of events so that they can be examined at extreme timescales in a similar way to the use of strobe lights in high-speed photography. Another aspect of these pulses — the high intensities that they produce — opened up access to a whole new area of physics. A typical ultrafast laser may run at 1 W and produce  $8 \times 10^7$  pulses per second, meaning each pulse has only 12.5 nJ of energy. Yet this energy is delivered within an interval of roughly 100 fs, resulting in peak powers on the order of hundreds of megawatts. Focusing these pulses can easily produce peak intensities of the order of  $10^{12}$  W/cm<sup>2</sup>, and at these intensities the optical media through which the pulses pass no longer behave in a linear fashion but instead exhibit nonlinear behaviour.

To understand the origin of nonlinear effects in optics, consider the polarisation field that is set up within an optical medium by an incident optical electric field. This is normally written as

$$P = \chi E \tag{1.5}$$

where  $\chi$  is the susceptibility, i.e. the polarisation is linear with the electric

field. In reality, though, the polarisation is described by a power series

$$P = \chi^{(1)}E + \chi^{(2)}E^2 + \chi^{(3)}E^3 + \dots \quad (1.6)$$

The reason that Equation 1.5 can normally be used in optical calculations is because the higher-order terms are insignificant at typically encountered intensities. When dealing with the high peak intensities produced by ultrafast lasers, however, it becomes a different matter. Derivations of the resulting effects from calculations involving the electric field and Equation 1.6 can be found in most textbooks on optics: This section shall only describe the results of these effects.

### 1.2.1 Linear polarisation

The linear polarisation  $P = \chi^{(1)}E$  is the one experienced at everyday ‘solar’ intensities. For an optical electric field with frequency  $\omega$ , the polarisation has the same frequency and affects the propagation of the optical electric field without affecting its frequency. As a result, optical fields of *different* frequencies do not affect each other during propagation through the media (the principle of superposition). The linear polarisation is responsible for linear effects such as refraction and (when  $\chi$  is complex) attenuation.

### 1.2.2 Quadratic polarisation

The first nonlinear term in Equation 1.6 is the quadratic polarisation.

$$P = \chi^{(2)}E^2 \quad (1.7)$$

The quadratic polarisation is the strongest of the nonlinear effects (since it is the first and most dominant nonlinear term). For an optical electric field with frequency  $\omega$ , expanding out the complex form of  $E^2$  produces

$$E = Ee^{i\varphi\omega t} + E^*e^{-i\varphi\omega t} \quad (1.8)$$

$$E^2 = E^2e^{i\varphi 2\omega t} + E^{*2}e^{-i\varphi 2\omega t} + 2EE^* \quad (1.9)$$

It can be seen that the quadratic polarisation contains a new optical electric field at  $2\omega$  and a constant (zero frequency) term. The  $\omega = 0$  is effectively a DC term, and so means that a static (or quasi-static) electric field is produced within the optical medium, an effect known as *optical rectification*. The  $2\omega$  term indicates clearly that light is produced with twice the frequency of the original optical electric field, an effect known as *second-harmonic generation* (SHG).

These two effects; the ability to optically create a DC electric field and to create light at a different wavelength, are impressive enough in themselves. But, since the polarisation now contains terms with frequencies other than  $\omega$ , it can interact with other optical electric fields that are present, causing even more interesting effects to occur. For instance, if two such fields of frequency  $\omega_1$  and  $\omega_2$  are present simultaneously within the medium, then the polarisation takes the complex form

$$P = \chi^{(2)}(EE^* + E^2e^{i2\omega_1t} + E^2e^{i2\omega_2t} + E^2e^{i(\omega_1+\omega_2)t} + EE^*e^{i(\omega_1-\omega_2)t} + c.c) \quad (1.10)$$

(where numerical factors have been ignored). As well as the terms that correspond to the rectification and the SHG of each field, there are two new terms. The first corresponds to a field with frequency  $\omega_1 + \omega_2$  and the second to a field with frequency  $\omega_1 - \omega_2$ . Therefore these fields are at the sum and difference frequencies of the two original fields, and (predictably enough) these processes are referred to as *sum-frequency generation* (SFG) and *difference-frequency generation* (DFG).

DFG is the mechanism behind an important effect known as *parametric amplification*, in which an intense field is used to amplify a weaker one. Referring, respectively, to the fields as the *pump* and the *signal* field with frequencies  $\omega_p$  and  $\omega_s$ , we find that within a suitable medium DFG produces a field at the difference frequency  $\omega_i = \omega_p - \omega_s$  called the *idler* field. This idler then reacts with the pump to produce another field with frequency  $\omega_p - \omega_i = \omega_s$ , and the amount of field produced is found to be proportional to the pump intensity. In essence, then, the signal field is amplified by the pump, the pump itself being reduced. (The idler can be considered a by-

product, though often it is as useful as the signal field, if not more so.)

On the subject of quadratic polarisation, a final mention is given to the situation where two fields are present in a medium and one has a frequency of zero, i.e. it is a DC electric field. As  $\omega_1 = 0$ , there will be a term with frequency  $0 + \omega = \omega$  (i.e. linear in  $\omega$ ) that is proportional to the magnitude of the DC field and the optical field. This linear term therefore corresponds to a change in the refractive index of the medium at frequency  $\omega$  by an amount proportional to the DC field. This is called the *Pockels effect* or the *linear electro-optic effect*, and is used in electro-optic sampling to measure electric fields.

### 1.2.3 Cubic polarisation

The third term in Equation 1.6 represents the cubic polarisation term. Having already seen the variety of effects associated with the quadratic polarisation, it should be clear that there is going to be a whole host associated with the cubic, involving combinations of terms for three fields  $\omega_1, \omega_2, \omega_2$  that produce a field at a fourth frequency. (For this reason, cubic polarisation effects are often described as *four-wave mixing* or 4WM.) These include self-diffraction, Raman effects, two-photon absorption, third-harmonic generation and many others too numerous to mention. One very important one, though, is the *optical Kerr effect* (also known as *quadratic electro-optic effect* or *induced birefringence*). This is analogous to the Pockels effect, in that one of the fields present is a DC electric field, but the induced polarisation is proportional to the square of the field. The optical Kerr effect is used in optical modulators and, in the case where the electric field changing the refractive index is due to the affected optical field itself (the *AC Kerr effect*), a *Kerr lens* can be formed in a laser crystal which can contribute to self-modelocking of the laser.

### 1.3 Applications of ultrafast pulses

Now that there has been a discussion of the properties of ultrafast pulses and some of their effects, it is worth considering them in the context of feasible applications. The ultrafast nature of the pulse obviously permits the investigation of events on ultrafast timescales — the trick is how to use the pulse to provide the information about the event. One example is electro-optic sampling (Section 3.5.1). This takes advantage of the Pockels effect to sense the electric field within an electro-optic crystal by measuring the induced change in planar polarisation of an ultrafast pulse. The ultrafast nature of the pulse means that it interacts with the crystal (and hence, indirectly, the electric field) over only a very short time, with the result that it samples the electric field on a timescale of the order of the pulse duration. If another pulse arrives at a slightly different time, this samples a different part of the electric field and so a picture of the electric field as a function of time can be built up with sub-picosecond resolution.

As described, nonlinear effects can generate pulses at previously inaccessible parts of the electromagnetic spectrum from those which are more easily obtained from typical laser media. Red light can be doubled to blue, or even tripled into the ultraviolet. Near-infrared pulses, ideal for telecommunications, can be produced by optical parametric effects. More recently, ultrafast pulses have been used to generate signals in the newly-named *terahertz* domain, which itself is finding application in medical and security imaging. The high intensities that give rise to non-linear effects also have more direct applications in fields such as machining and surgery, where the high peak power of a controlled pulse destroys a localised part of a material, while the low net power prevents heating and damage to the surrounding area. Other examples abound outside the laboratory or factory, and femtosecond lasers have been used, for instance, to remove graffiti, probe clouds for pollutants and clean leaves off railway tracks. They are even proving crucial to their successors, attosecond pulses, capable of probing the electron dynamics of the atom itself!

## 1.4 Thesis structure

This thesis is divided into four main chapters and two appendices.

- In **Chapter 2** the nature of ultrafast sources is described in some detail. The basic methods of modelocking are covered before moving on to the details of optical parametric oscillators (OPOs) and the design and construction of the NPL OPO system. Optical pulse shaping is then covered, together with the development of a pulse shaper based around a liquid-crystal modulator that enables changes to be made to the spectral amplitude and phase of pulses, resulting in temporal shaping. Additionally a genetic algorithm is developed to find mask settings for desired experimental outcomes. This pulse shaper/genetic algorithm arrangement is then used to introduce changes in the spectra of terahertz (THz) pulses by shaping the optical pulses used to generate the THz radiation.
- The subject matter in **Chapter 3** relates to quantum interference control (QUIC) and its possible application to ultrafast electrical metrology through the generation of  $< 100$  fs electrical pulses. A QUIC system is built based on the OPO of Chapter 2 and the presence of a QUIC signal confirmed by measuring the charge accumulation across a device and comparing the results with theory. Time-resolved measurements of the electrical pulse generated by QUIC are then made using the NPL electro-optic sampling system.
- **Chapter 4** contains a description of the measurement of optical ultrafast pulses, where the use of autocorrelation is discussed before leading into an outline of frequency-resolved optical gating (FROG) techniques that are employed for retrieval of amplitude and phase information from pulses. The design and construction of a compact FROG system is described and the FROG arrangement used to measure the optical pulses of the previous chapters, both shaped and unshaped.
- **Chapter 5** takes the form of a summary that covers the conclusions

of the three previous chapters and recommends improvements and extensions, as well as new avenues for exploration, for the QUIC, pulse shaper and FROG apparatus. Also included is a summary of a new NPL project which brings together elements of all three.

- **Appendix A** provides a listing of the software code developed during the research of this thesis for control of the pulse shaper and the FROG. The coding is in LabVIEW, a graphical as opposed to a textual language, and both the front-panel controls (the user interfaces) and the diagrammatical code are given. Not every last piece of code is presented, the intention being to give an overview of the software rather than a resource for duplication.
- **Appendix B** provides some of the further considerations that need to be made to develop the genetic algorithm for use with the pulse shaper, taking into account the constraints on pixel-to-pixel phase variations.

# Chapter 2

## Sources of ultrafast pulses

Femtosecond-scale optical pulses, once just a theoretical concept, are now a common feature in the laboratory. The extensive range of commercially available solid-state systems bring femtosecond technology within the reach of just about any research group. Lasers that produce 100–200 fs pulses are quite standard now and pulses an order of magnitude shorter are easily generated by the latest systems. The history of the development of ultrafast lasers has been well documented (e.g. [1]) and it is not the purpose of this thesis to repeat that task. It is useful, though, to review the principles behind these lasers in order to attain a good understanding of the expected nature of the pulses in relation to the application at hand, and this is the first aim of the chapter. Once the basics of generating ultrafast pulses are understood it will be seen how the high intensities generated enable nonlinear effects to be exploited to produce ultrafast pulses centred at new wavelengths.

### 2.1 Modelocked lasers

A typical continuous-wave (CW) laser produces a relatively narrow-band (or effectively monochromatic) beam, the spectral latitude of which depends mainly on the lasing medium used (e.g. 633 nm in the case of He-Ne). (Of course, strictly speaking it is the frequency of the light that is defined and the wavelength is dependent both on this and the medium in which the

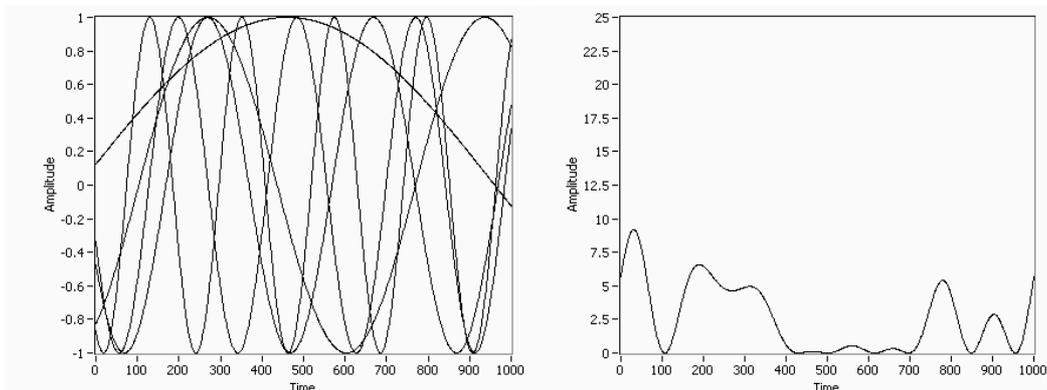


Figure 2.1: Model in LabVIEW showing five modes out of phase with each other (left) and the resulting noisy intensity at the output (right).

light is travelling. The convention that is adopted in this thesis though is to refer to the wavelength in vacuum when identifying a particular part of the visible or infra-red spectrum.) If the medium is capable of showing gain at multiple wavelengths, then multiple monochromatic waves can be present simultaneously within the laser cavity provided (i) they lie within the gain bandwidth of the medium and (ii) they follow  $\lambda_k = 2L/k$  where  $\lambda$  is the wavelength of the mode,  $L$  is the cavity length and  $k$  the mode number. These multiple modes will superimpose with random phases, generally resulting in a noisy output as in Figure 2.1. However, if the phases of different modes are aligned such that somewhere in the cavity the superposition of the modes is a superposition of peaks, a pulse will form (Figure 2.2). The more peaks that align this way, the more prominent the pulse compared to the optical noise, and the shorter its duration. If these modes maintain the same phase relation then the pulse will circulate around the cavity — the modes are locked, and the laser output is pulsed. The shortest pulses will be those with the maximum number of modes locked together, so clearly a laser medium with the broadest gain possible is required.

A comprehensive description of modelocking was derived by Haus [2] by

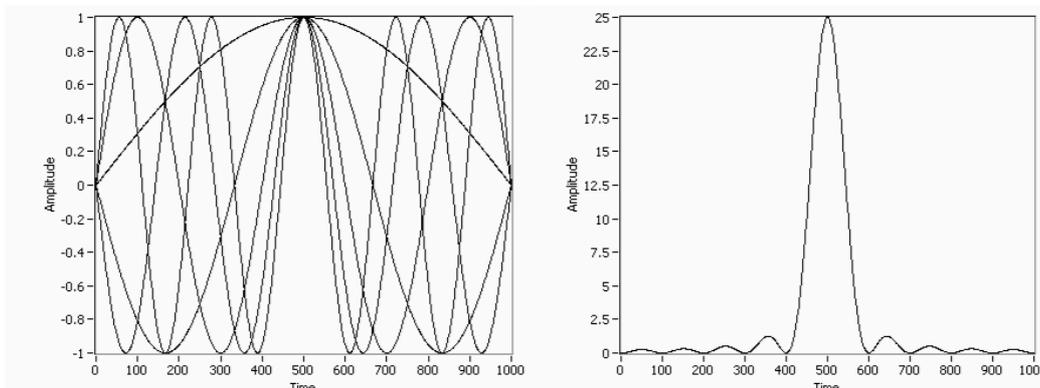


Figure 2.2: Model in LabVIEW showing five modes in perfect phase with each other (left) and the resulting pulsed intensity at the output (right).

considering modes in the frequency domain to obtain

$$- \left( m(t) \sum_{n=1}^l V_n e^{in\omega_M t} \right)_{k\text{th}} = \left( 1 + ib + \frac{ik\omega_M}{\omega_L} (\delta + g) - g \left( 1 - \left( \frac{k\omega_M}{\omega_L} \right)^2 \right) \right) V_k \quad (2.1)$$

where  $\omega_L$  is the linewidth of the lasing medium used,  $\omega_M$  the frequency spacing of the modes and  $l$  the total number of modes.  $V$  is the amplitude of the relevant mode.  $k$  is the cavity mode being considered. The terms  $b$ ,  $\delta$ ,  $g$  and  $m(t)$  relate to losses and dispersion within the cavity. (The finer details can be found in the reference.) Equation 2.1 is the modelocking equation in the frequency domain. A Fourier transform produces the modelocking equation in the time domain from which is obtained the evolution of the pulse during the time taken for one round trip of the cavity,  $T_R$ :

$$v(t-T_R+\delta T) = v(t-T_R) - \frac{\omega_0 T_R}{2Q} \left( 1 + ib - g \left( \frac{1}{\omega_L^2} \frac{d^2}{dt^2} - \frac{1}{\omega_L} \frac{d}{dt} \right) + m(t) \right) v(t-T_R) \quad (2.2)$$

where  $v$  is the complex amplitude of the pulse electric field.  $\delta T$  is the maximum time advance that the pulse can experience during one round trip in the steady state.  $Q$  is the quality factor and  $m(t)$ ,  $b$  and  $g$  again relate to losses, further defined in the reference. In his analysis [2] of the steady-state solutions to the modelocking equation Haus found that all solutions are un-

stable except for the ground state (Gaussian) solution. Any perturbations producing other states were found to quickly decay resulting only in a perturbation in phase of the carrier. Thus modelocked lasers will typically have a Gaussian spectrum (all modes being equal and locked) and a Gaussian shape in the time domain, which leads to the time-bandwidth product for the laser pulse:

$$\Delta\tau \cdot \Delta\nu = 0.44 \quad (2.3)$$

from which it is again clear that a broader spectral bandwidth is required to obtain shorter pulses. Note that it is not necessarily the case that all modelocked pulses have Gaussian profiles, merely that a Gaussian is a likely profile from a modelocked laser. Other profiles are also found to exist, such as  $\text{sech}^2$ , for example. (It is interesting to note that Gaussian and  $\text{sech}^2$  profiles share the fact that the intensity-time and intensity-frequency profiles retain the same shape.)

Having established the basic relations between locked modes and the resulting pulses it is worth briefly considering the three most common methods by which modelocking is achieved.

### 2.1.1 Active modelocking

In *active modelocking* some component within the laser cavity actively operates in a manner that produces modelocked output. The most common technique for active modelocking is placing an acousto-optic modulator (AOM) in the laser cavity to modulate the cavity losses of the laser. When an acoustic wave passes through an AOM crystal it induces refractive-index changes that diffract a large part (typically up to 70%) of any incident light energy, effectively attenuating the light along the original path. If an AOM in a laser cavity is subjected to a sinusoidally-varying acoustic wave then the amplitudes of the laser modes in the cavity will also vary sinusoidally due to the interaction with the AOM. If we consider the modulation at  $\omega_m$  of a single mode  $\omega$  and the resulting Fourier transform:

$$\int_{-\infty}^{+\infty} E(t) \cos(\omega_m t) e^{-i\omega t} dt = \frac{1}{2} \int_{-\infty}^{+\infty} E(t) (e^{i\omega_m t} + e^{-i\omega_m t}) e^{-i\omega t} dt \quad (2.4)$$

$$= \frac{1}{2} \int_{-\infty}^{+\infty} E(t) e^{-i(\omega - \omega_m)t} dt + \frac{1}{2} \int_{-\infty}^{+\infty} E(t) e^{-i(\omega + \omega_m)t} dt \quad (2.5)$$

$$= \frac{1}{2} \tilde{E}(\omega - \omega_m) + \frac{1}{2} \tilde{E}(\omega + \omega_m) \quad (2.6)$$

then it can be seen that the effect of the modulation is to introduce sidebands at  $\omega \pm \omega_m$ . The key to modelocking is to make the modulation frequency  $\omega_m$  equal to the cavity mode spacing  $\Delta\omega = 2\pi/(2Lc)$  (or some harmonic). The sidebands will then correspond to the cavity modes adjacent to the central one, and since these sidebands are in-phase with the central mode three cavity modes will be modelocked together. Since the above equations will apply to *all* modes in the cavity, it can be seen that each mode will couple energy into its adjacent modes and be phase-locked to them. The lowest loss operation for the system becomes a short pulse oscillating around the cavity.

Electro-optic modulators (EOMs) have also been used as an alternative to AOMs for active modelocking. In the case of an EOM an applied voltage changes the birefringence of the crystal, causing planar polarisation changes in the incident light. Amplitude modulation is achieved by coupling the EOM with a polariser. Frequency-modulation is also an alternative method of operation to amplitude-modulation when using an EOM, using phase-shifts to lock the modes. This tends to result in chirped pulses (Section 1.1), however. In general EOMs have not proved to be practical for modelocking.

Active modelocking is an effective method for short pulse generation but it suffers from a number of disadvantages. The necessity to match the modulation frequency to the cavity mode-spacing places severe demands on the modulator used. Far more important though is that the limited modulation speeds of AOMs restrict the width of pulses from actively-modelocked systems to 10–100 ps. To reach the femtosecond regime it is necessary to employ passive systems.

### 2.1.2 Passive modelocking

*Passive modelocking* differs from active modelocking in that it does not rely on some external signal for its operation (hence its passivity). Passive modelocking relies on the use of some optical element for which transmission

of light will vary with the intensity of that light. A saturable absorber (in current lasers this is typically a SESAM — SEemiconductor Saturable Absorber Mirror) is most often used. This absorbs light up to a saturation point, beyond which its absorption reduces. Thus a saturable absorber can be considered as a form of high-pass intensity filter, blocking light below some intensity but transmitting it above. When inserted into a laser cavity containing many randomly-phased modes, the effect will be to attenuate the initial, low intensity noise bursts. Eventually a random intensity fluctuation will occur that exceeds the saturable limit of the absorber and so passes through. This ‘intensity spike’ will experience further gain in the lasing medium and, provided the gain exceeds the intracavity losses, will continue to saturate and pass through the absorber while the other, lower intensity pulses are blocked. (It should be noted here that an important property of the saturable absorber is its relaxation time, i.e. the time taken to recover from absorption of a pulse intensity. The relation between the relaxation time and the pulse duration determines whether the saturable absorber is ‘fast’ (relaxation  $\ll$  duration) or ‘slow’ (relaxation  $\gg$  duration) Either can in fact prove suitable for passive modelocking.) The result is almost a form of ‘survival of the fittest’ – a single narrow high-intensity pulse oscillates around the cavity at the expense of all other pulse configurations, and the laser is modelocked. Passive modelocking is capable of exploiting the full bandwidth of lasing media to achieve optical pulse lengths in the femtosecond regime.

### 2.1.3 Kerr-lens modelocking

*Kerr-lens modelocking* is a form of passive modelocking, but is of enough importance to merit this separate section. It was first discovered [3] in 1989 while researchers were exploiting the coupled-cavity modelocking technique in a titanium-sapphire (Ti:S) laser<sup>1</sup>. Ti:S had been known to have a wide lasing bandwidth [4], making it an excellent candidate as a lasing medium for producing ultrafast pulses, but it was found to be ideal when it became apparent that *self-phase modulation* (SPM), crucial to coupled-cavity mod-

---

<sup>1</sup>D. E. Spence, PhD thesis, University of St Andrews, 1992

locking and usually obtained using an external optical fibre, was occurring within the Ti:S crystal itself. Pulses with  $\Delta\tau = 60$  fs could be produced from the laser with no additional optical element or other mechanism present to produce modelocking. An optical Kerr effect (c.f. Section 1.2.3) within the Ti:S crystal, where the refractive index that becomes intensity-dependent produces a longitudinal index profile that caused a frequency down-shift of wavelengths at the leading edge of a pulse travelling through the crystal and a frequency up-shift at the trailing edge of the pulse, leading to a spectral broadening that is a characteristic of SPM.

An associated feature was also occurring in the spatial domain. Since the higher intensities of the lasing beam within the cavity are at the centre of the cross-section of the beam (Gaussian profile), the refractive index of the Ti:S was higher at the centre of its cross section, causing it to behave like a positive lens, focusing the beam to emphasise the intensities at the centre in an effect known as *Kerr lensing*. By introducing into the cavity some form of aperture, such as a slit, that attenuated or blocked the weaker outer parts of the beam while passing the higher intensities at the centre, the crystal itself effectively became a saturable absorber. The use of a slit or other physical aperture in this way is known as a *hard aperturing*. There can also be an effect where the self-focusing of the beam leads to a better overlap with the pump beam for the higher intensities. This means those higher intensities experience greater gain compared to the weaker intensities, and so are again emphasised, known as a *soft aperturing*.

The spectral broadening of SPM could be further exploited by introducing some form of negative dispersion elsewhere in the cavity such as a prism pair [5]. This results in a shorter pulse than would be possible from the lasing of the medium alone. This was demonstrated by reducing the  $\Delta\tau$  of the pulses produced by the original Ti:S laser from 60 fs to 45 fs[3].

The broad spectral output of Ti:S, the action of SPM and the removal of the need for any additional components to maintain modelocking meant that Ti:S-based lasers rapidly became the preferred sources for femtosecond pulses. They are now available commercially in a variety of configurations offering different pulse energies, repetition rates and pulse widths from a few

hundred femtoseconds to under 7 fs for the latest systems<sup>2</sup>. It is a reasonable assumption that any laboratory working with ultrafast pulses will contain at least one Ti:S-based source.

If there is one disadvantage to the Ti:S laser, it is that the useable wavelengths are limited only by the gain bandwidth of Ti:S to the 650–1100 nm range. Many applications exist, though, that require femtosecond pulses at different wavelengths such as the  $\mu\text{m}$  region of the near-infrared. Therefore, other solutions need to be found, but it is seen that the Ti:S laser still has its part to play when employing solutions based on optical parametric effects — the high peak powers of modelocked lasers being the advent that made optical frequency conversion a practical method of accessing other wavelengths.

## 2.2 Optical Parametric Oscillators

Light can interact with materials in a number of ways, resulting for example in well-known effects such as refraction or reflection. These interactions are in the linear regime — the reflected or refracted light is of the same wavelength as the incident light. But as discussed in Section 1.2, if the incident light is of sufficient intensity then nonlinear effects will occur, producing light at different wavelengths to the original. Given sufficient power, scientists were no longer restricted to those regions of the electromagnetic spectrum where suitable lasing media could be found, but instead could access more of the spectrum using nonlinear crystals. High-power CW lasers made use of these processes (e.g. frequency doubling to obtain 533 nm<sup>3</sup>) but it wasn't until the advent of modelocked lasers, with their exceptionally high peak intensities, that such effects started to be properly exploited[6].

Typical nonlinear processes are second- or third-harmonic generation, but particularly of interest are *parametric processes* (Section 1.2.2). In a parametric process a strong light wave, usually referred to as the *pump* wave, produces two other waves through difference frequency generation. To see how this occurs, recall Equation 1.10 and consider only the term at the dif-

---

<sup>2</sup>e.g. Femtosource Rainbow

<sup>3</sup>Coherent Verdi

ference frequency for two waves, rewriting it as

$$P(\omega_2 - \omega_1) = \frac{1}{2}d_{\text{eff}}E_1E_2 \cos((\omega_2 - \omega_1)t) \quad (2.7)$$

where  $d_{\text{eff}}$  is the nonlinear coefficient that determines the strength of the effect. Considering a strong pump wave at frequency  $\omega_p$  in the presence of a much weaker wave designated the *signal* wave at frequency  $\omega_s$  it can be seen that a wave at the difference frequency  $\omega_p - \omega_s$  will be produced, commonly referred to as the *idler* wave with frequency  $\omega_i$ . However, the intensity of the pump wave can be such that the idler wave will grow in intensity itself to the point that it too interacts with the pump through difference frequency generation:

$$P(\omega_p - \omega_i) = \frac{1}{2}d_{\text{eff}}E_pE_i \cos((\omega_p - \omega_i)t) \quad (2.8)$$

and it can be seen that this produces more intensity at frequency  $\omega_p - \omega_i = \omega_s$  — the signal wave. The net result, then, is that the pump wave is depleted and its energy goes into amplifying the signal wave, with an idler wave by-product. Typically the idler wave is not used, so the rest of this discussion will generally refer to the signal wave only. (It should be clear that which frequency is termed ‘signal’ and which ‘idler’ is a matter of choice — conventionally,  $\omega_s > \omega_i$ .)

A system based on the process described above is referred to as an *optical parametric amplifier* (OPA). Although straightforward, it relies on some small fraction of signal wave being initially present, and also the gain only occurs on a single pass through the nonlinear crystal. Both of these restrictions can be overcome by placing the OPA within an optical cavity. If the cavity length is equal to an integral number of half-waves of the signal wanted then optical noise within the cavity can provide a signal input. If the gain exceeds the loss on each round-trip the signal wave will be progressively amplified, and as long as this gain is maintained some of the signal wave can be coupled out of the cavity for use. This system is known as an *optical parametric oscillator* (OPO). In the case where the pump is the output of a pulsed laser (such as a femtosecond laser) the cavity length is also an integral

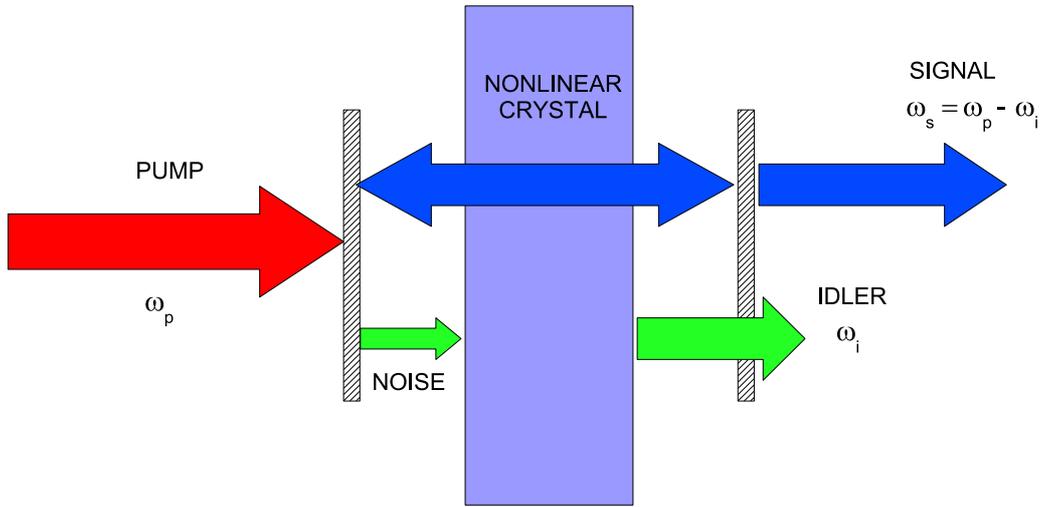


Figure 2.3: Schematic of synchronously-pumped OPO operation. A pump pulse is focused onto a nonlinear crystal placed within an optical cavity in the presence of optical noise at the idler wavelength. Parametric processes deplete the pump and couple the power into the idler and a signal pulse. By matching the cavity round-trip time for the signal pulse to the arrival of the next pump pulse the signal will progressively be increased. (Note that typically not all the pump energy will be converted.)

number of lengths of the pump cavity. Thus the signal pulse produced by one pump pulse traverses the cavity and arrives back at the nonlinear crystal at the same time as a new pump pulse arrives, resulting in amplification of the signal pulse. This arrangement is known as *synchronously pumped* (Figure 2.3).

### 2.2.1 Phase matching

Parametric generation, as seen from Equation 2.7, is the result of an induced oscillating polarisation within the nonlinear crystal in question. Typically the wavelengths of the light involved (pump, signal and idler) will be much less than the length of the nonlinear crystal in the direction of propagation. This means that, instead of there being just one source of signal within the crystal, there will be a number of sources. If these sources are out of phase

with each other, destructive interference reduces the signal and thus reduces the overall gain of the system. Maximum gain occurs when all sources are in phase and thus interfere constructively — the process for achieving this is known as *phase matching*.

Phase matching is thus an important consideration when selecting suitable nonlinear crystals for a given purpose. Dispersion within the nonlinear crystal is the main enemy of phase-matching. The general requirement for phase matching for parametric processes is

$$\omega_p n_p = \omega_s n_s + \omega_i n_i \quad (2.9)$$

where each  $n$  is the respective refractive index. Typical methods for phase matching make use of intrinsic properties of the crystal. Temperature changes may alter the refractive indices by different amounts or suitable birefringence may be exploited. An extrinsic method known as *periodic poling*, however, has had much greater success in producing crystals with good phase-matching properties.

In a periodically-poled (PP) material successive domain inversions are produced within the crystal. Each domain will therefore have an induced polarisation that is  $\pi$  out of phase with the domain immediately before it. Within a domain (or in the absence of domains), the induced polarisations along the propagation direction within the crystal become progressively out of phase with each other (Figure 2.4). Once the induced polarisation is  $\pi$  out of phase with the first induced polarisation the generated wave (created by the addition of all the induced polarisations) starts to decrease in amplitude. The crossover point (i.e. the point at which the generated wave has maximum amplitude) is known as the *coherence length*. The key to periodic poling is to make the domain lengths equal to the coherence length so that, due to the  $\pi$  phase change in the new domain, the new induced polarisation continues to add constructively to the generated wave (Figure 2.5). This method is known as *quasi phase-matching* (QPM), since the induced polarisations are not truly phase-matched (they do not interfere completely constructively within the domains). Although QPM cannot produce a generated wave as

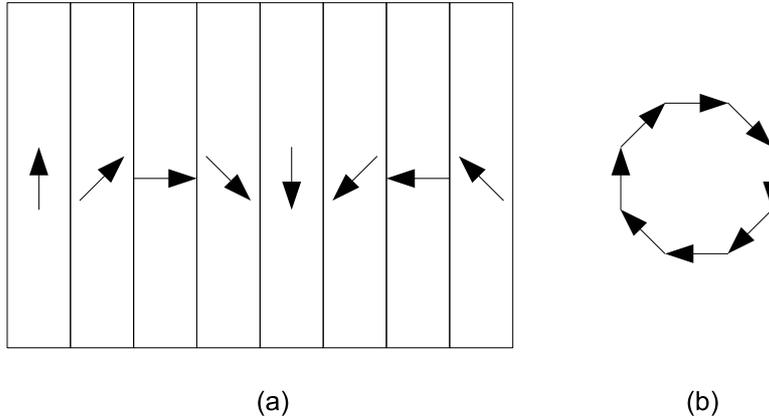


Figure 2.4: (a) Considering the progression of the generated wave along the crystal at eight different points, the induced polarisation (represented by the arrows) created by each successive source is seen to change. (b) The total wave will be the sum of sources; in this case there is total destructive interference.

large as true phase matching (TPM) theoretically would, its advantages lie in the fact that it can be applied in practice to situations where TPM is not possible. This allows freedom of choice of direction within a crystal (so that the direction with largest  $d_{\text{eff}}$  can be used), and development of parametric sources for pump and signal wavelengths where TPM is almost impossible to achieve.

## 2.3 The NPL OPO

Primarily for the purpose of quantum interference coherent control (QUIC — see Chapter 3), and also for additional applications[7], it was decided to develop an OPO at NPL capable of generating  $< 200$  fs pulses in the 1500–1550 nm range. A collaboration was started with St Andrews University, both because of the strong pre-existing links between the two organisations and moreover due to the extensive experience the St Andrews group have with these types of source. The initial stages of design and construction were to be carried out collaboratively at St Andrews, with subsequent development

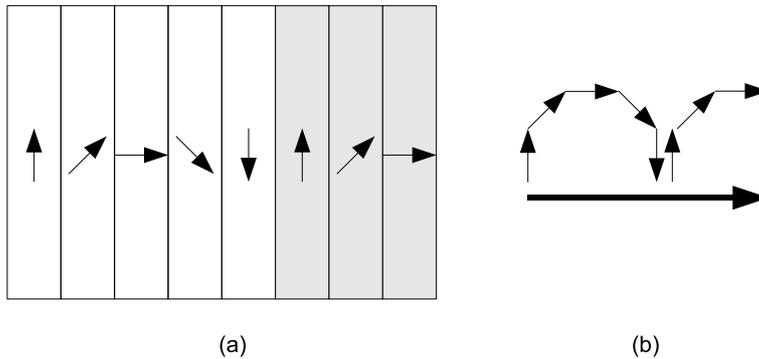


Figure 2.5: (a) This time a domain inversion (shaded area) has been produced. This results in an abrupt  $\pi$  change in the induced polarisation. (b) The total wave is now larger than it would be without an inversion, although still not at the maximum potential intensity if all sources were perfectly phase matched.

at NPL.

The initial brief for the OPO was straightforward. It should produce frequency-tunable femtosecond pulses centred around 1500 nm; it should be pumped by the existing NPL Ti:S laser<sup>4</sup>; it should be as compact as possible to minimise its footprint on the optical bench; and the 1500 nm pulses produced should be convertible to 750 nm for QUIC. Discussions with St Andrews established that the best approach would be to mount all the OPO components on an optical breadboard, employing folding mirrors to reduce the footprint as much as possible. A prism-pair[5] would introduce negative dispersion to the cavity for the purpose of transform-limiting the pulses produced, and finally the crystal for SHG of the 750 nm pulses would be placed in the cavity as well, for maximum SHG power. The first design of the OPO is shown in Figure 2.6.

When considering the choice of crystal, one extensive study[8] identified PPKTP (periodically-poled  $\text{KTiOPO}_4$ ) as suited to applications where particularly short pulses were required and reasonable pump powers were available. A PPKTP crystal would be able to provide pulses in the 1500–1550 nm range when pumped by 850 nm pulses from the NPL Ti:S. PPKTP

<sup>4</sup>Coherent Mira 900

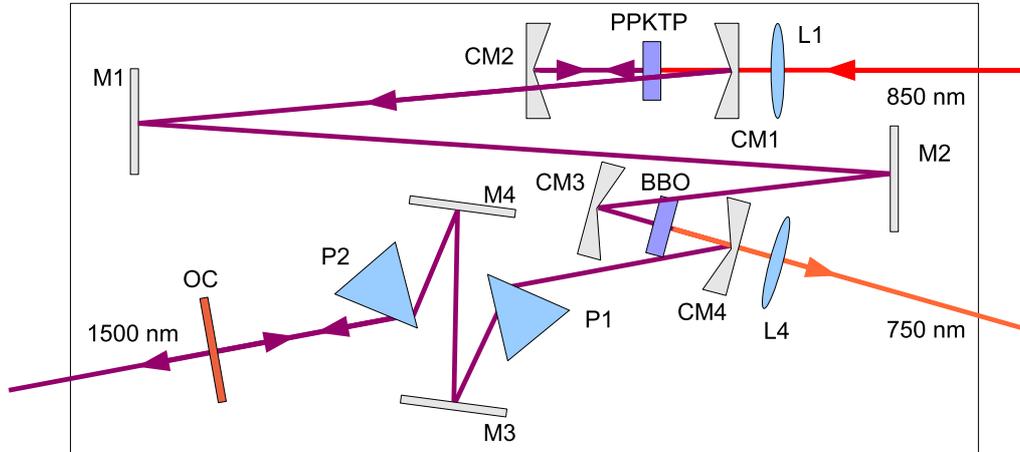


Figure 2.6: Schematic of the first design for the OPO. As well as the crystal for parametric conversion, the crystal for SHG and its associated focusing mirrors are also placed within the cavity to maximise the SHG power. L denotes a lens; M a mirror; CM a curved (focusing) mirror; P a prism; OC an output coupler.

also has the advantage that it achieves this conversion at room temperature and so there is no need to incorporate any form of temperature control into the OPO.

As there was a manufacturing delay in sourcing an anti-reflection coated PPKTP crystal the decision was made in the interim to build an initial version of the OPO around an available piece of PPLN (periodically-poled lithium niobate (LiNbO)). This would also provide suitable pulses when pumped by 800 nm pulses but required heating and temperature control to maintain the output wavelength[9], necessitating the incorporation into the OPO of a copper block heated via induction. The intention of using the PPLN was to enable issues such as focusing of the pump and matching of the cavity length to the pump repetition rate to be addressed in time for the delivery of the PPKTP, which should then simply replace the PPLN crystal.

The PPLN OPO was constructed<sup>5</sup> according to the design of Figure 2.6 and can be seen in Figure 2.7.

The green light visible is due to an additional mixing process within the

<sup>5</sup>Joint work with P. Loza-Alvarez

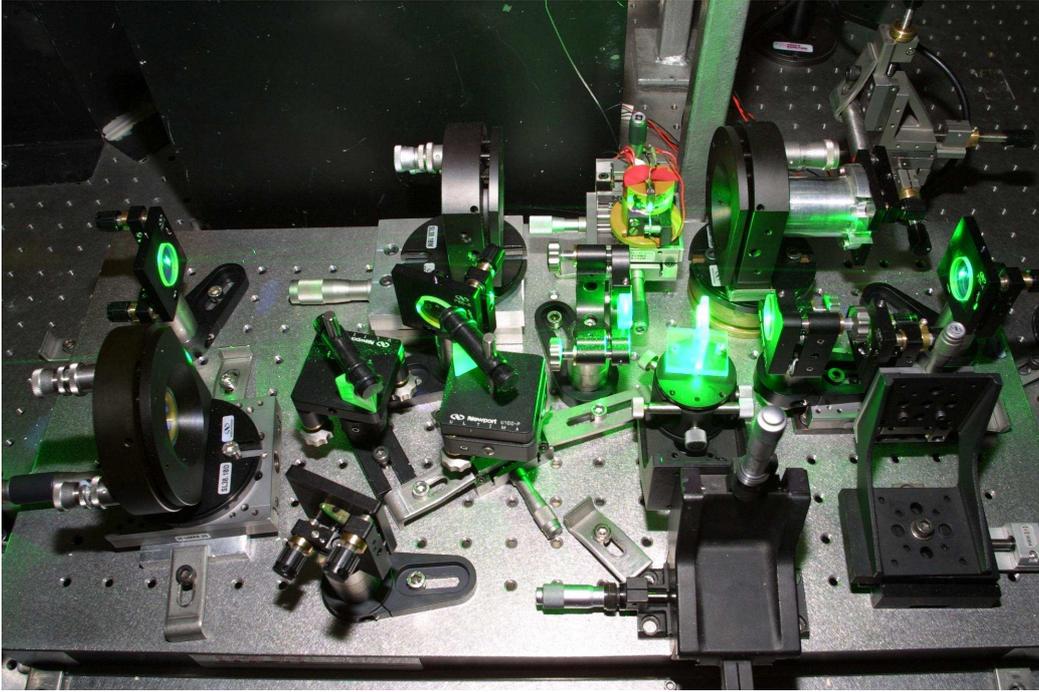


Figure 2.7: Photograph of the PPLN-based OPO when based at St Andrews. The PPLN is attached to the copper block seen at the top of the photo. The green light is due to a mixing process between the pump and signal pulses.

OPO between the pump ( $\omega_p$ ) and the signal ( $\omega_s$ ) pulses, where

$$\frac{1}{800 \text{ nm}} + \frac{1}{1500 \text{ nm}} \rightarrow \frac{1}{520 \text{ nm}} \quad (2.10)$$

i.e. sum-frequency mixing between the pump and signal pulses generates additional light in the green part of the visible spectrum. This green light follows the path of the signal pulses as far as the first prism.

### 2.3.1 SHG generation

The principal experiment for which the OPO was constructed (Chapter 3) required pulses at both the signal wavelength and also at its second harmonic. This would be achieved using basic SHG within a BBO type I crystal. (The ‘type’ describes the phasematching employed within the crystal when two

waves are present. Type I implies that the input waves have the same planar polarisation and the output has orthogonal polarisation. Type II implies that the input waves have orthogonal polarisations and the output has the same polarisation as one or the other.) As the process is relatively inefficient it was deemed that a suitable way to maximise the SHG power would be to incorporate the crystal within the OPO cavity itself, where the intensities available were much higher (since the output of the OPO is only a small fraction of that inside the cavity). Initial measurements using a pump power of 1.5 W indicated powers of 9 mW for the fundamental and 8 mW for the second harmonic.

However, further calculations into the power requirements for the proposed experiment (Section 3.4.2) determined that in fact only a very small second harmonic power was required in proportion to the fundamental, and moreover that higher power in the second harmonic was counter-productive. Given that the current configuration proposed by the instigators of the project produced relatively high second harmonic powers at the expense of the fundamental it was clear that the design was completely at odds to the needs of the experiment. Therefore the BBO crystal and its associated mirrors were removed from the OPO cavity (Figure 2.8) and SHG achieved with a basic lens-crystal-collimation arrangement external to the OPO.

### 2.3.2 PPKTP OPO

Once the PPKTP crystal was ready the PPLN and associated heating were removed and replaced with the PPKTP within a five-axis optical mount. Initial oscillation using just the crystal, lens, curved mirrors and a single external mirror was achieved easily, verifying that the focal positions and cavity lengths used with the PPLN crystal were tested suitably for the PPKTP. Unfortunately, all attempts to produce oscillation through synchronous pumping with the prism pair in the cavity were to prove frustrating and, ultimately, fruitless.

Operation of an OPO requires that the pump path and signal path overlap as closely as possible between the focusing mirrors surrounding the crystal

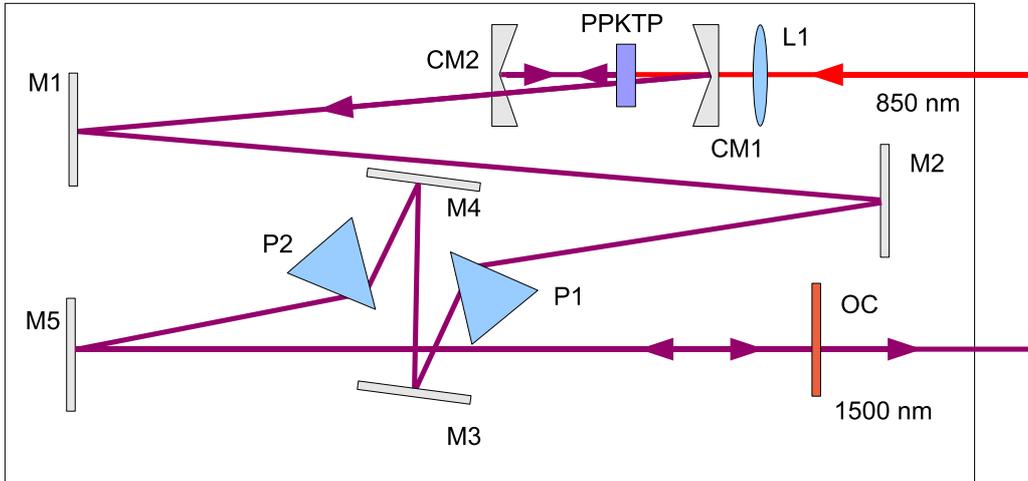


Figure 2.8: Revised layout of the NPL OPO. M denotes a planar mirror, CM a curved (focusing) mirror, L a lens and P a prism. OC indicates the output coupler.

and also that the signal return path from the end output coupler or mirror overlaps the outward path as closely as possible. With a non-prism arrangement this alignment is relatively easy as SHG of the pump within the crystal provides visible (425 nm) light that follows the same path as the signal pulses and so can be used to carry out the alignment. Once suitably aligned it is a matter of fine-tuning the cavity length to obtain synchronous pumping and then fine-adjusting the optics to optimise the alignment[10]. When introducing a prism-pair within the cavity, however, the SHG of the pump pulses cannot be used for alignment as diffraction causes it to follow a different path to the signal pulses after the first crystal. In this situation it is necessary to use an IR viewer to observe the outward and return path of the weak signal pulses (weak since there is no synchronous pumping) and align them appropriately before adjusting the cavity length. Often a mirror will be used at the end of the cavity to reflect maximum signal, being replaced by an output coupler once the alignment is sufficiently correct.

When introducing the prism pair as described, signal pulses were seen with the IR viewer and used for alignment purposes. Despite extensive efforts while changing the cavity length, this would not result in synchronous

pumping. Careful checks of the cavity length determined that it lay within the middle of the 25 mm range of the translation stage on which the end mirror was mounted, suggesting the issue was one of alignment, most likely that of the return beam. Synchronous operation without the prism pair was achieved repeatably and reliably. The prisms were even sent to St Andrews for testing within a second OPO there and found to operate correctly but still no progress was made with the PPKTP OPO at NPL. Eventually a design was selected to eliminate the prism pair and operate the OPO without them (Figure 2.9). A photograph of the PPKTP OPO operating is shown in Figure 2.10. Again the green light due to mixing between the pump and signal is much in evidence, as well as additional wavelengths in the orange.

### PPKTP characteristics

In the configuration shown, the OPO provided optical pulses that were tunable between 1500–1550 nm, the tuning being carried out by adjusting the cavity length. Figure 2.11 shows a typical spectral measurement of the OPO<sup>6</sup>. The asymmetry of the spectrum in the wavelength domain is actually desired since this is expected when the spectrum is symmetrical in the (more correct) frequency domain. The asymmetry arises since conversion is determined by  $\lambda = c/f$ , suggesting a Gaussian spectral output. Measurement in the time domain using an autocorrelator<sup>7</sup> (c.f. Section 4.1) shown in Figure 2.12 indicates a  $\Delta\tau$  of approximately 150 fs, assuming a Gaussian pulse. A power meter<sup>8</sup> used to measure the output of the OPO indicated powers of 80 mW.

## 2.4 Optical pulse shaping

So far in this thesis, the production of femtosecond pulses has been considered from suitable lasing sources or, in the case where no suitable lasing medium exists for a desired wavelength, from parametric effects. The pulses produced are, ideally, transform-limited (i.e. all modes correctly coupled) and the

---

<sup>6</sup>Taken using an Ist-Rees E202 LSA

<sup>7</sup>Femtochrome Research FR103

<sup>8</sup>Coherent Labmaster

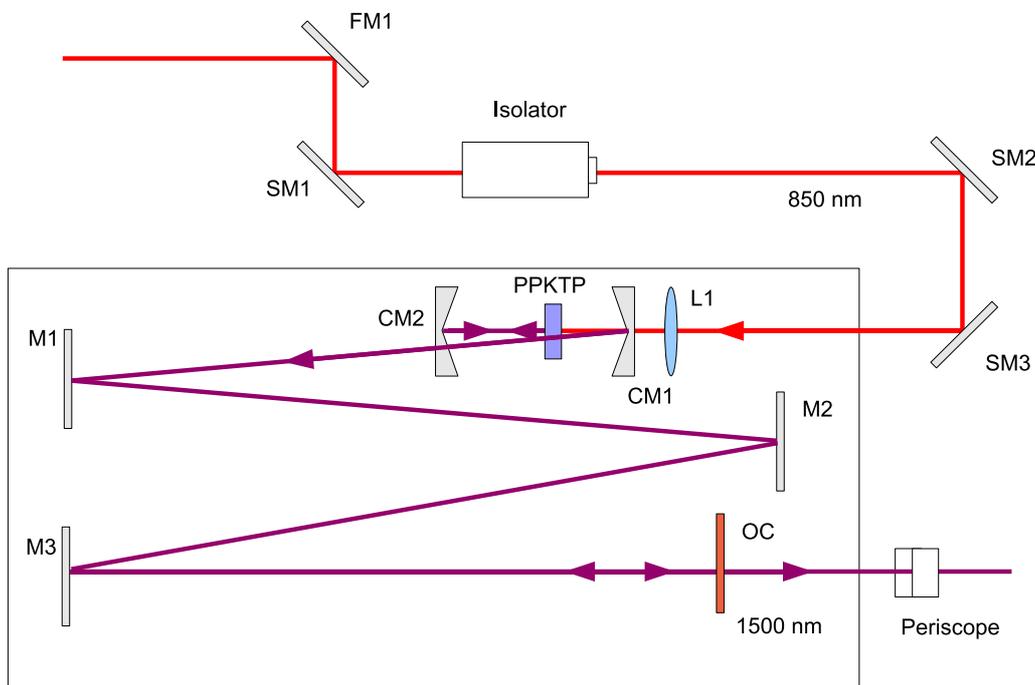


Figure 2.9: Layout of the NPL OPO configuration, revised to remove the prisms. FM denotes a ‘flipper’ mirror (i.e. one that can be moved in or out of the beam). SM denotes a steering mirror for the pump beam, M a planar mirror and CM a curved (focusing) mirror. L denotes a lens. The arrangement for SHG outside the OPO cavity is not shown. A periscope arrangement was required to raise the output to a height used in existing experimental work.

durations are as short as can be. But there are occasions when a Gaussian or  $\text{sech}^2$  transform-limited pulse *is not* the ideal. Sometimes it is more useful to have, say, a double-pulse for a pump-probe experiment, even if this means the pulses are slightly longer. Or sometimes the optics between the source and the experiment have a negative effect on the pulse through broadening, and some way of correcting for this is desired. These possibilities and more have driven the area of *pulse shaping*, where the intensity (and phase) of a pulse are altered for a desired outcome. For the research at NPL it was envisaged that pulse shaping would be a useful tool to optimise the processes and provide even greater control of electrical pulses, and so a it was decided

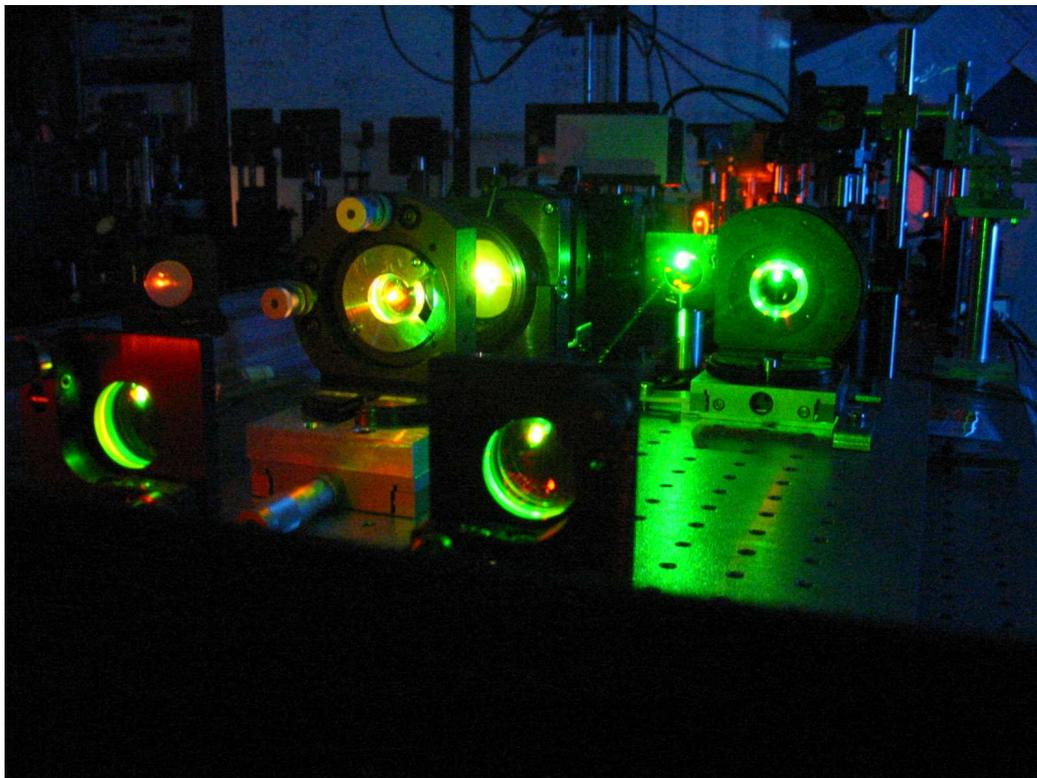


Figure 2.10: Photograph showing the operation of the PPKTP OPO. Without prisms the path of the green light matches that of the signal and is clearly visible. The crystal is the bright source to the left of the centre.

to design and construct a suitable pulse shaper.

### 2.4.1 Theory of LCM pulse shaping

As discussed in Section 2.1, an optical pulse can be described by its intensity and phase either in the time domain or in the spectral domain, and the relation between the two means that defining one will necessarily define the other. From this it follows that the profile of the pulse in the time-domain can be shaped by either temporal or spectral manipulation. For instance, we could employ some temporal function  $g(t)$  so that

$$E_{\text{shaped}}(t) = E_{\text{in}}(t)g(t) \quad (2.11)$$

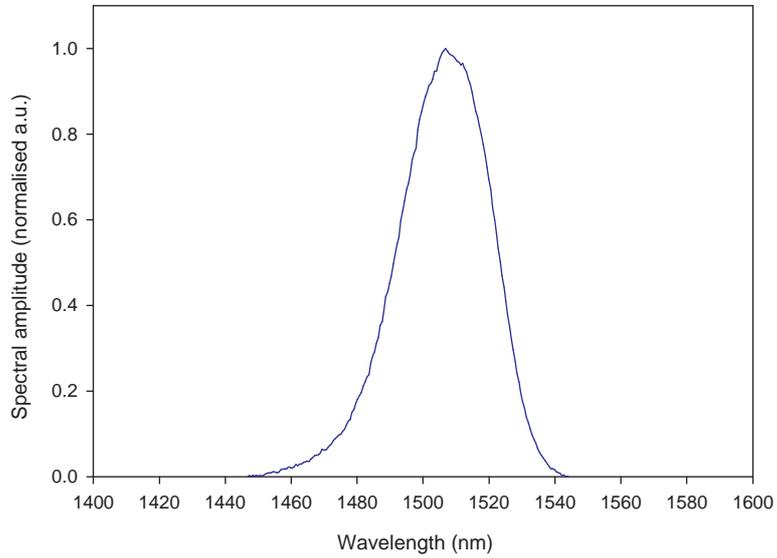


Figure 2.11: Spectral measurement of the OPO.

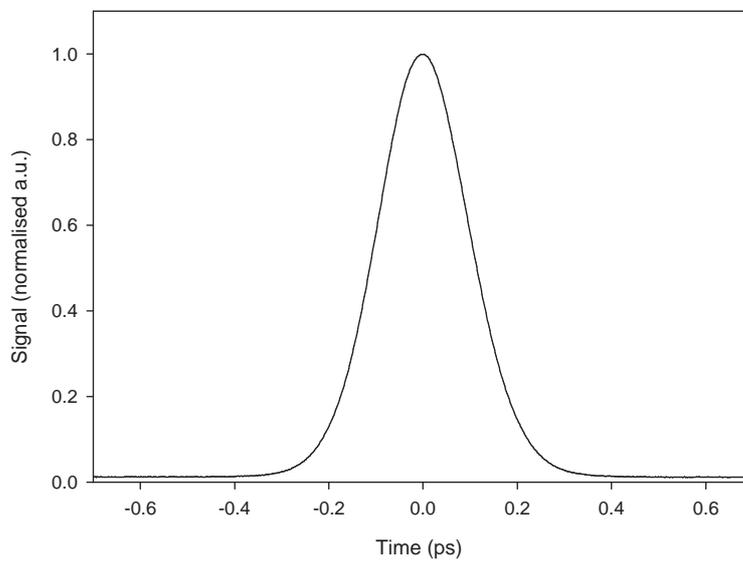


Figure 2.12: Autocorrelator measurement of the OPO output. The FWHM of the autocorrelation is deduced to be 220 fs. If the profile of the pulse were Gaussian this would correspond to a  $\Delta\tau$  of  $220/\sqrt{2} = 150$  fs.

Alternatively we could apply a spectral function:

$$E_{\text{shaped}}(\omega) = E_{\text{in}}(\omega)h(\omega) \quad (2.12)$$

Thus in the time domain the ability to shape the pulse is limited by the temporal resolution of the function used, whereas in the spectral domain it is limited by the spectral resolution of the function. Therefore, manipulation of pulses of 15 fs and 150 nm FWHM would require a system capable of reacting either at speeds in the attosecond regime or at wavelengths in the nanometre regime (unconstrained by time). The latter is typically much more achievable and so optical pulse shaping is carried out almost entirely using spectrally-based systems.

Conversion of an optical pulse into the spectral domain and back again is a relatively simple matter (Figure 2.13). The pulse is first dispersed spatially into its various wavelength components, typically using a grating. A lens or parabolic mirror collimates the wavelengths to produce the Fourier transform of the pulse. An identical lens/grating arrangement then transforms the pulse back into the time domain. If set correctly, the system will introduce no additional dispersion. (There will necessarily be some attenuation in the case of gratings due to the loss of the non-first order diffractions, but this should be uniform across the spectrum and so have no effect on the shape of the pulse.) As the correct setting normally requires that each grating be a focal distance from its lens and that the lenses be two focal lengths apart, this arrangement is commonly known as a  $4f$  pulse shaper. Typically the mid-point between the lens, where the spectral components are focused (assuming convex lenses), will be referred to as the Fourier plane of the pulse.

The actual shaping of the pulse occurs by altering the phase and amplitudes of the Fourier transform, and within a  $4f$  pulse-shaper this shaping can, in theory, be completely arbitrary within the limits of the available spectral power. An early method[11] achieved this using fixed masks, etching fused silica for phase changes or depositing Ti:Au on silica for amplitude changes. This was extended to a programmable method[12] using liquid crystals (LCs). An array of 32 LC cells was placed at the Fourier plane, each controlled by a

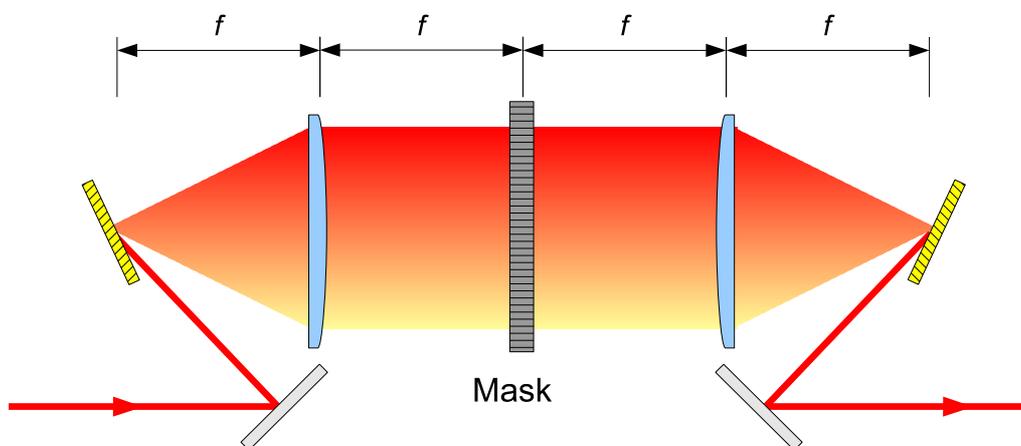


Figure 2.13: Schematic of a  $4f$  pulse shaper. A grating and lens arrangement Fourier-transforms a pulse into the spectral domain and a second identical arrangement transforms it back. By placing some form of mask at the Fourier plane to alter the spectral amplitude and phase of the pulse, the temporal amplitude and phase can be shaped.

voltage independently of the others. The voltage applied to an LC cell would cause it to rotate, inducing a proportional change in refractive index and so a phase change in the wavelengths passing through the LC in question. Thus the liquid-crystal array becomes a phase modulator, or LCM (liquid crystal modulator) for short. (Sometimes the term spatial light modulator or SLM is also used.) Using this LCM the authors of [12] were able to demonstrate conversion of a pulse to a double-pulse and dispersion-compensation of a frequency-chirped pulse.

As well as phase modulation, amplitude modulation is possible by placing a polariser at the exit of the LCM so that the planar polarisation change induced by the change in refractive index of an LC converts to an attenuation. This was extended by others[13] to simultaneous control of both phase and amplitude. LCMs are now widely available commercially for pulse shaping use, boasting many more pixels for finer spectral resolution, high-bit resolution for finer adjustment of the LCs and full computer control for integration with software. LCMs offer many advantages for pulse shaping as they introduce little unwanted additional attenuation into the  $4f$  system

and do not place any limits on the repetition rate of the pulse source. There is also no requirement on synchronisation of the LCM to the optical pulse. However, there are some disadvantages too: they can be slow to update, and the ‘dead-space’ between pixels can cause errors in the pulse-shaping due to the resulting imperfect control of the spectrum or ‘bleed-through’, where wavelengths are affected by more than one pixel. There is also a *Nyquist requirement* placed on the system, whereby adjacent pixels cannot vary in the phase change introduced by more than  $2\pi$  from each other. (This holds whether the aim is to change phase or amplitude of the spectrum.) Despite these constraints, the flexibility of an LCM-based  $4f$  system made it the most attractive choice for the NPL pulse shaper.

## 2.4.2 The NPL pulse shaper

### The CRi LCM

The LCM used in the NPL pulse shaper was a CRi (Cambridge Research and Instrumentation) SLM-640. This self-contained unit consists of two 640-pixel masks allowing simultaneous control over both phase and amplitude.

Pixel specifications	
Height	5.0 mm
Pitch	100 $\mu\text{m}$
Interpixel gap	2.0 $\mu\text{m}$
Inter-mask alignment	$\pm 2.0 \mu\text{m}$
Inter-mask separation	1.02 mm
Pulse damage threshold	100 $\mu\text{J}/\text{cm}^2$
Maximum modulation	$3\pi$ radians at longest wavelength

The drive voltage to each pixel can be set with 12-bit (2.44 mV step) resolution, i.e. to a value between 0 and 4095. Interfacing between the LCM and a computer is via a USB connection. Frames (640 12-bit values) are loaded into the on-board LCM memory buffer and from there can be transferred into either or both of the masks. The buffer is capable of holding up to 32 frames. A BNC trigger allows the LCM to load predefined frames

into the masks in sequence on a TTL trigger signal should some form of time-varying control be required.

The operation of a dual-mask pulse shaper, as explained in Section 2.4.1, modulates amplitude and phase through the action of two LC arrays and an exit polariser. It must be realised that it is not a simple matter of one mask array controlling the phase and the other the amplitude. Both arrays will introduce a phase change, and the combination of polarisation change introduced by each will determine the planar polarisation at the exit polariser and thus the total amplitude transmitted. The total effect is given by

$$E_{\text{out}} = E_{\text{in}} \cos(\delta\phi_A - \delta\phi_B) e^{i(\delta\phi_A + \delta\phi_B)/2} \quad (2.13)$$

where  $\delta\phi_A$  and  $\delta\phi_B$  are the phase changes introduced by masks  $A$  and  $B$ , respectively ( $A$  being the mask closer to the input).

The constructed pulse shaper is shown in Figure 2.14. Commands are issued to the LCM in ASCII format using whatever software is capable of transmitting over the USB interface. The commands themselves are fairly simplistic, based on loading and altering the data contained in frames and selecting frames to be applied to the two masks. Software was written in LabVIEW (Appendix A) to perform these functions and initial tests were carried out to confirm operation. No phase measurement was possible at the time so instead there was simple monitoring of the amplitude modulation using an optical spectrometer<sup>9</sup>.

Figure 2.15 shows an extreme mask application to the spectrum of a 100 fs laser source<sup>10</sup>. (The trace marked ‘Original’ is the spectrum measured when all pixels are set to provide maximum transmission i.e. when both masks are equal). The pulse shaper mask has been set so that alternate pixel pairs are set to either maximum or minimum transmission. It is apparent that neither level of transmission is achieved due to the Nyquist limit on the phase settings of the pixels. That it is possible to produce a mixture of maximum and minimum transmission is demonstrated in Figure 2.16. Here the pixel

---

<sup>9</sup>Thorlabs SP2

<sup>10</sup>Coherent Mira 900

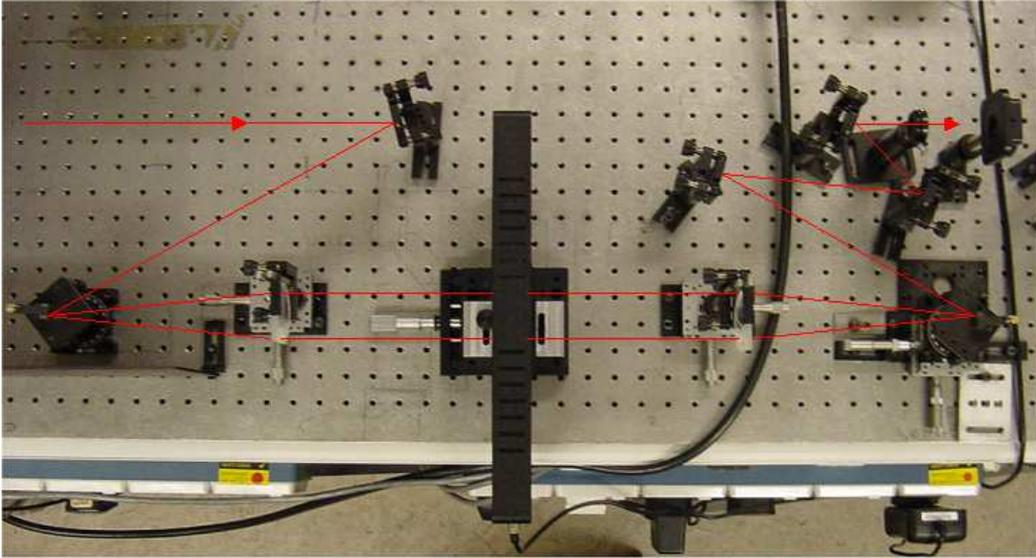


Figure 2.14: Photograph of the NPL  $4f$  pulse shaper based around a twin-array 640 pixel LCM. The first and final mirrors are mounted on flipper mounts to enable the pulse shaper to be removed from the pulse path if required. Two additional mirrors are used at the output to enable correct alignment of the exit pulses.

pair affecting 845 nm has been reduced to minimum transmission while the Nyquist requirement has been met in the rest of the mask. This required the setting of four pixel-pairs, effectively the minimum number required to switch from maximum intensity to minimum intensity and back. Note that this is partly due to the gratings used, which were chosen for operation with a different laser<sup>11</sup> with a much broader spectral output. As a result the laser spectrum here does not fill the LCM array and so the resolution is reduced. From the point of view of the test this is not an issue, as the only concern here is to confirm correct operation of the LCM using the software. It also serves to highlight that control resolution is not only a function of the pixel size but also the Nyquist requirement. Figure 2.17 demonstrates the fastest spectral oscillation achievable using this grating pair.

<sup>11</sup>see Section 2.7.1

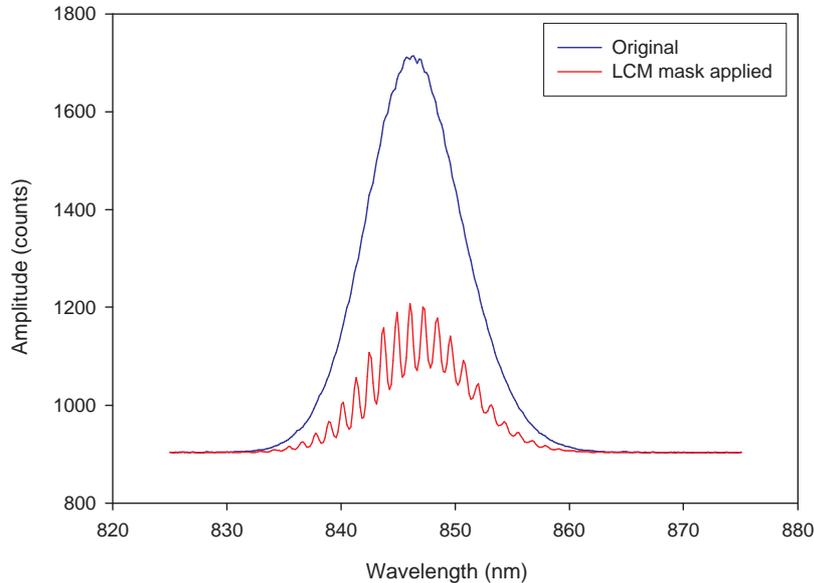


Figure 2.15: Effect on the laser spectrum when the LCM pixels are set to maximum and minimum transmission alternately.

Although the pulse shaper has been demonstrated to offer control of the spectrum of a laser source, there remains the issue of deciding what changes should actually be made. A theoretical calculation, taking the knowledge of the existing pulses as they arrive at the pulse-shaper, calculating the amplitude and phase mask required for a specific outcome and from there the corresponding settings for the LCM, can be rewarding but except for the simplest cases can be expected to be time consuming, fraught with error and possibly even futile. All aspects of the relationship between the pulse shaper and the pulse, and more importantly the relationship between the pulse and how its shape affects the outcome of the experiment to be formed, must be understood. It may even be the case that what is believed to be a desired solution (e.g. the shortest pulse possible) may not be the optimum for what is required (e.g. a pulse with a particular variation of phase may provide a better result) yet the researcher could be unaware of this possibility.

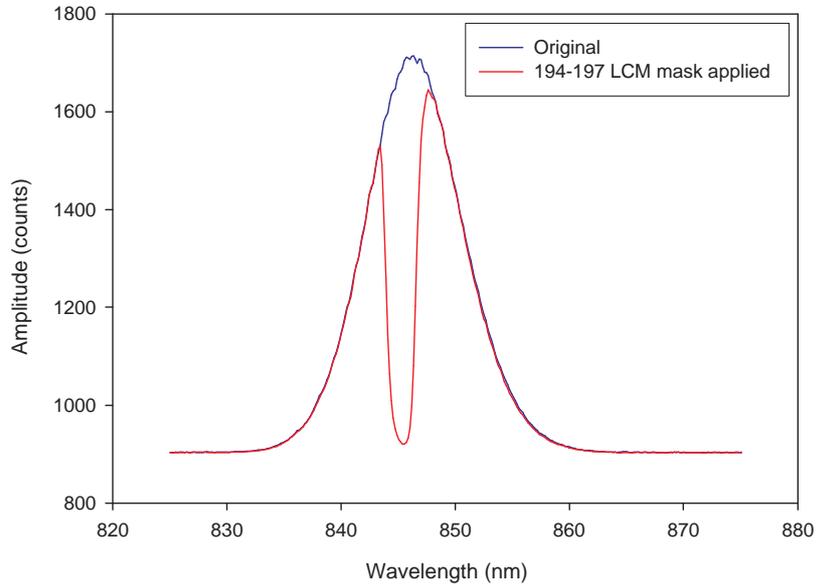


Figure 2.16: Minimisation of the spectrum at 845 nm. The Nyquist restriction requires this to use four pixel pairs (in this case 194 to 197) to maintain maximum spectral power elsewhere.

Given the complications involved in setting the pulse shaper and the uncertainty that can exist in the knowledge of what is the ideal, what would be of greatest use is some form of system where various mask settings can be tested and the outcome at the experiment monitored to determine which are better and which are worse. It would be even more useful if this system could continually refine the settings until the optimal result is reached, and this led to idea of using a genetic algorithm to seek out the optimal settings of the pulse shaper for any given experiment.

## 2.5 Genetic Algorithms

In many areas of science and technology, the particular task at hand is to find an unknown solution to a known problem, and typically that solution will be subject to some form of constraint. Sometimes an equation exists which sup-

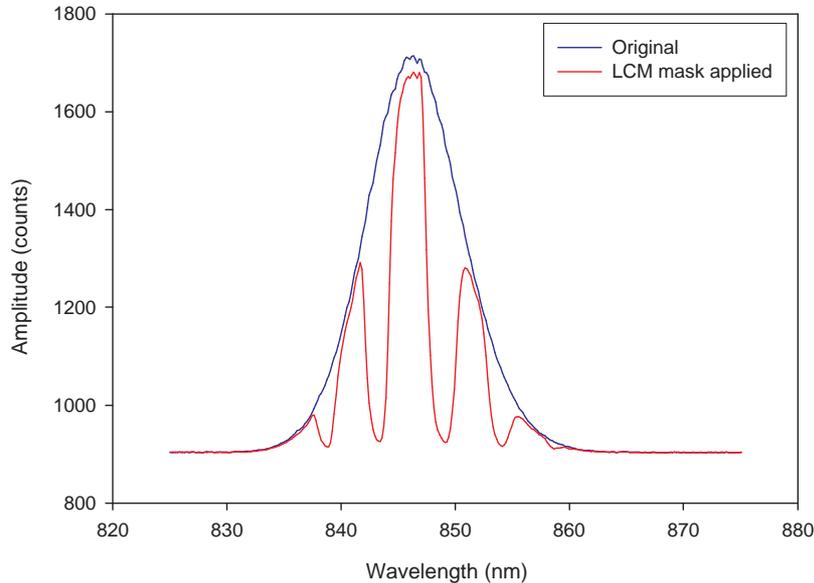


Figure 2.17: Minimisation of the spectrum at various points. This demonstrates the limit of the pulse shaper in introducing minima with the given laser/grating combination. Since the gratings do not fill the LCM array with this spectrum better performance would be possible with different gratings.

plies the solution when given the known parameters; sometimes an analytical approach provides the solution; and at other times a numerical approach or computational model can be used to perform the necessary calculations that result in the desired outcome. What such methods have in common, though, is a pre-existing understanding of the likely form of the solution or the approach that should be taken. Other, less obvious yet better, solutions may be missed altogether, and in the case where no clear approach is apparent, scientists can find themselves at a loss. Yet a number of techniques exist to solve such problems, and one of the most popular is the *genetic algorithm* (GA)[14].

The GA takes its inspiration from the biological successes of nature. All around can be found cases where the ‘problem’ is survival in a very particular kind of environment and nature appears to have arrived at ‘solutions’

of specialised creatures through random chance, without any calculation or mathematical manipulation. The reality, of course, is that these ‘solutions’ have been arrived at through the process of evolution, where a combination of chance and natural selection produces seemingly perfect adaptations over thousands or millions of years.

The application of a GA to a scientific problem has, in essence, two prerequisites. Firstly, each possible solution to the problem must be able to be represented by a string of numbers, known as the *individual* (or more commonly the *chromosome*, though this thesis will use the former term). All possible values of these strings (i.e. all possible individuals) form what will be termed the *universe*. The individual that represents the ideal (or closest physical ideal) solution must exist somewhere within this universe, so the ability of the algorithm to access the entire universe is important (Section 2.5.3). Secondly, there must be some function for assigning a *fitness* to an individual, i.e. a number that represents the quality of that individual as compared to the desired solution. If these two requirements can be fulfilled by the experiment or problem at hand, a GA should be capable of arriving at a solution, given enough time.

A GA typically follows the process outlined in Figure 2.18. A starting population is created and the fitness of each individual evaluated. ‘Breeding’ occurs between individuals to produce *offspring*, which then have some random element introduced through *mutation*. The fitness of the offspring is then evaluated and the offspring inserted into the population, replacing the less fit individuals. Each cycle of breeding/mutation/insertion is termed a *generation*. These stages are covered in greater detail below.

### 2.5.1 Population

The population is the collection of individuals for the current generation. At the beginning the individuals making up the population will often be randomly generated, but other possibilities exist such as using a previous population known to contain individuals of superior fitness, or generating random individuals with a statistical leaning towards desired features that

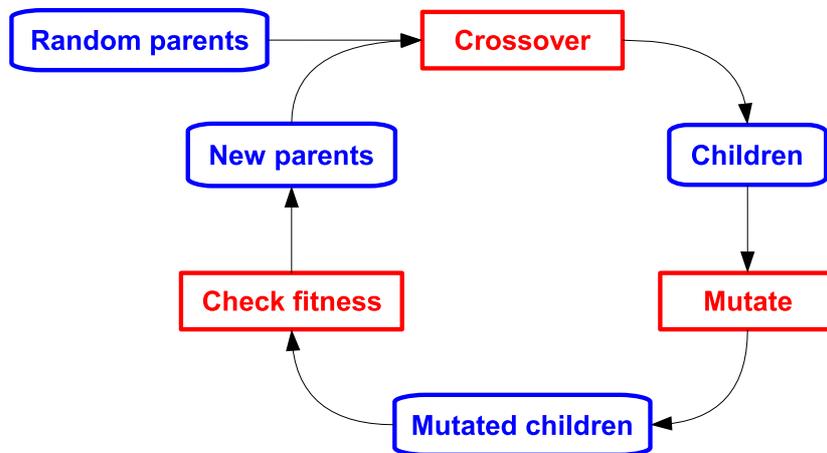


Figure 2.18: Flowchart demonstrating the operation of a genetic algorithm. Starting with a (typically) random population, individuals are bred (crossover) to produce children which are mutated, evaluated and inserted back into the population according to their fitness. Successive repetitions of this process will eventually produce the best possible solution.

produce good fitness values. (Note that this would require some knowledge of the form of the ideal solution, not always available in situations where genetic algorithms have been chosen for use.) As the algorithm advances through the generations, the overall fitness of the population is expected to improve.

### 2.5.2 Breeding

Breeding is the process of producing new individuals (offspring) that contain features of existing individuals (parents). A typical method for breeding is the cross-over technique[14]. Two individuals are selected and a random point within their length chosen. Each individual is split at this point, and their halves swapped to produce two new offspring. Of note is the method by which the parents are selected for breeding. One way is to select them completely at random, another is to select them based purely on fitness. While the latter seems a reasonable way of ensuring those traits that improve fitness survive, too much emphasis should not be given to dominant individuals in the early

stages of the population's evolution, as this can often prevent convergence towards the ideal solution. As such, a better approach is to combine fitness and random selection.

### 2.5.3 Mutation

If offspring are only produced by cross-over as described above, the parts of the universe accessible to the algorithm will be limited to those individuals that can be produced through combinations of the original population instead of the whole (for instance, if none of the individuals start with the value 5, no individual starting with the value 5 can ever appear). It is unlikely that the starting population will contain all the elements needed for the individual representing the ideal solution, so some method of accessing individuals with features not present in the current population is required. This is typically done through mutation, i.e. randomly altering values within an offspring to produce new features. As in nature, mutation often results in a less-than-fit individual which quickly disappears from the gene pool. However, occasionally a mutation will occur that provides a beneficial feature that increases the individual's fitness, and this resulting feature should eventually be shared throughout the population as it evolves. A key consideration when applying a mutation is the probability of occurrence. It is undesirable to mutate every single offspring produced yet too little mutation means the algorithm will be slow to approach a suitable solution.

### 2.5.4 Insertion

Once offspring have been created and evaluated, they (and their corresponding fitnesses) should be inserted into the current population. Genetic algorithms differ from nature in that their population is generally fixed in number. The typical course of action is to replace the least fit members of the population with the offspring. Note that these older individuals will be replaced by the newer individuals even if the newer individuals have lower fitness; this gives beneficial traits which may not yet be apparent the opportunity to appear.

## 2.6 Algorithm coding

The core routines comprise ranking of the individuals (**Ranking**), selection for breeding (**Select**), breeding (**Recombine**), mutation (**Mutate**) and insertion of the offspring into the original population (**Reinsert**). There will also be a routine for evaluation the fitness of an individual (**Fitness**) that will vary from case to case. The eventual code produced in LabVIEW (Appendix A) was based in part on routines in MATLAB provided by Reading University during collaboration, although the underlying principles are widely known. LabVIEW was chosen for the implementation because it makes it easier to integrate the code with the control software developed for the LCM. Additionally, the act of translation also ensures a deep understanding of what can be quite abstract code, and indeed a couple of errors in the original MATLAB code were uncovered and corrected during this process. Finally, subsequent testing of the algorithm (Section 2.6.1) demonstrated that the LabVIEW implementation of the algorithm ran faster than the MATLAB version.

### Individuals

As described in Section 2.4.2 the LCM of the pulse shaper contains two masks consisting of two liquid-crystal arrays of 640 pixels. Each pixel is driven by a voltage with 12-bit resolution, i.e. the drive level applied to each voltage is an integer between 0 and 4095. Therefore for the purposes of the genetic algorithm one way of defining an individual is as a string of 1280 numbers, each having a value of between 0 and 4095. This will then cover all possible settings of the LCM (the universe). (Note that for phase-only pulse-shaping, both masks should have the same settings. In such cases individuals of 640 numbers will suffice.)

### Fitness

The question of fitness is determined by the experiment in hand, what can be measured and the desired outcome. For instance, if the pulses act with some form of detection circuit that generates a voltage, the voltage is influenced

by the pulse shape and a maximised voltage is required, then the measured voltage provides a simple fitness evaluation for each individual. More complicated schemes might call for producing a particular spectral output, in which case some form of error function between the wanted spectrum and the measured spectrum could suffice.

### 2.6.1 Testing of the algorithm

To first test the ability of the algorithm to obtain solutions, individuals consisting of 50 elements having the range 0–99 were used. The requested solution was a simple one: The sum of all the elements in the individual should be zero. A population size of 100 with 10% offspring was used (i.e. 10 offspring per generation) and the *threshold* (the fitness at which the algorithm stops seeking) set to zero, i.e. a perfect solution. When run with these parameters the GA found the ideal solution in 9740 generations. A plot showing the evolution of the algorithm is shown in Figure 2.19. What is seen here is typical behaviour — initially large improvements are made very quickly, but as the population evolves the improvements become gradually more incremental. This is the main purpose for using a threshold value in normal experiments; to decide a cut-off point at which a solution is deemed good enough. Any further increases beyond this will normally provide small return compared to the time invested.

The linear plot is difficult to interpret by eye, so instead a logarithmic plot is preferable. Figure 2.20 shows the same data plotted on a log scale. This makes other features more clear, such as ‘plateauing’, where no improvement is seen for a number of generations before some chance mutation results in a number of quick improvements. Such features will be seen in the physical experiments described in later sections (Figure 2.38).

The speed and efficiency of the LabVIEW code compared to the MATLAB version was also checked to confirm the translation was both valid and appropriate. Using the settings described above, each version was given five runs and the number of generations required to reach a solution recorded.

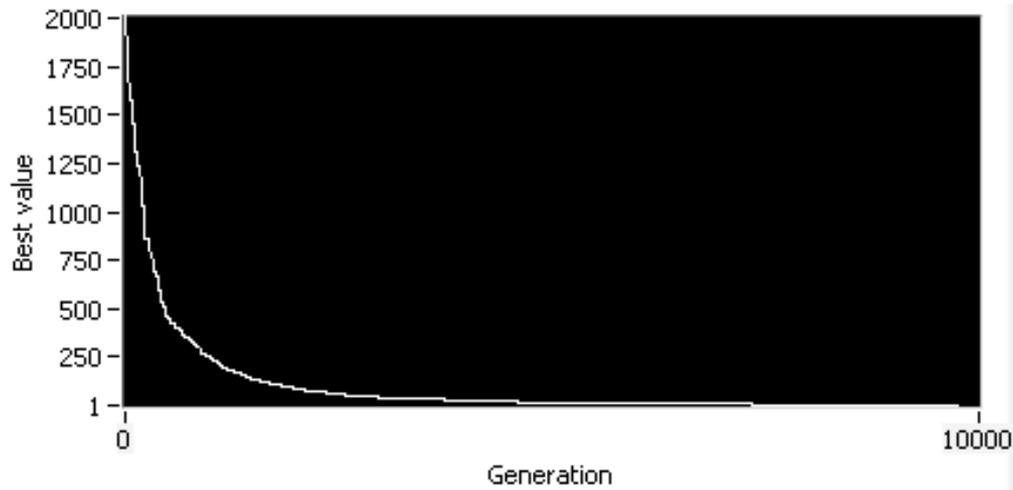


Figure 2.19: Plot showing the evolution of the genetic algorithm towards a solution. The value of the fittest individual in each generation is plotted against generation number. The vertical scale here is linear.

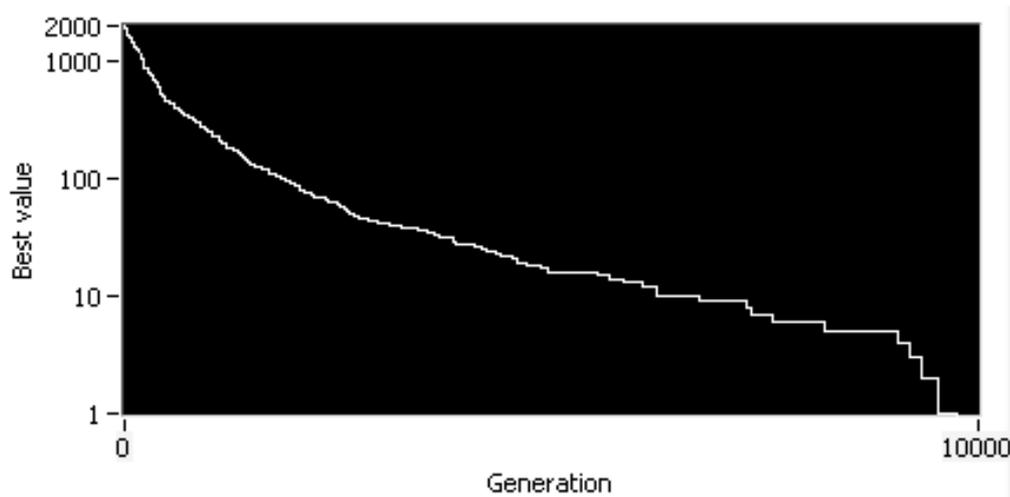


Figure 2.20: Plot of Figure 2.19 using a logarithmic scale for clearer viewing.

Version	Generations	Version	Generations
MATLAB	10166	LabVIEW	10133
	10019		8766
	7251		10652
	9451		9433
	5701		8144

It is clear that the variable elements of the genetic algorithm mean that the number of generations taken to reach a solution can vary widely, even when using the exact same software. On the small sample above, there is no real evidence that, as far as the number of generations required is concerned, the algorithm performs any better or worse in LabVIEW than in MATLAB.

As the algorithm is the same in each case, no difference should be expected so, rather than increase the number of samples to confirm this the issue of actual processing time was considered. Each algorithm was slightly adapted to record the time taken for a run from start to finish, and amended where necessary so that no data or information was plotted to the screen. Both algorithms were run on the same PC (Pentium 4 2.8 GHz with hyperthreading enabled, 2.5 GB of RAM)<sup>12</sup>. The results are shown below.

Version	Gen.	Time (s)	Time/gen. (ms)
MATLAB	6643	23.92	3.60
	5823	21.13	3.63
	5866	20.97	3.57
	7177	23.94	3.34
	11725	43.55	3.71
Version	Gen.	Time (s)	Time/gen. (ms)
LabVIEW	9367	22.02	2.35
	10382	24.98	2.41
	12325	29.74	2.41
	9237	22.33	2.42
	9637	23.27	2.41

---

<sup>12</sup>For those of you reading this some years from now, yes, it does seem like an antiquated machine!

Again the wide variation in the number of generations taken, and thus the overall time taken, can be clearly seen. However, the time taken to evaluate each generation is far more consistent, especially in the LabVIEW implementation. Also it is quite clear that the LabVIEW version processes each generation more quickly, in roughly two-thirds of the time in fact. Given that the number of generations taken to reach a solution is expected to average out as the same for both versions, the LabVIEW version can be expected to be faster, all else being equal.

Before implementing the genetic algorithm together with the pulse shaper it is necessary to have an understanding on the most appropriate starting parameters. These parameters are the population size and the number of offspring produced each generation. Although, theoretically, a correctly-designed algorithm should always converge on an optimal solution this is given an infinite amount of time. As the experimenter does not have an infinite amount of time to run an experiment the initial parameters should be set so as to reach a satisfactory solution within some reasonable time frame. It is found that the actual times taken for processing are small, as in the timings given for the previous example in the table below.

Module	Number of runs	Average time (ms)
Recombine	9741	0.760
Mutate	9741	0.433
Select	9741	0.372
Ranking	9741	0.332
Reinsert	9740	0.091
Fitness	9742	0.006

The types of experiments for which the pulse shaping system is intended are mainly of the electro-optic sampling type (Section 3.5.1), where each evaluation will require a number of mechanically-delayed probe pulse measurements and thus will take far longer than a millisecond. Thus it will be the number of evaluations of individuals that determines the overall experimental time, and it is this which should be minimised. For this reason a number

of tests were carried out to determine the relationship between population number, offspring proportion and the total number of evaluations.

### 2.6.2 Minimising the evaluations

The first test was to determine how the offspring proportion affects the number of generations required to reach a result. To more closely represent the environment that the algorithm will find in conjunction with the pulse shaper, individuals were used consisting of 1280 elements each ranging from 0–4095. The same ideal solution was used as previously, i.e. sum of values = 0. Because a threshold of 1 would take prohibitively long (in one test no solution was found after two hours of searching and over 200,000 generations) a higher threshold is used. In this case it is the threshold that indicates the average element is  $\leq 1/3$  of the maximum — the actual threshold value is unimportant, as long as it is the same for each test.

A population of 100 was used and various offspring proportions used ranging from 5% to 90%. For each proportion the algorithm was run eight times and the number of generations recorded. Standard deviations of the number of generations taken were found to be 3–6% of the mean with no clear relation to the proportion value. The results are plotted in Figure 2.21. It can be seen that increasing the offspring proportion significantly reduces the number of generations required to reach a result, up until a proportion of around 40–50% (i.e. half the population replaced each generation). Beyond this the improvements are negligible or non-existent.

As has been pointed out, though, from the point of view of the proposed experiments it is the number of evaluations that will determine the overall time taken. Since the GA must evaluate each new individual the total number of evaluations carried out (ignoring the evaluation of the original population) is given by multiplying the offspring proportion by the population (100 in this case) and by the number of generations taken to reach a solution. This is plotted in Figure 2.22, indicating a proportion of between 10% and 30% is optimal.

The other parameter that will influence the number of evaluations re-

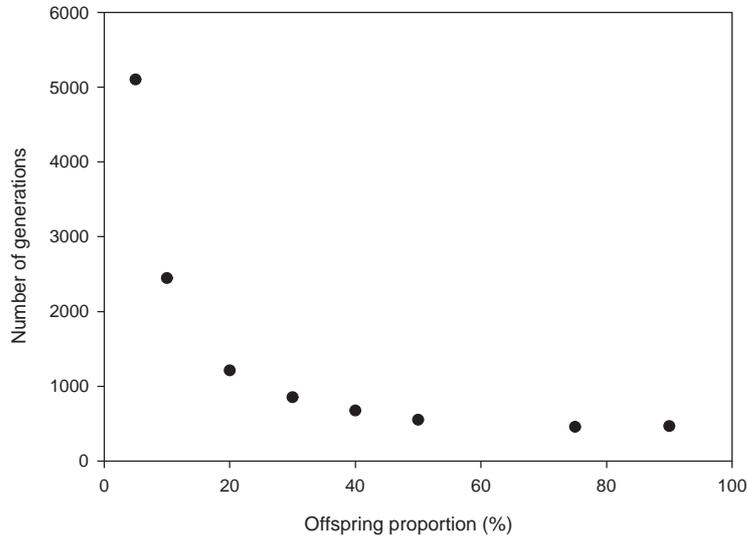


Figure 2.21: Plot showing how the number of generations taken to reach a result varies with the proportion of children introduced each generation.

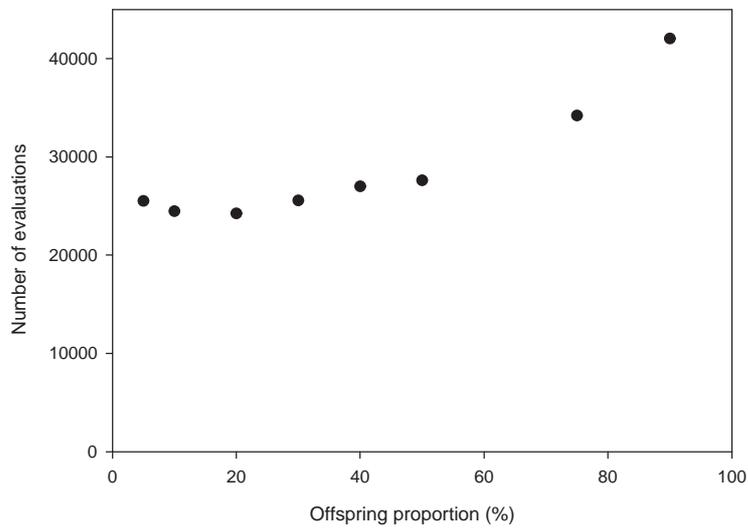


Figure 2.22: Plot showing how the total number of evaluations made by the GA to reach a result varies with the proportion of children introduced each generation.

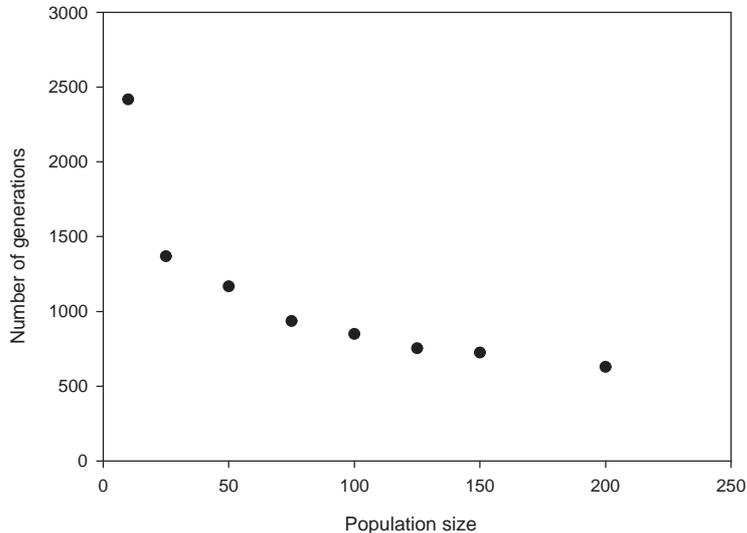


Figure 2.23: Plot showing how the number of generations taken to reach a result varies with population size.

quired to reach a result is the population. Clearly for any given offspring proportion a larger population will mean more evaluations to be carried out each generation, but a larger population should also mean less generations required. To determine these relations an offspring proportion of 30% was used (fitness timing not being an issue during testing) with varying populations from 10 to 200. Again for each population a number of repeat runs of the GA were used and the average taken. The same threshold was used and the standard deviations were found to vary from 2.5–6.5% of the mean with no clear relation to the population size. The results are plotted in Figure 2.23.

It can be seen that larger populations reduce the number of generations required to reach a result and that this effect decreases with increasing population size. Again, though, it is the number of evaluations made that is of more importance. This is calculated from the offspring proportion (30%) multiplied by the population size multiplied by the number of generations taken to reach a solution. This is plotted in Figure 2.24

This result is altogether more surprising. It is clear that although a

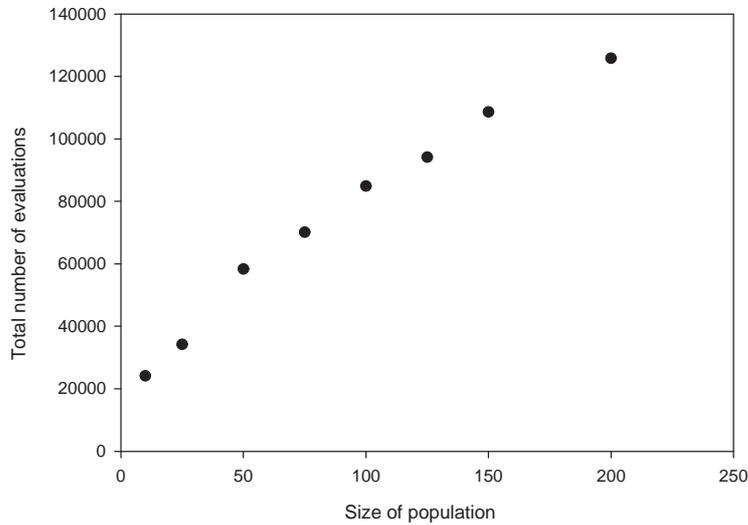


Figure 2.24: Plot showing how the population size affects the number of evaluations the GA is required to make to reach a result.

larger population reduces the number of generations taken to reach a result, this is more than offset by the correspondingly larger number of evaluations that must be made at each generation. This looks almost linear, suggesting that the greater benefit is provided by making the population as small as possible. Clearly there must be some cut-off point (a population of 1, for instance, could not even breed) but it can be seen that even a population of 10 in which only 3 offspring are produced each generation results in a lower total number of evaluations overall. This is counterintuitive to much of the standard thinking related to GAs, where the focus is on reducing the number of generations required to reach a solution rather than considering the number of evaluations.

In conclusion, then, it has been found that although large populations and high offspring proportions will provide solutions more quickly in terms of numbers of generations, in situations where the time taken to evaluate the fitness of an individual is significant a relatively tiny population and small percentage of offspring will return a solution more quickly in terms of actual time.

## 2.7 Practical application of the pulse shaper

### 2.7.1 Experimental outline

A suitable experimental test of the combined pulse shaper and genetic algorithm required a system based on ultrafast optical pulses and with an output that could depend on the pulse shape and would provide a scientifically-interesting result. An appropriate possibility was identified based on apparatus located in the same laboratory as the pulse shaper. This apparatus had been developed at NPL for generating terahertz pulses with the aim of developing wavelength standards.

#### **Terahertz**

Terahertz (THz) is the name given to a section of the electromagnetic spectrum ranging in frequency from roughly 0.1 THz to 10 THz, i.e. the region between the far-infrared and millimetre waves. This region is of considerable interest as THz radiation has a number of useful properties, mainly as a spectroscopic tool or as an imaging method for medical or security applications. This is a relatively new area of research due to the only recently developed methods of generating THz radiation easily. The principal methods are based around ultrafast lasers, converting the energy into THz frequencies through either optical rectification in a suitable nonlinear crystal or photoconductive effects in semiconductor-based devices. (Generally speaking, the broader the spectrum and the narrower the  $\Delta\tau$  of the optical pulses, the greater the bandwidth of the generated THz pulses.) This use of ultrafast pulses suggests an application of pulse shaping to alter the pump optical pulse in such a way as to influence the resulting THz radiation generated and some work has indeed been done in this field[15]. More complicated or involved control of the THz radiation requires a detailed understanding of the relationship between the pumping optical pulse and how it reacts with the medium used to generate THz wavelengths, an understanding which is not always present. In this situation adaptive control seems a useful tool as is borne out by some papers[16], but currently no application of a genetic algorithm is known to

have been used in this context.

A THz scheme exists at NPL that has been designed and developed for the purposes of creating wavelength and spectroscopic standards. A schematic is shown in Figure 2.25 and is typical of many such arrangements. A Ti:S laser source<sup>13</sup> provides ultrafast pulses with  $> 120$  nm bandwidth. These ultrafast pulses are focused onto a ZnTe crystal to generate THz pulses that are collimated and then focused onto a sample under test by a parabolic mirror pair. The focused THz pulses then interact with the sample before an identical mirror pair collect and focus the THz pulses onto a second ZnTe crystal. The purpose of the second ZnTe crystal is to perform a temporal measurement of the THz radiation using electro-optic sampling (EOS - described more fully in Section 3.5.1) using probe pulses created by splitting the output of the ultrafast laser. Changes in the planar polarisation of the probe pulses, proportional to the THz field present at the point of interaction, are measured by a differential detector. Lock-in amplification and an optically-chopped pump beam improve the signal to noise. An optical delay line changes the arrival of the THz pulses relative to the probe pulses by altering the path length of the pump pulses, enabling a temporal waveform representing the THz pulse to be built up (for example, Figure 2.26). The corresponding spectrum is then obtained by performing a Fast Fourier Transform (FFT) (see Figure 2.27).

By measuring spectra with and without material samples in the THz beam path the absorption characteristics or THz ‘fingerprint’ of different materials can be established, enabling their identification in other THz detection systems, such a security scans. In order to check the accuracy of these spectral signatures some set of known standards is required, standards that absorb in specific regions of the THz spectrum and enable such a system to be calibrated. This is one of the principal uses of this system at NPL, to develop wavelength standards in the THz frequency range. A measurement of a filter that could form such a standard has been provided by NPL in Figure 2.28, consisting of a metallic structure that was found to absorb around 0.7 THz.

---

<sup>13</sup>Femtolaser Scientific Pro

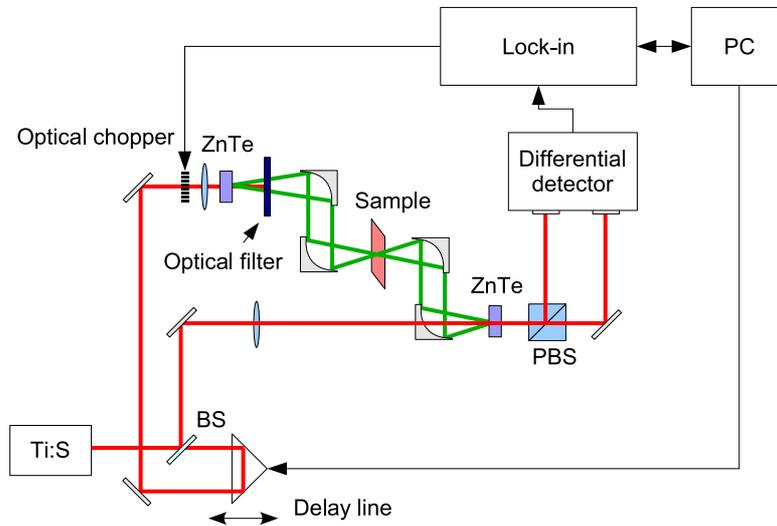


Figure 2.25: Schematic of the NPL THz system. Full details are given in the text. The probe beam is focused onto the ZnTe through a hole drilled in the fourth parabolic mirror.

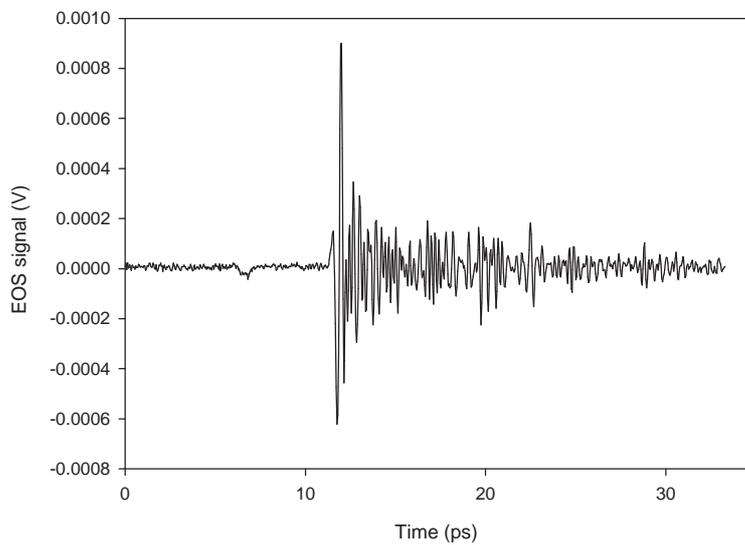


Figure 2.26: Electro-optic sampling was used over a stepper translation of  $5000 \mu\text{m}$  to produce this temporal measurement of a THz pulse. The speed of the scan was  $0.1 \text{ mm/s}$  and the sample rate  $8 \text{ Hz}$ . The time constant of the lock-in amplifier was  $30 \text{ ms}$ .

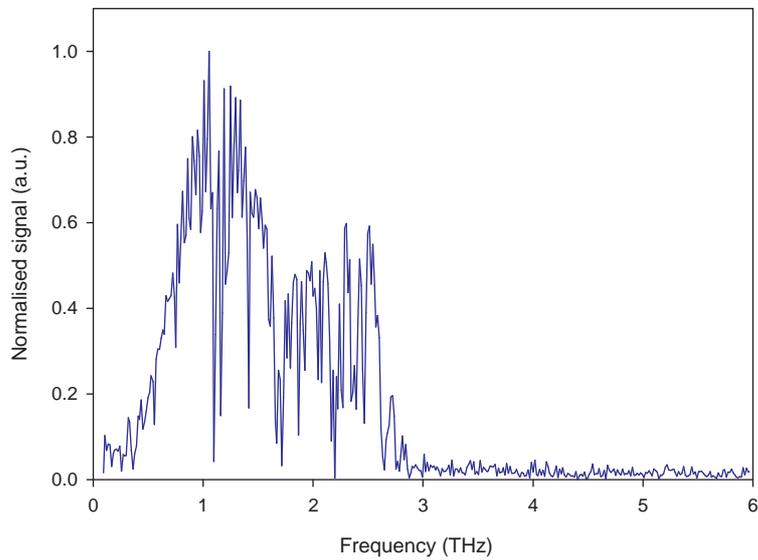


Figure 2.27: Spectrum of the pulse from Figure 2.26 obtained using an FFT. Spectral content can be seen up to 2.5 THz. The minima are mainly due to water absorption.

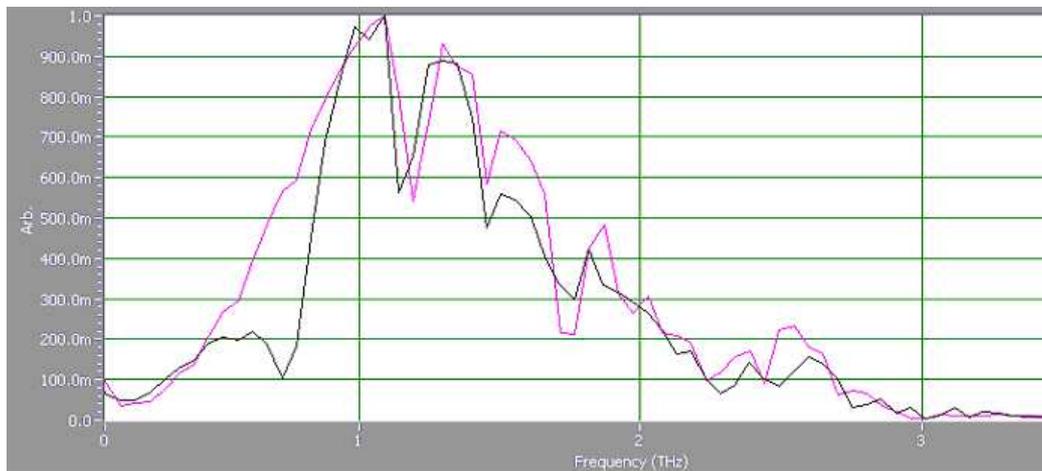


Figure 2.28: Plot of the THz spectrum with (black) and without (pink) a metallic filter structure. This spectrum is computed from the FFT of a time domain measurement. Note the absorption around 0.7 THz, indicating filtering at this frequency. (Plot provided by the THz Section at NPL.)

Standards such as these could be used to provide calibration for THz systems but there are a number of issues. Such standards require to be transferred to the system being calibrated, and there can be issues with maintaining the correctness of the standard during such a move as well as the risk of damage. Also these standards need to be compared to some reference and that itself will require calibration, making it difficult to establish a golden standard without defining it arbitrarily. For reasons such as these, and given that the THz pulses are generated from ultrafast pulses, it was proposed to investigate the possibility of introducing spectral gaps in the THz spectrum by shaping the ultrafast pump pulses. If these could be used to influence the spectral features, such shaping could provide a non-physical means of establishing a calibration standard for a system that could in turn be used to calibrate lesser, physical standards.

The relationship between the pump pulse and the THz pulse can be complicated, and the calculation of the optical pulse required to produce a given THz spectra may be difficult, time consuming, and even impossible. The proposal therefore was to instead use the genetic algorithm to search for the required settings for the pulse shaper. The input to the algorithm would be the THz spectra and the fitness of each mask given to the pulse-shaper established by comparing the THz spectrum produced to the desired outcome.

### **Proposed experiment**

Figure 2.29 illustrates the proposed experimental set-up to investigate the changing of THz spectra using pulse shaping of the pump optical pulses. The THz layout is generally as described in Section 2.7.1. The pulses from the source ultrafast laser are divided by a beamsplitter to provide both pump and probe pulses. The pump pulses are taken from the reflected portion to limit dispersion before the pulse shaper. In contrast to the original set-up, it is the probe pulses that pass through the optical delay line. The reasons for this are mainly due to the layout: Sending the pump pulses via the delay line as well as the pulse shaper would require an excessive amount of corresponding delay to be added to the probe pulse.

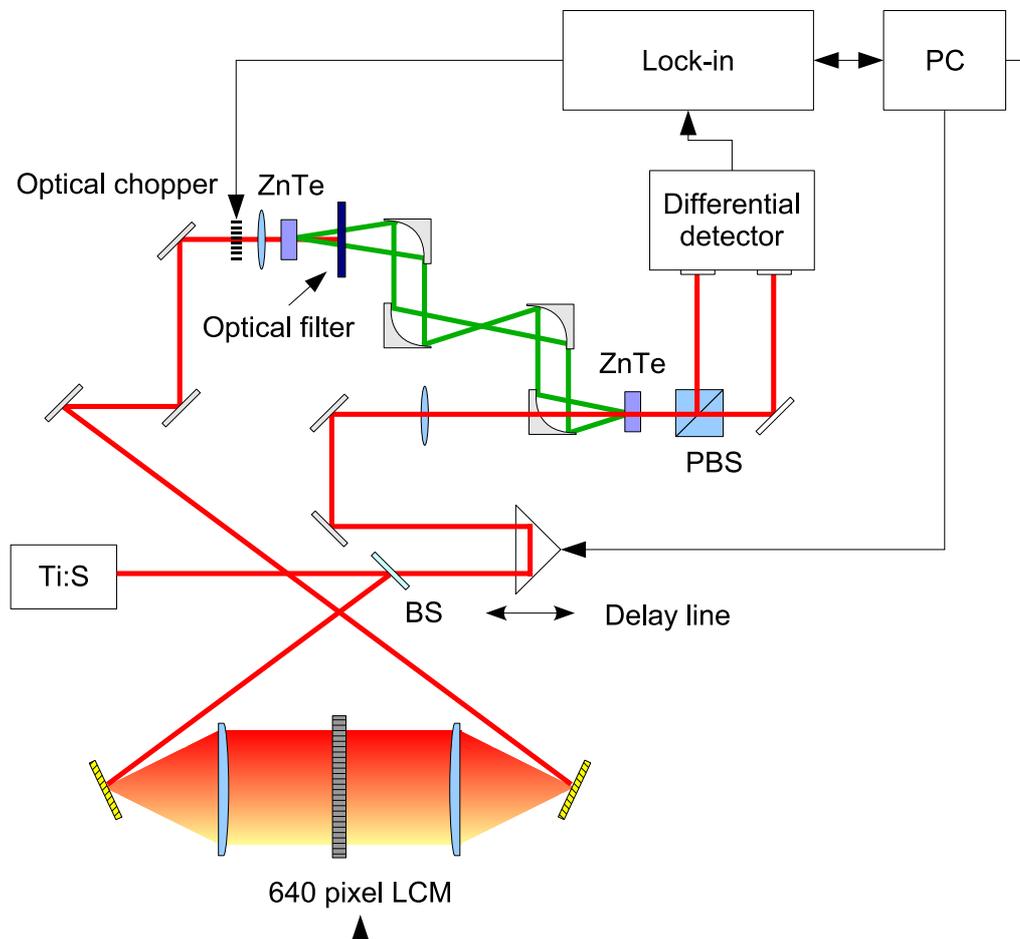


Figure 2.29: Schematic of the integrated THz and pulse-shaper layout. The pulse shaper influences the THz pulse indirectly through shaping of the pump pulse. A PC monitors the THz shape and alters the LCM setting according to a genetic algorithm.

Initial measurements were carried out without the LCM present to confirm the generation of THz pulses and to ensure that the spectra were as expected. (The reason for removing the LCM was so that there would be no issue as to whether errors in the LCM amplitude and phase settings were contributing to unexpected results. The grating/lens pairs should have no effect if set correctly.) A typical THz measurement is shown in Figure 2.30. It is apparent that the spectral power with the pulse shaper is significantly

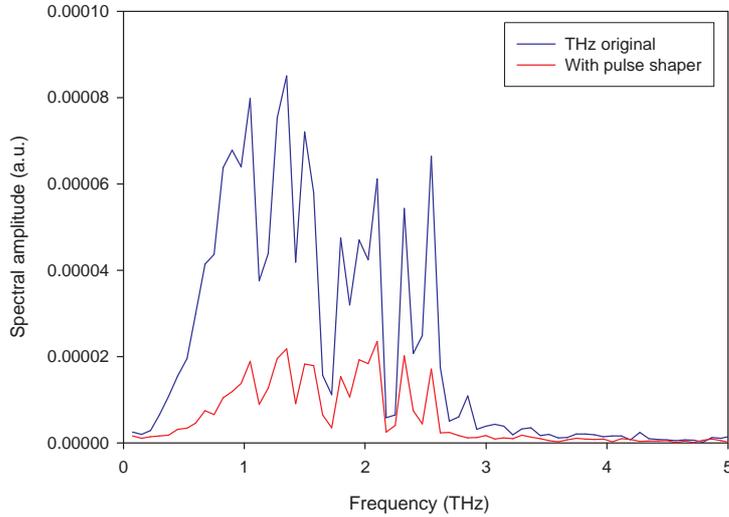


Figure 2.30: Plot showing the THz spectrum obtained with the pulse shaper in place (but no LCM) compared to the spectrum produced by the original system.

lower than in the original system. (The approximate value is 25% of the original.) This is mainly due to the loss from the first grating, where only the first diffracted order is used. To confirm whether the actual spectral content is the same the spectra are normalised to each other (Figure 2.31). Here it can be seen that the approximate shapes are the same, suggesting the (unshaped) pulses produced by the system with the pulse shaper in place are comparable to those of the original system except for a reduction in power.

## 2.7.2 Experimental setup

The software was written and tested so that genetic algorithm control of the pulse shaper, measurement of the THz pulse and calculation of the THz spectrum by FFT would all be integrated. The time taken to evaluate an individual supplied by the algorithm is determined primarily by the time taken to make a THz EOS scan. This can be adjusted within reason, in that the size and speed (i.e. the resolution) of the scan can be set. Since

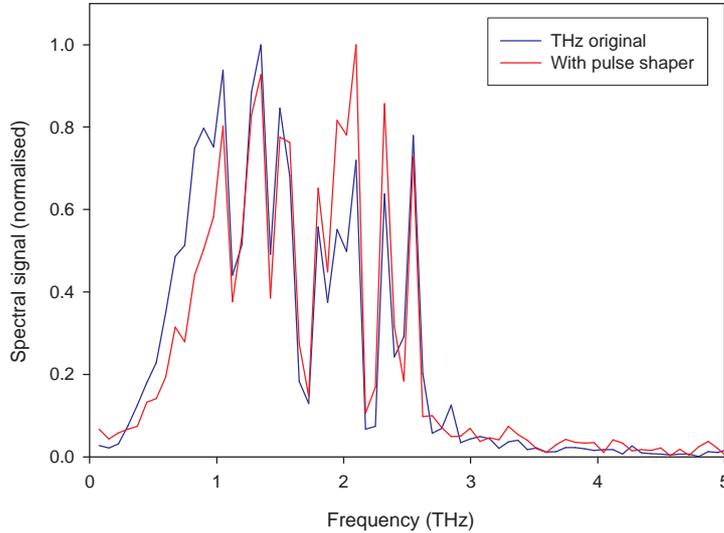


Figure 2.31: Figure 2.30 with the traces normalised for better comparison of the spectral content.

the frequency and time domains are related through a Fourier transform, decreasing the time taken by reducing the size of the scan will reduce resolution in the spectrum. Increasing the speed will reduce the spectral width that can be accessed. Therefore it is necessary to balance the required spectral range and detail with the need to reduce the time taken for the genetic algorithm to reach a result within a reasonable time. After a number of tests a scan size of  $2500 \mu\text{m}$  at a speed of  $0.2 \text{ mm/s}$  was decided to be an acceptable compromise. An example measurement taken using these settings is shown in Figure 2.32.

For the algorithm individuals of  $1280 \times 4096$  were deemed to produce too large a universe and provide excessive detail for the current level of work. To reduce the universe size the element levels were set to 7-bit (i.e. 128 possible values, multiplied by 32 when applied to the pixels). It was determined by inspection that the spectrum from the laser source only filled 540 of the 640 available pixel pairs so the individuals were set to 540 in length with extra values added as padding when uploaded to the LCM. Finally, these initial

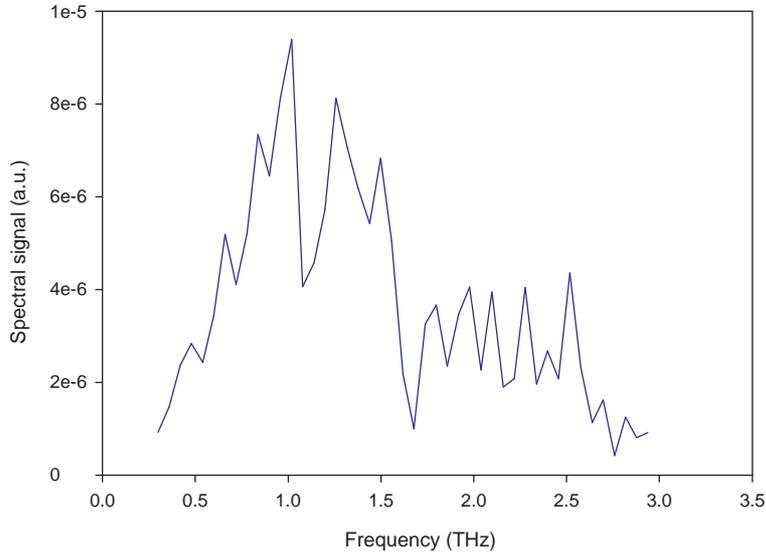


Figure 2.32: Example spectrum taken with the LCM in place using the same EOS scan parameters as for the genetic algorithm. This provides a balance between speed of acquisition and spectral resolution.

experiments were carried out using phase-only pulse shaping (both masks loaded with the same data) to further reduce the universe. Given that it was established (Section 2.6.2) that the smallest number of evaluations to reach a given solution occurs when the population and the offspring proportion are comparatively low, a population of 20 with an offspring proportion of 30% was chosen. It was soon found that even with the various settings chosen as best to minimise the number of evaluations made in each generation, the process was time-consuming. A progress rate of 7–8 generations per hour was typical and runs were normally made overnight or over weekends. When evaluating a result returned at the end by the GA, care was always taken to compare it to a measurement made at the same time with a neutral setting of the LCM to eliminate changes in laser output as the cause of differences. Due to the time taken for the population to evolve a stopping point was usually decided by laboratory time constraints rather than any threshold value.

### 0.7 THz filtering

The first test was to produce a THz spectrum with a hole at 0.7 THz, similar to the physical filter of Figure 2.28. A perfect result would be a spectrum identical to the unshaped spectrum but with a gap at 0.7 THz, and one way to aim for this would be to gauge each mask by performing a root-mean-square subtraction of the measured spectrum from the desired. However, this can be a severely limiting solution and it is not even guaranteed that the THz spectrum can be controlled in this way. Instead the fitness was evaluated by comparing the spectral power at 0.7 THz with that at 0.5 THz and 0.9 THz with the aim being to maximise the differences between them (in the sense of reducing the 0.7 THz power). To reduce the effects of noise the evaluation for each individual was made by making five EOS measurements, averaging them and then performing an FFT. The FFT was normalised before the relative positions of the 0.7 THz to the 0.5 THz and 0.9 THz points were compared. The reason for normalising the FFT is because the effect wanted is suppression of the 0.7 THz as much as possible. If un-normalised FFTs were evaluated a spectrum with 60% reduction at 0.7 THz could be rated as better than one with 100% reduction but half the power. The GA settings were the same as given in Section 2.7.2.

Figure 2.33 shows the result from a 135 generation run. The spectra from the reference (unshaped) and mask (shaped) pulses have been normalised to each other.

A significant dip in the spectral power can clearly be seen but the dip is at 0.75 THz and not 0.7 THz as requested. The spectral power at 0.7 THz is in fact close to the unshaped result. Inspection of the code used to determine the fitness revealed the cause of the error. The spectrum was stored as an array and the elements representing 0.7 THz, 0.5 THz and 0.9 THz extracted. However, the incorrect index was used to address the array and so the points 0.75 THz, 0.55 THz and 0.95 THz were used instead with the fitness routine assessing the individuals according to their ability to minimise the 0.75 THz. In this respect, then, the actual result is a success.

The code was amended to correctly assign fitness according to the relative

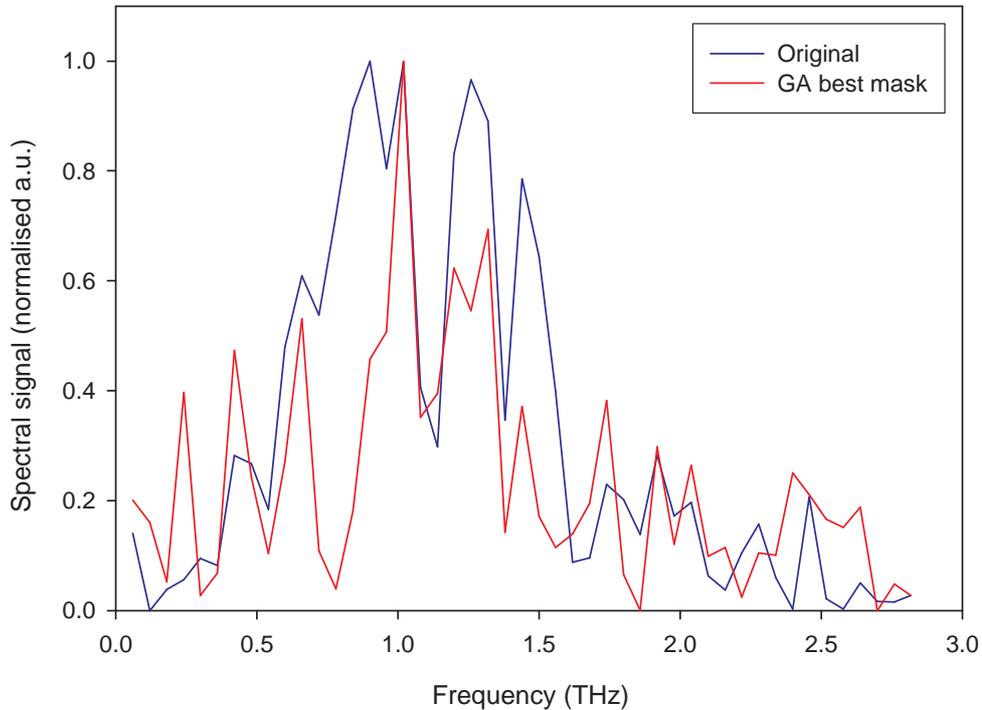


Figure 2.33: Comparison of the best mask returned by the genetic algorithm after 135 generations when attempting to produce a spectral hole at 0.7 THz. The hole is seen at 0.75 THz instead. This was found to be due to a programming error.

strength of the 0.7 THz signal and run again. The result from a 171 generation run is shown in Figure 2.34. Here a definite dip can again be seen, but this time centred at 0.7 THz as requested.

### 1.5 THz filtering

To test the flexibility of the algorithm, it was run again but this time with the request for a dip at 1.5 THz. After 263 generations the best mask returned was used to produce the spectrum shown in Figure 2.35.

No dips are seen this time at 0.7 THz, suggesting that the results of the previous section were not a coincidence. Moreover, a clear dip is seen at

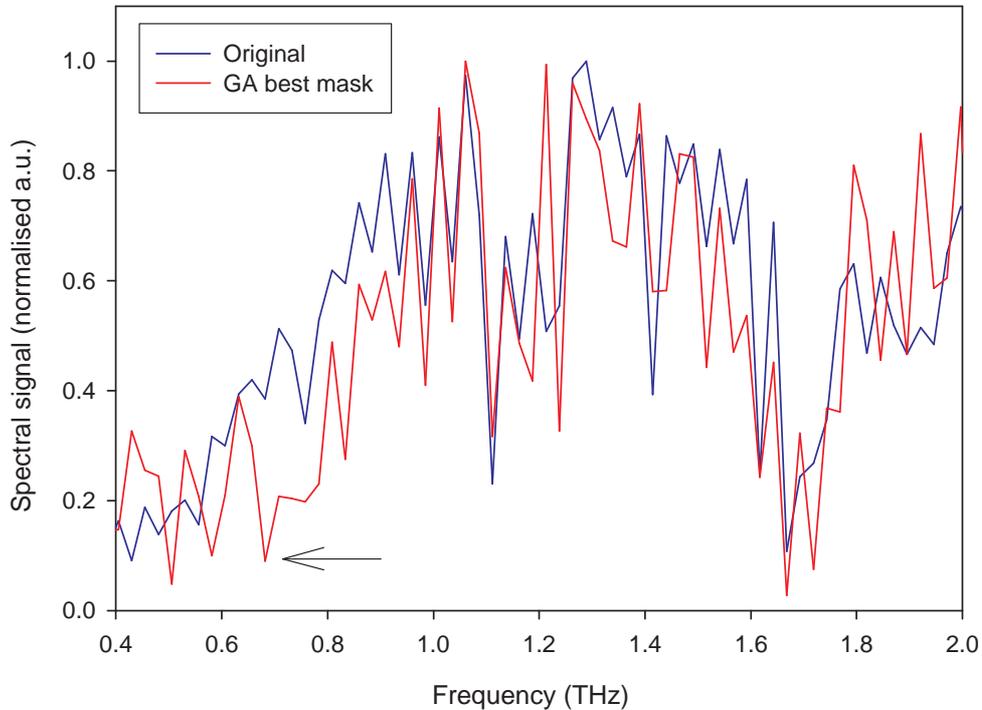


Figure 2.34: Comparison of the best mask returned by the genetic algorithm after 171 generations when attempting to produce a spectral hole at 0.7 THz (indicated). A longer epoch was measured in the time domain to increase the spectral resolution.

1.5 THz, again as requested. The region between 0.9 THz and 1.3 THz looks lower as a whole, however, so the shaped result was rescaled to confirm the presence of the dip at 1.5 THz. This is shown in Figure 2.36, and demonstrates that the spectral power at 1.5 THz has indeed been reduced as a result of pulse shaping.

### 0.7 THz maximisation

The ability of the pulse shaper to increase power at a specified wavelength was also investigated. When deciding how to program the fitness routine appropriately the first approach was to simply maximise the 0.7 THz value.

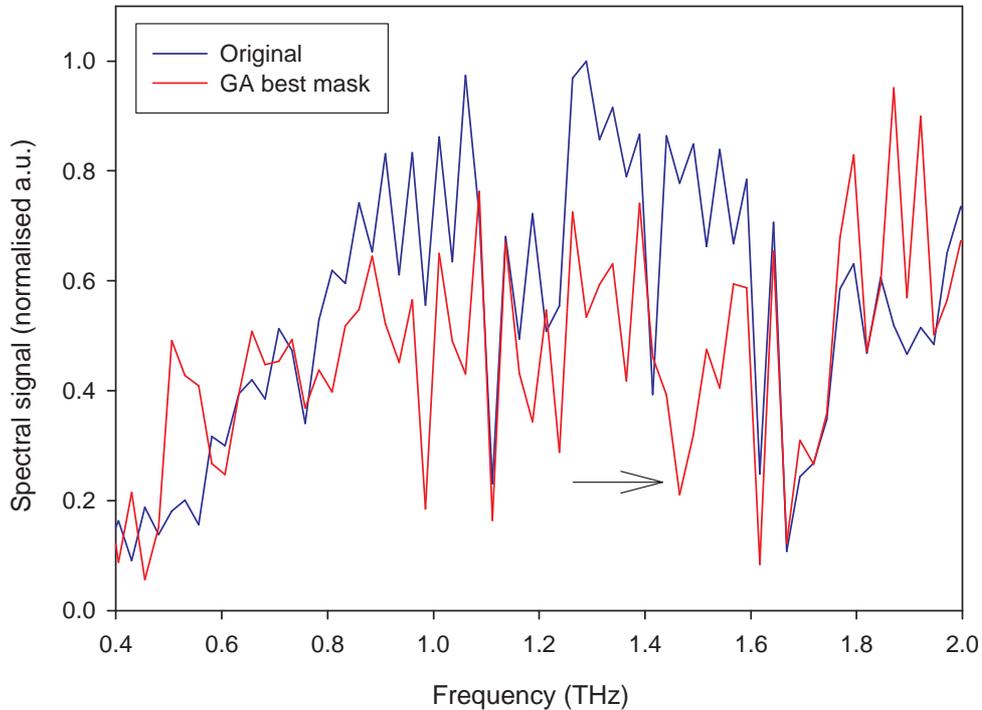


Figure 2.35: Comparison of the best mask returned by the genetic algorithm after 263 generations when attempting to produce a spectral hole at 1.5 THz (indicated).

This was not found to produce any suitable solutions, and it was speculated that, given that the spectrum is normalised, spectra where the THz power was minimal and thus dominated by noise could be mistaken for suitable solutions. Removing the normalisation (i.e. inspecting the absolute power at 0.7 THz) would be unlikely to provide a mask where there was any noticeable enhancement at 0.7 THz relative to the other powers. In this case the original spectrum, for instance, would be evaluated with higher fitness than a slightly lower powered spectrum with a strong peak at 0.7 THz.

As an alternative, the routine was amended to ascribe fitness to increased differences between 0.7 THz and 0.5 THz/0.9 THz in the favour of 0.7 THz. (Effectively a sign reversal of the evaluation used in Section 2.7.2.) The best

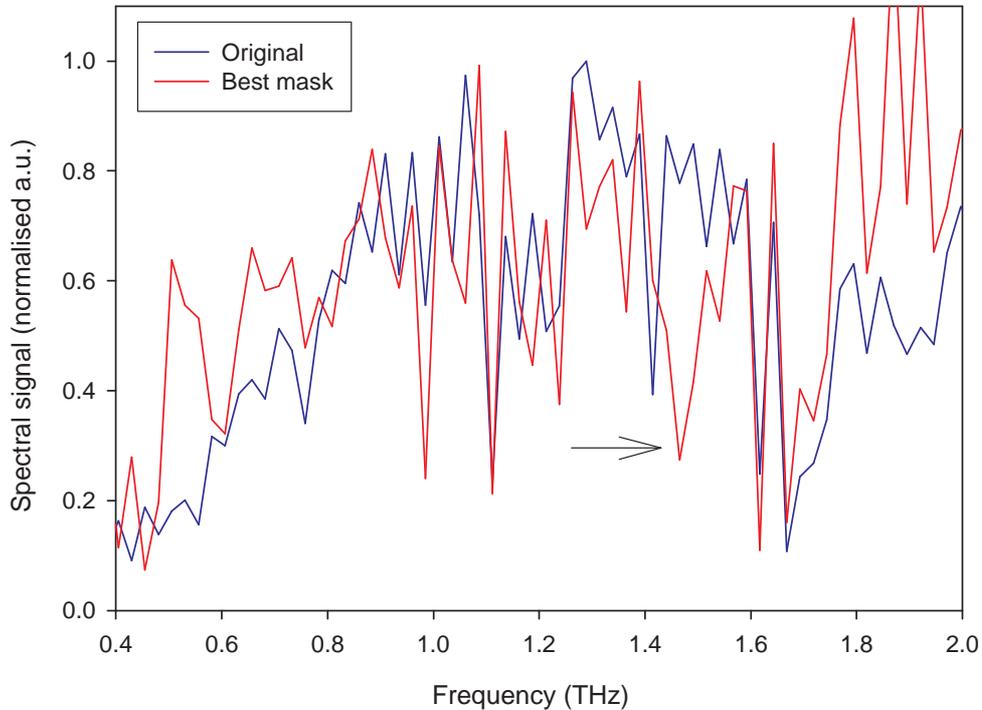


Figure 2.36: Figure 2.35 rescaled to match the powers in the 0.9–1.3 THz range. The reduced power at 1.5 THz (indicated) is still apparent.

mask achieved after 301 generations was used to produce the spectrum seen in Figure 2.37. It seems that the routine *has* done what was requested of it, but the result is not quite the effect hoped for. Rather than produce a strong peak at 0.7 THz, the spectral power at 0.9 THz has been reduced instead.

## 2.8 Conclusions

An optical parametric oscillator has been designed and constructed on an optical breadboard for the production of 1500 nm ultrafast pulses. Based around a PPKTP crystal pumped by a Ti:S laser the OPO was found to produce pulses tunable from 1500 nm to 1550 nm with an average power of around 80 mW. An measurement of the OPO autocorrelation indicated that

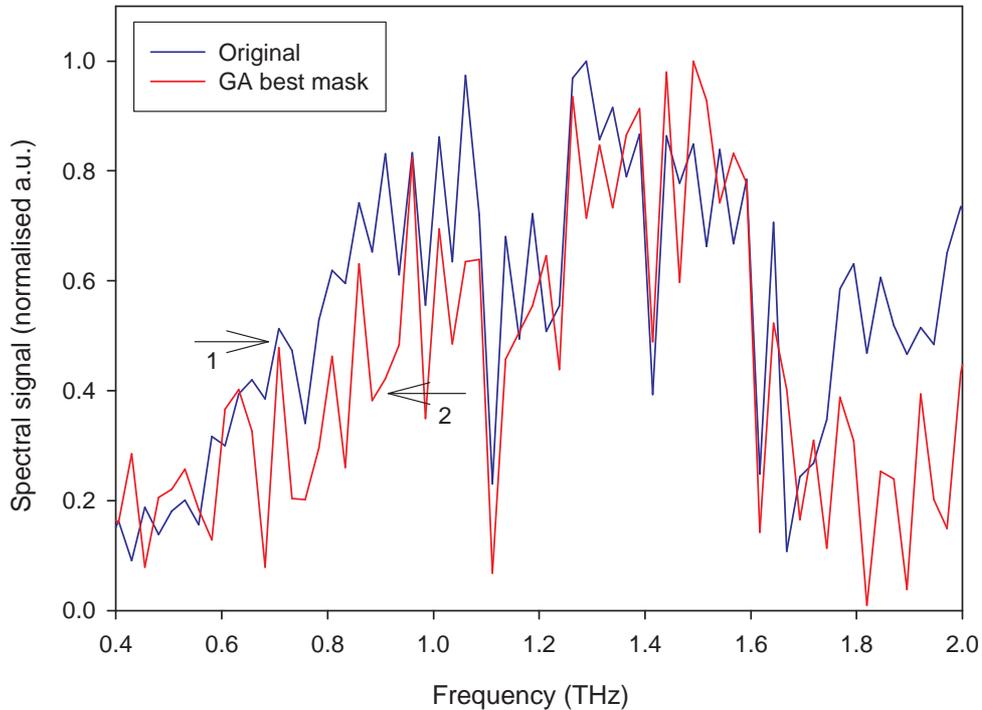


Figure 2.37: Comparison of the best mask returned by the genetic algorithm after 301 generations when attempting to produce a spectral peak at 0.7 THz (1). The effect instead has been to suppress the 0.9 THz (2).

the output had a  $\Delta\tau$  of 150 fs assuming Gaussian output. The original OPO scheme incorporated a SHG crystal for producing 750 nm pulses and a prism pair for dispersion compensation. The SHG crystal was removed when it was found that mounting it externally would produce more appropriate powers in the fundamental and second harmonic for the specific experiment proposed. The prism pair was omitted when oscillation proved impossible to achieve with the prisms in place.

A  $4f$  pulse shaper was constructed based around a dual mask 640 pixel LCM and control software written for it. It was demonstrated to be capable of altering the spectral content of ultrafast optical pulses from a femtosecond laser and a genetic algorithm was then coded and integrated with the pulse

shaper software to enable the LCM settings to be optimised according to some external experimental parameter. Tests of the genetic algorithm found that, in cases where the evaluation of the external parameter is the most significant contribution to the run-time, the number of evaluations should be minimised by using a small population and a small percentage of offspring per population even if the number of generations taken to reach a solution is large. This pulse-shaping system was then integrated with a THz pump-probe system and a number of experiments carried out to investigate whether pulse shaping could be used to emulate physical filters. The results demonstrate that it is indeed possible to influence the THz spectrum by pulse shaping and in such a way as to introduce artifacts that resemble the effects of filters. This is the first time that genetic algorithms and pulse shaping have been applied to THz spectra in this way and is of interest to the wider terahertz community[17].

The difficulties encountered when attempting to maximise a section of the spectrum highlight the importance of choosing the correct method of evaluating fitness and also the correct implementation. Care should be taken to consider possible unwanted effects (for example, simply asking that the spectral power be set to a minimum at a desired wavelength could result in mask settings that totally prevent the production of THz pulses getting the highest evaluations).

One of the most crucial issues is the time taken to perform an optimisation. For pump-probe experiments such as the one described even relatively short, fast scans with a minimum of averaging can mean a couple of minutes are required to evaluate an individual. Even with only three new individuals per generation this can mean barely ten generations per hour, and when it is seen that a couple of hundred generations are generally required to produce a reasonable result it becomes clear that methods of speeding up the process are desirable if these results are to be improved on. Figure 2.38 illustrates the evolution of the genetic algorithm that produced the mask settings that produced Figure 2.35. As seen in modelling (Section 2.6.1) the greatest gains are seen at the start and progress then begins to slow. It is clear that there are large ‘plateaus’ of over fifty generations where no progress is made before

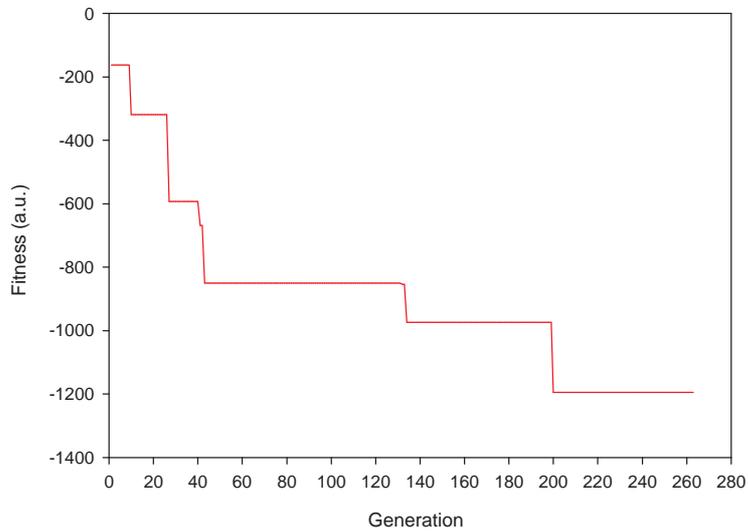


Figure 2.38: Evolution of the genetic algorithm towards a mask that emulates a 1.5 THz filter. Plotted is the fitness of the population's best individual against generation number. It can be seen that the evolution follows a pattern of long periods without improvement punctuated by quick gains.

another sharp improvement. If no further method exists for speeding up the evaluation, the alternative is to reduce the size of the universe that the algorithm can access to avoid the evaluation of useless individuals (Appendix B).

# Chapter 3

## Quantum Interference Control

### 3.1 Basic theory

*Quantum Interference Control* (QUIC) is one of the more exciting applications of ultrafast optical pulses to be developed in recent years. In simplest terms QUIC refers to the ability to influence two or more quantum processes in a system through causing an interference between them. For instance, consider a quantum system in an initial state  $|i\rangle$  for which there exist two different perturbations with Hamiltonians  $H_1$  and  $H_2$  that can take it to a final state  $|f\rangle$ . In general,

$$|\langle i | H_1 + H_2 | f \rangle|^2 \neq |\langle i | H_1 | f \rangle|^2 + |\langle i | H_2 | f \rangle|^2 \quad (3.1)$$

i.e. the response of the system when both perturbations are present is not equal to the sum of the responses to the individual perturbations. There is an interference term present which is dependent upon the relative phase between  $H_1$  and  $H_2$ , and so by controlling this phase relation it is possible to control the likelihood of the system moving from state  $|i\rangle$  to  $|f\rangle$ .

(The term *coherent control* is often used to describe this technique. However, some papers have also used this other term to describe different methods where a coherent source such as a laser affects some system. This thesis shall only use the term QUIC for clarity's sake.)

One of the earliest proposed applications involving QUIC was published

in 1967 by Manykin and Afanas'ev [18]. The authors considered the effects of coherent electromagnetic waves passing through a resonant medium and showed that a suitable superposition of waves could reduce the probability of photon absorption to such an extent that the medium became effectively transparent at that wavelength. In the case where waves were used with frequencies at  $\omega_1$  and  $\omega_3$  in a medium with response at  $2\omega_1$ , they found that the medium became effectively transparent when the following conditions were met:

$$3k_1 = k_3 \quad (3.2)$$

$$\frac{R_3 D_{\mu\nu}(-\omega, 3\omega)}{R_1 D_{\mu\nu}(\omega_3, \omega)} = 1 \quad (3.3)$$

$$\varphi_3 - 3\varphi - \pi = 0 \quad (3.4)$$

where  $R$  is the wave amplitude and  $D_{\mu\nu}$  are summations over dipole moments (defined in [18]), and  $k$  and  $\varphi$  are the wavevector and phase, respectively. If Equation 3.2 and Equation 3.3 are satisfied, then the remaining dependency is on the phase between the waves. By varying the relative phases the degree of absorption can be controlled, from opaque to transparent — quantum interference control.

Although Manykin and Afanas'ev discussed QUIC for a bulk macroscopic effect, much of the current interest has been in controlling atomic and molecular processes. Asymmetric photoelectron distribution in rubidium [19] was one early demonstration of experimental QUIC. Photoionisation of a rubidium atom was considered in the cases of light with  $\hbar\omega > \text{IP}$  (where  $\text{IP} =$  ionisation potential),  $\text{IP}/2 < \hbar\omega < \text{IP}$  and with bichromatic light — i.e. one-photon absorption and two-photon absorption, separately and together. In keeping with the simple-system approach, QUIC has been applied to quantum well systems [20] and gas phase molecules [21]. As with an atom, such systems lend themselves to QUIC due to the simplicity of the interaction. Each atom/molecule/quantum well can be considered a discrete unit in that its contribution to the output according to the process being controlled is independent of the other members of the system. The total output from the system is simply the sum of the outputs from each individual 'sub-system',

and interactions between the sub-systems that affect the QUIC process can be ignored. This approach underlay much of the initial work looking at QUIC in semiconductors discussed in the following section.

### 3.1.1 QUIC in semiconductor

Traditionally, photocurrents are produced by a two stage process in which photon absorption promotes an electron to the semiconductor conduction band and then a bias voltage applied across the semiconductor accelerates the electron in the desired direction. The vector of the photocurrent is therefore dependent on the voltage applied (and, to a lesser extent, on the properties of the semiconductor). With no voltage, the free electrons are generated with random momenta and no net current can result. Even if anisotropy is present in the semiconductor, the probabilities of photocurrent in the forward or backward direction of a given crystal axis will still be equal and the net current zero. In investigating QUIC for photocurrents, Kurizki *et al* [22] considered a superposition of the two bound states  $1s$  and  $2p_0$  in a donor atom. These states were then photoionised by two mutually phase-locked lasers operating at slightly different frequencies but with the same polarisation axes. This was therefore photoionisation of the superposition state

$$|\psi\rangle = c_1 |1\rangle + c_2 |2\rangle \quad (3.5)$$

by a  $z$ -polarised, two-colour source with electric field

$$E_z(t) = E_1 \cos(\omega_1 t + \phi_1) + E_2 \cos(\omega_2 t + \phi_2) \quad (3.6)$$

They found that the magnitude and the sign of the photocurrent in the  $z$  direction was dependent upon  $E_1$  and  $E_2$ ,  $N$  (the donor concentration),  $\omega_1$  and  $\omega_2$  and, most significantly, on the phase difference  $\phi_1 - \phi_2$ . Thus for a given piece of semiconductor and two suitable phase-locked lasers a photocurrent could be created and controlled simply by adjusting the relative phase difference between the lasers without the need for any external bias whatsoever. Using some reasonable numbers for a piece of Si measuring

$0.1 \times 10 \times 10 \mu\text{m}$  they predicted currents of 10 to 100 mA, enough to be of practical use and prompting suggestions of a number of possible applications, including relative phase measurement and information transmission by phase modulation.

Other suggestions for QUIC in semiconductors were put forward [23] but all did their best to consider small, atomically-analogous systems such as quantum wells or, as above, donor atoms. This seemed reasonable as all QUIC work had been carried out in the context of relatively simple quantum systems such as atoms or molecules. However, in their 1996 paper [24] Atanasov *et al* calculated one- and two-photon photocurrent injection for bulk intrinsic semiconductor. They found that when both processes were present simultaneously, the probability of an electron transition to the conduction band was asymmetric in momentum space  $\mathbf{k}$ . This asymmetry meant a net photoelectron momentum in a given direction and hence a net photocurrent without an external bias. The current density injection rate was calculated to be

$$\dot{\mathbf{J}}^I = \hat{\eta}:\mathbf{E}(\omega)\mathbf{E}(\omega)\mathbf{E}(-2\omega) + c.c. \quad (3.7)$$

where  $\hat{\eta}$  is a four-tensor calculated from the semiconductor band structure. This meant the possibility of QUIC-generated photocurrents in standard semiconductor devices.

### Calculation of injected QUIC current in GaAs

Writing Equation 3.7 in terms of its general components gives

$$\dot{J}_\mu = \eta_{\mu\alpha\beta\gamma} E_\alpha^\omega E_\beta^\omega E_\gamma^{-2\omega} \quad (3.8)$$

using  $E^\omega$  to denote  $E(\omega)$  for clarity. According to [24] the only non-zero values of  $\eta$  in LT-GaAs are  $\eta_{xxxx}$ ,  $\eta_{xyxy} = \eta_{xyyx} = \frac{1}{2}\eta_{xxxx}$  and  $\eta_{yyyz} \approx 0$ , these relations holding true for all exchanges of  $x, y$  and  $z$  due to the cubic symmetry of GaAs. Since  $\eta_{xxxx}$  is the largest component the authors recommended use of both beams polarised in the (100) direction in GaAs for optimum QUIC current generation, claiming that the (110) direction would result in only half

the injected current since  $\eta_{xyxy} = \frac{1}{2}\eta_{xxxx}$ . However, here it is demonstrated that this is not in fact the case, and that the (110) direction yields the *same* injected current as the (100) direction.

If the two beams are polarised in the  $xy$  plane so that  $E_z^\omega$  and  $E_z^{-2\omega}$  are both zero, calculating the injected current density in component form obtains

$$\dot{J}_x = \eta_{xxxx}(E_x^\omega)^2 E_x^{-2\omega} + \eta_{xyxy} E_x^\omega E_y^\omega E_y^{-2\omega} + \eta_{xyxy} E_y^\omega E_x^\omega E_x^{-2\omega} + \eta_{xyyx}(E_y^\omega)^2 E_x^{-2\omega} \quad (3.9)$$

$$\dot{J}_y = \eta_{yyyy}(E_y^\omega)^2 E_y^{-2\omega} + \eta_{yyxx} E_y^\omega E_x^\omega E_x^{-2\omega} + \eta_{xyxy} E_x^\omega E_y^\omega E_x^{-2\omega} + \eta_{yxyx}(E_x^\omega)^2 E_y^{-2\omega} \quad (3.10)$$

Taking the coordinate system where (100) is the  $x$  direction and (010) the  $y$  direction, having both beams polarised along (100) gives

$$\dot{J}_x = \eta_{xxxx}(E^\omega)^2 E^{-2\omega} \quad (3.11)$$

$$\dot{J}_y = 0 \quad (3.12)$$

(As an aside, having the beams cross-polarised, the  $\omega$  beam being polarised along the (010) direction, gives

$$\dot{J}_x = \eta_{xyyx}(E^\omega)^2 E^{-2\omega} \quad (3.13)$$

$$\dot{J}_y = 0 \quad (3.14)$$

Since  $\eta_{xyyx} = \frac{1}{2}\eta_{xxxx}$ , cross-polarising the beams clearly results in half the injected current, but it is interesting to note that a current is still present.)

If the beams are both polarised in the (110) direction

$$E_x^\omega = \frac{1}{\sqrt{2}}E^\omega, E_y^\omega = \frac{1}{\sqrt{2}}E^\omega \quad (3.15)$$

$$E_x^{-2\omega} = \frac{1}{\sqrt{2}}E^{-2\omega}, E_y^{-2\omega} = \frac{1}{\sqrt{2}}E^{-2\omega} \quad (3.16)$$

Then, taking for GaAs

$$\eta_{xxxx} = \eta \quad (3.17)$$

$$\eta_{xxyy} = \eta_{xyxy} = \frac{\eta}{2} \quad (3.18)$$

$$\eta_{xyyx} = 0 \quad (3.19)$$

and inserting into Equation 3.9 gives

$$\dot{J}_x = \eta \frac{1}{2} (E^\omega)^2 \frac{1}{\sqrt{2}} E^{-2\omega} + \frac{\eta}{2} \frac{1}{2} (E^\omega)^2 \frac{1}{\sqrt{2}} E^{-2\omega} + \frac{\eta}{2} \frac{1}{2} (E^\omega)^2 \frac{1}{\sqrt{2}} E^{-2\omega} \quad (3.20)$$

which condenses down to

$$\dot{J}_x = \frac{\sqrt{2}}{2} \eta (E^\omega)^2 E^{-2\omega} \quad (3.21)$$

Equivalent expressions will be found for  $\dot{J}_y$ , and so the resultant injected current will be in the  $xy$  (110) direction and have magnitude given by

$$|\dot{J}| = \sqrt{\dot{J}_x^2 + \dot{J}_y^2} \quad (3.22)$$

$$= \sqrt{\frac{2}{4} \eta^2 (E^\omega)^4 (E^{-2\omega})^2 + \frac{2}{4} \eta^2 (E^\omega)^4 (E^{-2\omega})^2} \quad (3.23)$$

$$= \eta (E^\omega)^2 E^{-2\omega} \quad (3.24)$$

which is equal to Equation 3.11. Therefore both beams polarised in the 110 direction produces the same magnitude of injected current as both beams polarised in the (100) direction. This is contrary to the assertion [25] that the (100) direction is most favourable since  $\eta_{xxxx} = 2\eta_{xxyy}$ .

(It should be noted that if  $\eta_{xxxx} \gg \eta_{xyyx} \neq 0$  there is an increase in  $|J|$  for both beams in the (110) direction, but this increase is negligible.)

## 3.2 QUIC and ultrafast electrical metrology

It has been shown in the previous sections that is possible to produce a photocurrent in a GaAs semiconductor without requiring an electrical bias. This section now considers how this can be of importance to electrical metrology.

A typical electrical measurement instrument that requires calibration is the oscilloscope, which provides a measurement of voltage against time. The

characteristic that is most commonly calibrated is the impulse response or the related rise time. (In the frequency domain, the equivalent is the bandwidth.) The impulse response effectively determines the fastest electrical signals that the oscilloscope can correctly recover. Signals containing fast changes (high frequency content) beyond the response of the oscilloscope are recorded incorrectly.

To determine the oscilloscope response an electrical pulse of duration comparable to the response time, if not faster, is required. Typical commercial photodiodes pumped by ultrafast optical pulses can currently provide electrical pulses similar in duration to the faster oscilloscope responses, around 5–7 ps. Yet technology moves on all the time, and the faster the electrical pulse the greater the resolution with which the next generation of oscilloscopes can be measured. For this reason the Ultrafast Measurements Section at NPL employ the use of a *photoconductive switch* for oscilloscope calibrations. A photoconductive switch consists of a metal-on-semiconductor device, typically Au on LT-GaAs. A typical such device is a transmission line formed by two ground planes and a central conductor that contains a gap which forms the switch. If this gap is illuminated by an optical pulse of suitable wavelength then photoelectrons will be produced or ‘injected’ into the semiconductor. By applying a voltage bias between the central conductor and the ground planes these photoelectrons are accelerated and produce an electric current. When the pulse has gone no more photoelectrons are injected and those that had been soon recombine with their corresponding holes. The current therefore returns to zero creating an electrical pulse. (This is the main reason for using LT-GaAs, which has a faster recombination time than normal GaAs.) The electrical pulses produced in this way have been measured to have durations as short as 650 fs [26].

So what benefits could using QUIC provide? The key here is that, when QUIC is used to inject a photocurrent into the semiconductor, the injected photoelectrons come into existence with their momenta already established according to their location within the conduction band. In the standard photoconductive switch method described above, the initial distribution of injected photoelectrons will be random within momentum space and so the

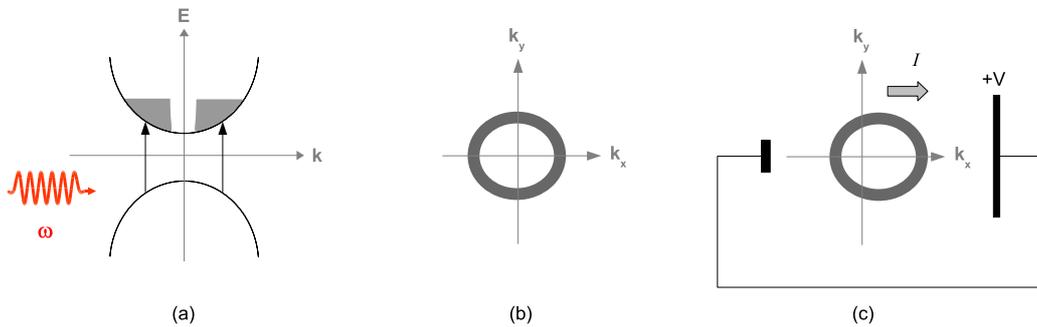


Figure 3.1: (a) One-photon absorption in a semiconductor promotes electrons from the lower valence band to the higher conduction band. (b) Since the bands are symmetric in momentum ( $\mathbf{k}$ ) space the distribution of the electrons is also symmetric and there is no net current. (c) Applying a bias shifts the distribution in  $\mathbf{k}$  space, producing a current.

net photocurrent is zero. It is not until the external bias accelerates the photoelectrons that the distribution is distorted and a net photocurrent is produced (Figure 3.1). It is clear then that there are two inefficiencies at work here: The first is the creation of photoelectrons with momenta that act *against* the photocurrent, and the second is the delay between photoelectron creation and acceleration. QUIC eliminates both these inefficiencies. By correctly controlling the phase between the two quantum processes photoelectrons will be created preferentially with momenta in the desired direction, reducing the cancelling-out effect of having photoelectrons with opposite momenta. Furthermore, because no acceleration is required due to the generation of a distribution with a net momentum (Figure 3.2) there is no delay between the injection of photoelectrons and the appearance of a photocurrent. If the generation of the photocurrent is linked in time only to the absorption of the photons then it can be expected that the rise-time of the photocurrent will be of the order of the generating optical pulse — typically 80 fs for a  $\Delta\tau = 100$  fs Gaussian pulse. The fall-time, and hence  $\Delta\tau$ , of the photocurrent will still be dominated by recombination processes within the LT-GaAs. Even so, it is clear that QUIC has the potential to provide much faster electrical pulses than conventional methods. There is

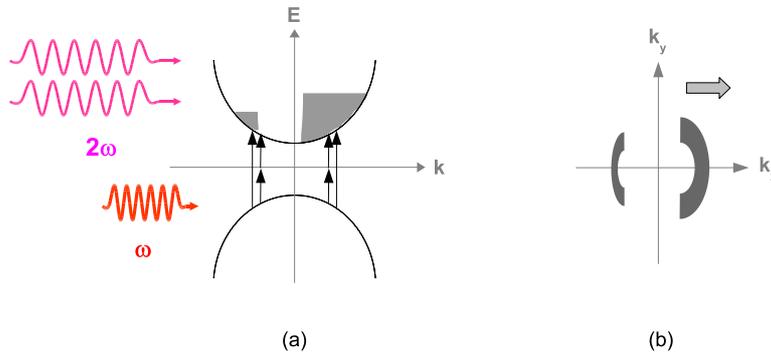


Figure 3.2: (a) One-photon and two-photon absorption occurring simultaneously in a semiconductor. Since the two processes link the same quantum states, interference between the two can affect the probability of promotion of an electron from the lower valence band to the higher conduction band. (b) Since this interference is not necessarily symmetric in momentum ( $\mathbf{k}$ ) space, the distribution of the electrons can be asymmetric, leading to an injected current.

even the possibility that the phase relation between the optical pulses for the two quantum processes could change within each pulse due to pulse shaping (Section 2.4). This would enable the injection of a photocurrent in one direction to be quickly cancelled by injection in the opposite direction, bypassing the need for the slower recombination processes and reducing  $\Delta\tau$  of the electrical pulse. The possible applications of such ultrafast electrical pulses to electrical metrology prompted NPL to invest in a Strategic Research project to design, construct and test a system capable of generating electrical pulses using QUIC.

### 3.3 Principle of the work

To integrate most easily with the existing NPL ultrafast electrical measurement system it was decided to generate QUIC within an LT-GaAs photoconductive switch similar to those used for traditional ultrafast electrical pulse generation (Section 3.2).

The two linked quantum processes chosen for QUIC in LT-GaAs were

two-photon and one-photon absorption of 1500 nm and 750 nm pulses, respectively. The reasons for choosing these processes were the applicability of an ultrafast 1500 nm source to telecommunications applications and the readily available 750 nm pulse source that would result by frequency doubling of some of the 1500 nm pulses using a nonlinear crystal (Section 1.2.2). After conversion of part of the 1500 nm into its second harmonic the two outputs would then be separated along two different paths to enable polarisation, amplitude and phase (temporal delay) changes to be introduced in either output according to need. They would then be recombined collinearly and focused onto a photoconductive switch to generate QUIC injected current. This current should take the form of an electrical pulse which would then be measured using the EOS system.

To confirm the presence of QUIC injected current prior to measurement with the EOS system it was decided to first look at charge accumulation across the photoconductive switch and check for behaviour representative of QUIC. Charge accumulation measurement is achieved simply by applying some form of voltmeter across the photoconductive switch. The assumption is that in the steady state the voltage across the switch will be proportional to the magnitude of the QUIC injected current. Once satisfied that the QUIC is occurring inside the device EOS measurement can be carried out.

### 3.3.1 Device design

Due to project constraints, the decision was made to use an old device mask for preliminary work rather than design a new one. ESE at Essex were chosen to make the devices, having previously done work of this kind for NPL. The devices were to be made in Au, on both SI-GaAs and LT-GaAs.

The mask used had a total of nine devices, alphabetically marked, and the most promising was device B (Figure 3.3). This consisted of two earth planes with a transmission line between them. At the edges the lines are of a suitable width for connection to a Wiltron test fixture (described below), each tapering towards a narrower width before meeting in the centre, separated by a 10  $\mu\text{m}$  gap.

The mask was available in both positive and negative versions. Both were sent to ESE with an explanation of what was wanted, allowing them to make their own choice of a suitable resist. The LT-GaAs wafer (M1852) had been grown in the (100) direction, a large flat indicating (011) and a smaller flat indicating  $(0\bar{1}1)$ . As it had been established that (110) was as suitable as (100) (see Section 3.1.1), it was requested that the devices be made with the transmission line parallel to the equivalent (011) direction.

Two SI-GaAs devices and two LT-GaAs devices were received from ESE, the Au film grown to a thickness of 0.5–1.0  $\mu\text{m}$ . Both devices were delivered side-by-side on their respective substrates, requiring separation using a Load-point wafer saw. After careful cutting to minimise removal of metallisation the result was four devices each measuring approximately 10.5 mm  $\times$  6.0 mm.

Due to the fragility of the semiconductor substrates, each device was mounted on quartz for added support and strength. Araldite Rapid was used to attach them and the wafer saw used to cut the quartz to the required sizes to produce four devices.

The mount used to hold the semiconductor devices was a Wiltron 3680V test fixture (NPL label X000472). This consisted of two clamps, one fixed and the other capable of movement along two metal rails with a locking nut to hold it in place. Each clamp grips the device on the earth plane and presses a connector pin to the central signal line. The signal and earth are connected within each clamp to a female V coaxial connector, allowing transition from coaxial line to transmission line. For alignment with the QUIC-generating pulse paths the fixture itself is mounted on an xyz Martock mount using a right angle block. This gives movement of approximately 2 mm in three dimensions, with fine adjustment.

### 3.3.2 Experimental Layout

A schematic of the experimental arrangement is shown in Figure 3.4. The 1500 nm output pulses (horizontally polarised) from the OPO are first raised by a periscope arrangement to a height matching the EOS system (Section 3.5.2) before being focused onto a type-I BBO crystal for second har-

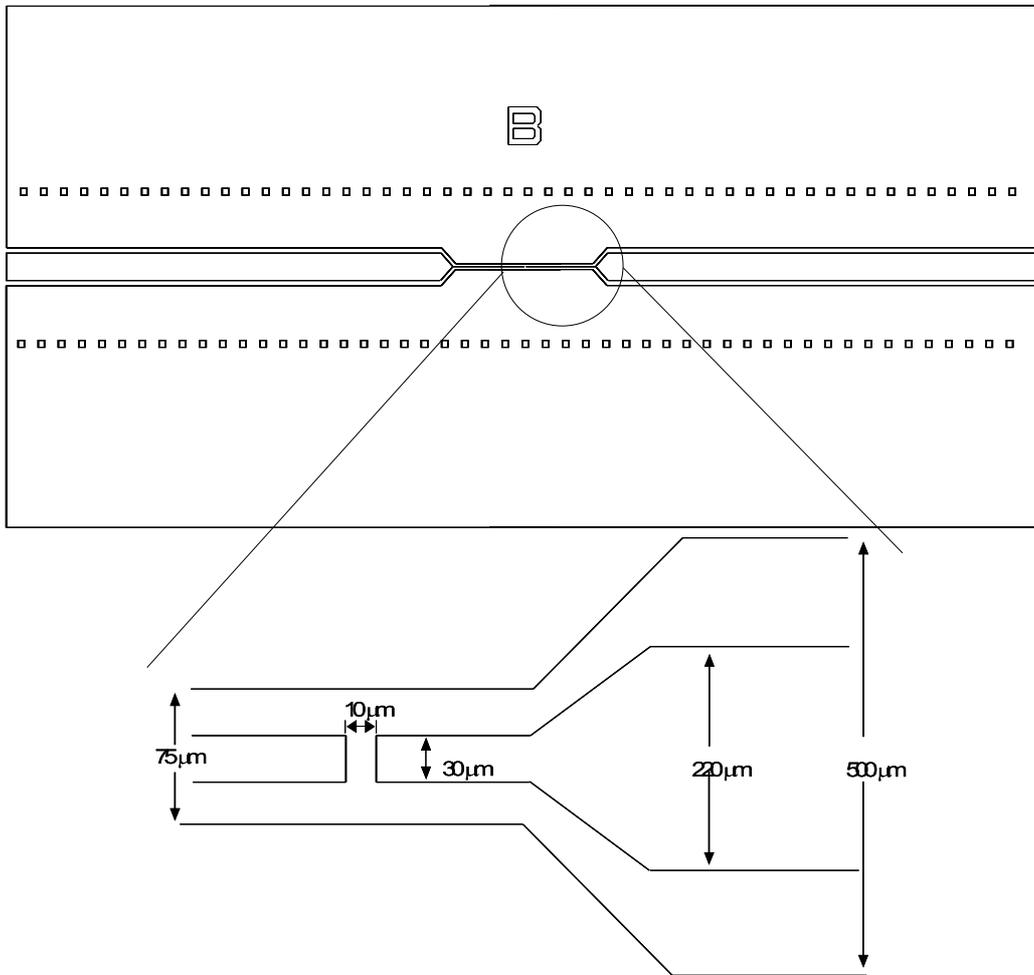


Figure 3.3: The devices made for the initial QUIC experiments were taken from the design shown here, and consisted of Au on LT-GaAs or SI-GaAs. Two ground planes sit either side of a central strip  $220\ \mu\text{m}$  wide that tapers down to  $30\ \mu\text{m}$  near the centre, where there is a  $10\ \mu\text{m}$  gap. The small squares in each ground plane serve as markers.

monic generation. The resulting 750 nm pulses and the remaining collinear 1500 nm pulses are collimated by a second lens. Next a dichroic beamsplitter separates the 1500 nm and 750 nm pulses along two separate arms.

A delay line in the 1500 nm arm, consisting of an Au-coated air retroreflector mounted on a translation stage and an additional piezoelectric transducer (PZT)-controlled stage, provides a means for adjusting the relative lengths of the two paths, changing the temporal delay and hence the phase relation of the pulses. A half-waveplate in the 750 nm arm allows the 750 nm polarisation to be adjusted. The 750 nm polarisation is orthogonal to the 1500 nm polarisation on leaving the SHG crystal, so this also allows the planar polarisations to be made equal as well as enabling investigation of polarisation-dependent QUIC behaviour. Variable neutral density filters (VNDFs) can be inserted in either arm as required (remembering to compensate for any additional delay produced by the insertion of material in the optical path). (A neutral density filter is a piece of glass or other substrate that is coated so as to attenuate all wavelengths by an equal amount. A variable NDF normally takes the form of a glass disc with different levels of coating. Rotating the disc varies the level of coating in the beam path, providing a form of adjustable attenuation.)

Both pulses travel along equal path lengths (see below) before being recombined using an identical beamsplitter. They are then focused onto the MSM device using a microscope objective lens.

A pellicle beamsplitter (Section 4.3.1) and CCD camera arrangement provide a view of the device for determining the location of the focus. Collimated light from a white-light source is reflected via a glass slide and then reflected again by the pellicle beamsplitter to illuminate the circuit. (A pellicle beamsplitter is used to minimise dispersion in the optical pump and probe pulses.) The image of the circuit is reflected back by the beamsplitter and through the glass slide and recorded by a CCD camera arrangement which is also sensitive to the 750 nm light reflected from the circuit. By this means the position of the focused pulses on the circuit can be monitored.

The MSM device itself is mounted in the Wiltron test-fixture, enabling coaxial connections to be made to the device. One such connection was made

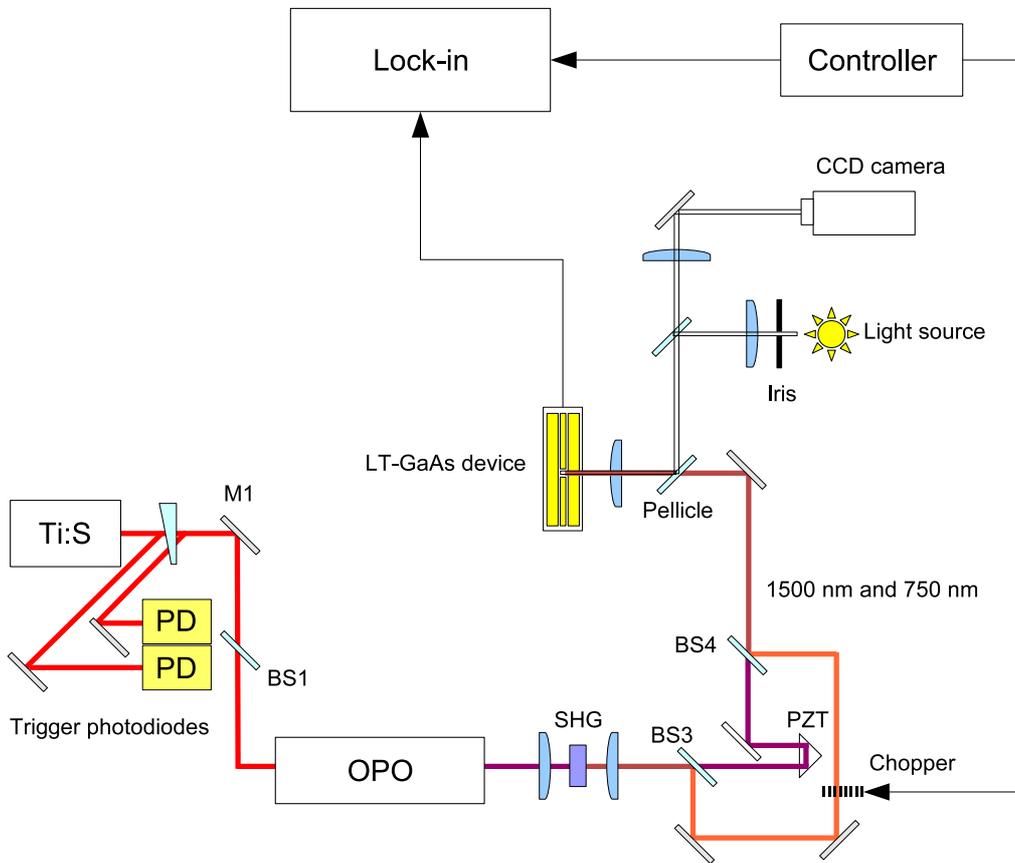


Figure 3.4: Experimental layout for the QUIC charge accumulation measurements. BS3 and BS4 are optics coated to transmit 1500 nm but reflect 750 nm.

between the device and a lock-in amplifier (SRS 810) for improved detection of any charge accumulation. The frequency modulation was provided by an optical chopper placed in the 1500 nm arm — modulation of the 1500 nm pulses prevents any QUIC signal being obscured by the single-photon absorption that would be detected when modulating the 750 nm pulses.

To ensure equal path distances of the 750 nm and 1500 nm arms a gallium-arsenide-phosphide (GaAsP) photodiode was used. The key feature of this photodiode is that the bandgap is such that two-photon absorption is induced by 750 nm + 750 nm light and also by 750 nm + 1500 nm light. In the case of pulses the latter will only occur if the pulses arrive simultaneously. By

using lock-in amplification, and modulating only the 1500 nm pulses, only the 750 nm + 1500 nm signal will be detected (a little four-photon absorption of the 1500 nm may also be detected, but this will generally be minuscule). Adjusting the 1500 nm path length to maximise the signal detected at the GaAsP photodiode therefore ensures equal path lengths between the two arms.

With both pulses focused on the MSM gap and no bias applied a voltage signal of between 3 mV and 6 mV across the gap was measured by the lock-in amplifier. With only the 1500 nm present a signal of around 0.5 mV was measured, but the origin of this was unclear. To confirm that the voltage measured with both pulses present was due to a QUIC effect a series of experiments were carried out as described in the following sections.

## 3.4 Charge Accumulation Measurements

### 3.4.1 Variation with 750 nm polarisation

It was seen in Section 3.1.1 that the QUIC effect was dependent on the orientation of the electric fields of the optical pulses relative to the semiconductor crystal. It follows that there will be a variation in the QUIC injected current with polarisation of one or both of these beams. By calculating this variation and checking experimentally for it, evidence of QUIC can be obtained.

Theoretically, the injected QUIC current density in a direction  $\mu$  is given by

$$\dot{J}_\mu = \eta_{\mu\alpha\beta\gamma} E_\alpha(\omega) E_\beta(\omega) E_\gamma(2\omega) \quad (3.25)$$

where  $\mu$  is the  $x$ ,  $y$  or  $z$  direction in relevant geometry and  $\alpha$ ,  $\beta$  and  $\gamma$  are summed over  $x$ ,  $y$  and  $z$ . For the LT-GaAs device, taking  $x$  to be the (100) direction the only non-zero components of  $\eta$  are  $\eta_{xxxx}$ ,  $\eta_{xyxy}$  and  $\eta_{xyyx}$  where  $x$ ,  $y$  and  $z$  are cyclic. As seen in Section 3.1.1, these values are related by

$$\eta_{xyxy} = \eta_{xyyx} = \frac{1}{2}\eta_{xxxx} \quad (3.26)$$

$$\eta_{xyyx} \approx 0 \quad (3.27)$$

For our MSM the direction across the gap is the (110) direction, and it can be shown that the current density in this direction is given by

$$j = \sqrt{j_x^2 + j_y^2} \quad (3.28)$$

where

$$\dot{J}_x = \eta E_x(\omega) (E_x(\omega)E_x(2\omega) + (E_y(\omega)E_y(2\omega))) \quad (3.29)$$

$$\dot{J}_y = \eta E_y(\omega) (E_y(\omega)E_y(2\omega) + (E_x(\omega)E_x(2\omega))) \quad (3.30)$$

where  $\eta = \eta_{xxxx}$ .

These equations can be combined to produce a function  $\dot{J}(\theta)$ , where  $\theta$  is the polarisation of the 750 nm beam relative to the (110) direction. Taking the assumption that the voltage  $V$  across the MSM gap is  $V \propto \dot{J}$  then a plot of voltage against 750 nm polarisation can be used to check for the expected dependency.

Figure 3.5 shows the results from two measurements of voltage for various settings of a polarisation wheel containing a half-waveplate suitable for 750 nm placed in the 750 nm arm of the QUIC experimental set-up as shown in Figure 3.4. These measurements were taken with both 1500 nm and 750 nm pulses present on the MSM gap and the solid curves show the dependence that would be expected if the signal variation were due to a QUIC effect as described above. In each case this expected dependence takes the functional form  $\dot{J}(\theta)$ , scaled and offset to align with the measured results. It can be seen that the measured results clearly follow the expected dependencies.

From examination of the equations of the theoretical curves given in Figure 3.5 it can be seen that both are offset from the expected value by  $5^\circ$ . This is presumed to be due to a slight error of the half-waveplate alignment within the polarisation wheel used to perform the rotation. That the offset is the same in each case gives weight to this possibility.

The differences in magnitude between the two signals is assumed to be due to a difference in the OPO output power when the readings were taken on two different occasions. It can be noted that the minimum values in each

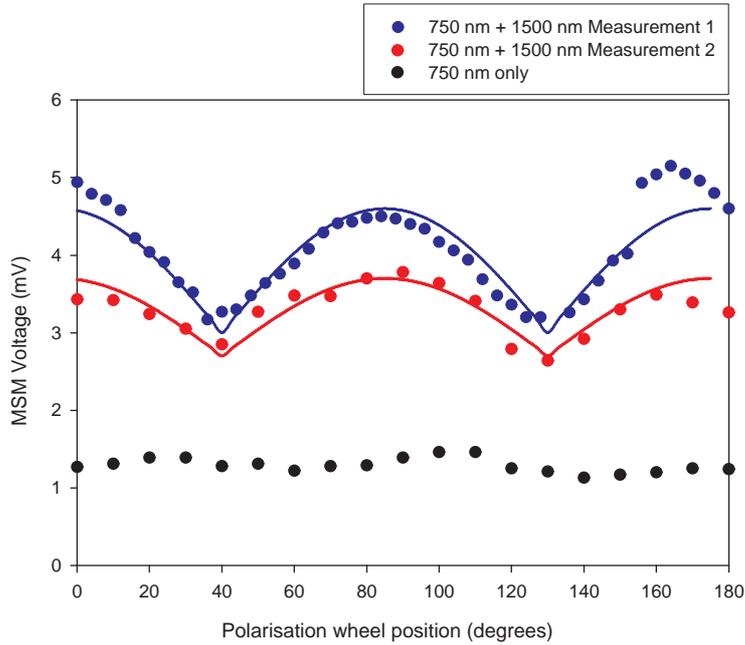


Figure 3.5: Measurements of the voltage across the MSM gap for varying polarisations of the 750 nm pulses. The solid lines show the dependence that would be expected if the voltage was due to QUIC. The blue curve is given by  $V = 1.6J(\theta - 5^\circ) + 3.0$  mV and the red curve is given by  $V = J(\theta - 5^\circ) + 2.7$  mV

case (i.e. the vertical offsets), which should represent no QUIC signal present, are almost equal.

If this polarisation dependence is due to QUIC then such a dependence would not be expected when no 1500 nm pulses are present. To confirm this a similar set of measurements was taken with the 1500 nm pulses blocked and optically chopping the 750 nm pulses (Figure 3.6). A number of features are apparent here that demonstrate the same dependence is not present. The closest-fit theoretical curve based on  $\dot{J}(\theta)$  has a different phase-shift to the  $-5^\circ$  of the earlier curves, More importantly, it shows that any apparent maxima and minima within the data lack the same periodicity of that seen in the data taken with the 1500 nm present. What variations are seen are most likely due to variations within the amplifier reading.

In summary, when 1500 nm and 750 nm pulses are present simultaneously

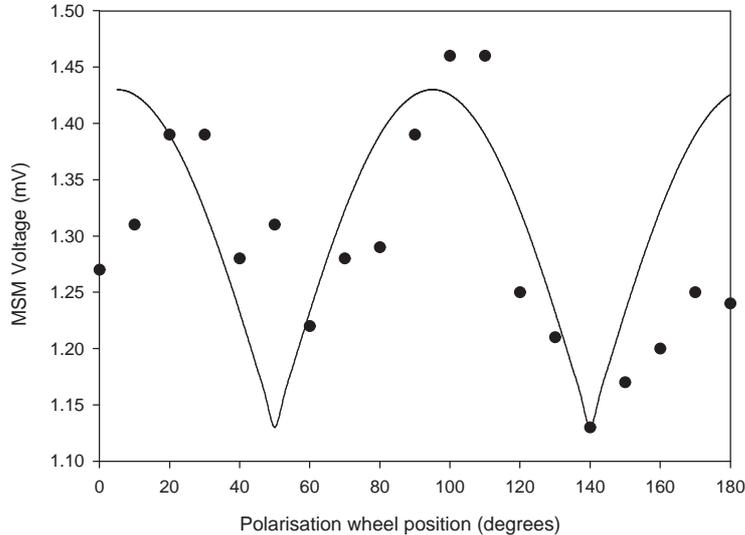


Figure 3.6: Measurements of the voltage across the MSM gap for varying polarisations of the 750 nm pulses when no 1500 nm pulses were present. The solid line shows the closest fit found for a curve given by a  $V = xJ(\theta - \phi) + y$  function.

the voltage across the MSM junction is seen to vary with the polarisation of the 750 nm pulses in a way that corresponds to that expected from theory if the voltage was due to a QUIC effect. When only the 750 nm pulses are present the variation is different and arguably random. Therefore it can be inferred that the signal seen with both pulses present is due to QUIC current injection. To lend further support to this several other aspects were also investigated.

### 3.4.2 Variation with 750 nm power

The macroscopic speeds of the QUIC-injected photocarriers (corresponding to the drift speed of a conventional current) is given by  $v = \dot{J}/e\dot{N}$ . This can be shown[25] to have the following dependence upon the intensities of the

750 nm and 1500 nm pulses:

$$v \propto \frac{I(\omega)\sqrt{I(2\omega)}}{\alpha I(2\omega) - \beta(I(\omega))^2} \quad (3.31)$$

where  $\alpha$  and  $\beta$  are, respectively, the one- and two-photon absorption coefficients of the semiconductor in question. From this it follows that the maximum speed (and hence maximum current) occurs when  $\alpha I(2\omega) = \beta(I(\omega))^2$ . Therefore a second test of correct QUIC behaviour is to check the power dependence of the measured signal.

The power from the OPO was measured to be approximately 80 mW. Assuming the spot sizes of the 750 nm and 1500 nm pulses at the MSM junction to be equal and approximately 20  $\mu\text{m}$  as judging from the image of the circuit seen with the CCD camera, the pulse  $\Delta\tau$  to be roughly 100 fs and the repetition rate to be 80 MHz, a calculation gives  $I(2\omega) \approx 333 \text{ MW/cm}^2$ . Using the values of  $\alpha$  and  $\beta$  for LT-GaAs to determine the  $I(2\omega)$  that gives maximum current according to the above relation and then calculating back, again using the above assumptions for pulse  $\Delta\tau$  and spot size, produces an optimal power for the 750 nm pulses of 8  $\mu\text{W}$ .

The power of the 750 nm pulses was attenuated using a VNDF and both the 750 nm power and the resulting voltage across the MSM gap was measured for different levels of attenuation. A plot of these measurements is shown in Figure 3.7 along with the theoretical carrier velocity as a function of  $P(2\omega)$  based on Equation 3.31, where it has been assumed that  $V_{MSM} \propto v$  and  $P(2\omega) \propto I(2\omega)$ .

It can be seen that although the measured and the theoretical results do not overlap, there are clear correlations between the two. The signal rises steeply with increasing power but not quite linearly, the gradient decreasing slowly as the power increases. It then reaches a peak before dropping away with increasing power. Unfortunately (for this experiment at least) the maximum 750 nm power available is not much more than that required for the peak, making it difficult to confirm that the signal dependence follows the expected form beyond the peak. However, the non-linear relation between

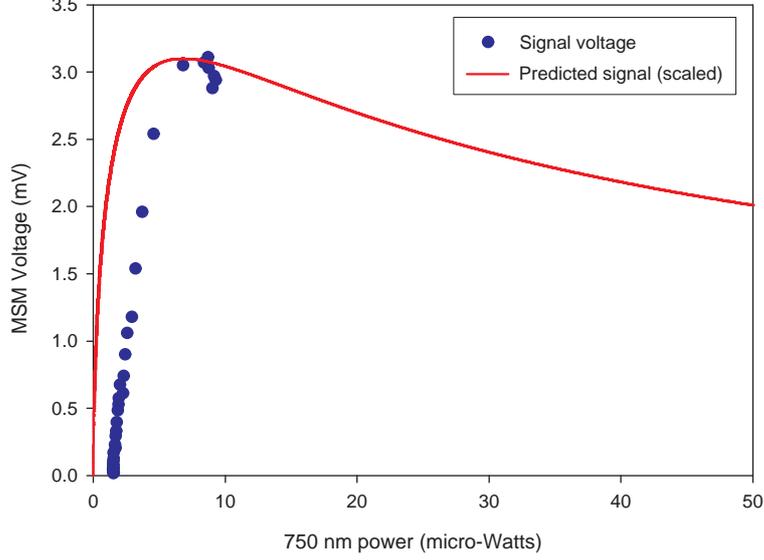


Figure 3.7: Measurements of the voltage across the MSM gap for varying powers of the 750 nm pulses.

the power and the signal, with a gradient similar to that predicted by theory, the location of the peak that also corresponds to theory and the reduction of the signal at higher powers all lend weight to the conclusion that the signal seen is the result of QUIC injected current. Nevertheless, further experiments were carried out on other aspects to provide more evidence.

### 3.4.3 Variation with pump power

From Equation 3.25 it can be seen that, in the case where the two beams have matching polarisation,

$$|\dot{\mathbf{J}}| \propto (E_\omega)^2 E_{2\omega} \quad (3.32)$$

Since the OPO output is used to generate the second harmonic,  $E_{2\omega} \propto (E_\omega)^2$  and as a result

$$|\dot{\mathbf{J}}| \propto (E_\omega)^4 \quad (3.33)$$

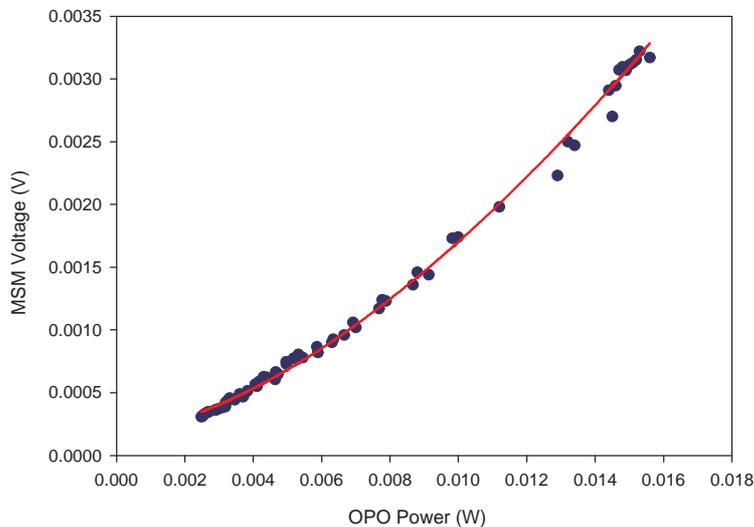


Figure 3.8: Measurements of the voltage across the MSM gap for varying powers of the 1500 nm pulses prior to the SHG crystal. The solid line has the equation  $V = 2.5P^{1.61} + 0.0002$

$$\propto (P_{\omega})^2 \quad (3.34)$$

$$\propto (P_{\text{OPO}})^2 \quad (3.35)$$

Taking the assumption that  $V_{MSM} \propto |\mathbf{j}|$  then  $V_{MSM} \propto P_{\text{OPO}}^2$  and so it would be expected that the voltage across the MSM gap should vary quadratically with the power of the OPO. Therefore the experiment described in this section was carried out to check for such a quadratic dependence.

To measure the variation of the suspected QUIC signal with pump power a VNDF was placed between the OPO output and the SHG crystal and used to attenuate the 1500 nm power. For a number of positions of the VNDF the 1500 nm power was first measured and then the voltage across the MSM measured after removal of the power meter. These measurements are plotted in Figure 3.8. It can be seen that the best fit shown by the solid line represents a dependence on the OPO power of  $P_{\text{OPO}}^{1.61}$ , which is not the expected quadratic variation.

At first sight this failure to match the theoretical result looks disappoint-

ing. However, researching the literature provides two examples where a similar measurement was carried out[25][27] that determined dependences of  $V \propto P_{\text{OPO}}^{1.6}$  and  $V \propto P_{\text{OPO}}^{1.7}$  for signals that were confirmed (by other methods) to be QUIC generated. This agreement with other verified instances is more encouraging, even if the reasons behind the non-quadratic dependence are not readily understood. The important result here is that the signal dependence on OPO power was not found to be cubic or quartic, dependences which would have suggested that it was due to an optical rectification effect. Even so, one last series of experiments were planned to verify the presence of a QUIC-generated signal, and they were based on the most importance dependence of all, that on the delay between the signals.

#### 3.4.4 Variation with relative delay

From Equation 3.7 it can be shown that the phase dependence of the QUIC signal is given by

$$|\dot{\mathbf{J}}| \propto \sin(2\phi_\omega - \phi_{2\omega}) \quad (3.36)$$

where  $\phi_\omega$  and  $\phi_{2\omega}$  correspond to the phases of the 1500 nm and 750 nm pulses, respectively. From this it follows that a QUIC signal should vary as a sine function with a period corresponding to successive delays of 750 nm in either pulse. This dependence is one of the most characteristic features of QUIC, and confirming it provides substantial confirmation of its presence.

To look for this dependence the length of the 1500 nm path was varied using an actuator upon which was mounted an Au hollow-cube retroreflector. To improve signal-to-noise the fluctuations in the OPO output were accounted for by using a glass wedge to pick off a small amount of the OPO output and measuring the power simultaneously with the voltage across the MSM gap. Using this system several sequences of voltage measurement across the MSM gap were carried out for delays ranging from -4000 nm to +2000 nm on several different occasions. These are plotted in Figure 3.9.

The black sine curve in Figure 3.9 is a modelled curve scaled to match the data and has a period of 750 nm. It is important to note that since 750 nm is the period expected if the variation is due to QUIC this was the period

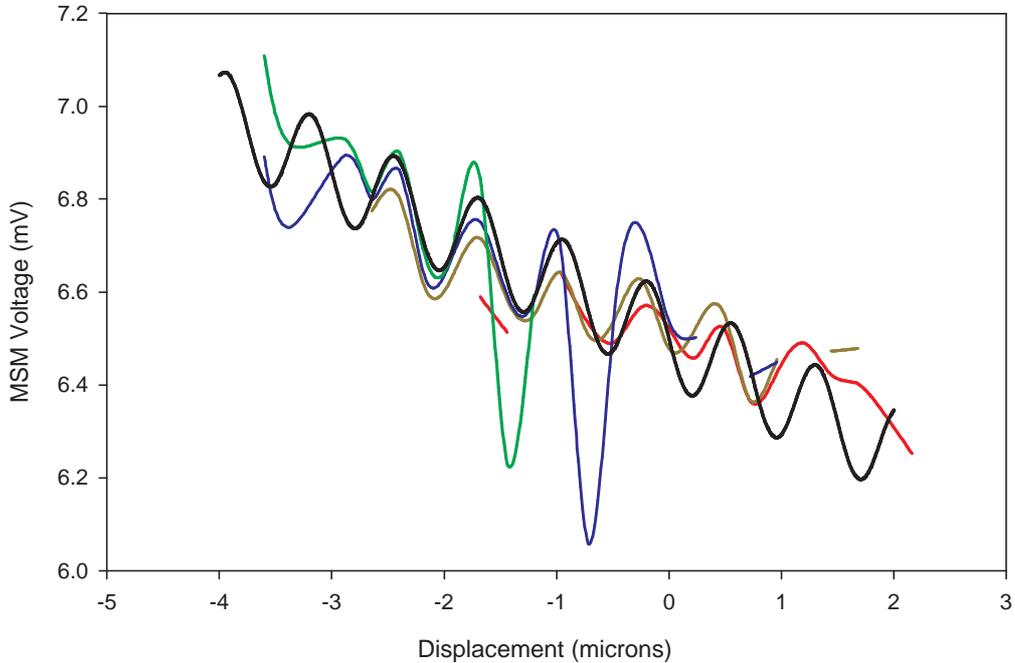


Figure 3.9: Measurements of the voltage across the MSM gap for varying delays of the 750 nm pulses relative to the 1500 nm pulses. The black curve is a sine function with period 750 nm.

to which the sine curve was set — the period is *not* a best-fit to the data. The amplitude and phase, however, were adjusted to provide a best-fit to the data, which is reasonable. It was also found to be necessary to apply an increasing offset with voltage.

It can be seen that there is a significant amount of agreement between the data and the sine curve in the range -2500 nm to +500 nm, but beyond this range discrepancies start to occur. There are a number of possible explanations for this. It could be that errors are occurring in the displacement as the actuator moves away from the central point, through hysteresis or some nonlinear translation. Another possibility is the reduced temporal overlap of the pulses. A 3000 nm total delay corresponds to 10 fs, approximately ten percent of the estimated  $\Delta\tau$  of the pulses. As the temporal overlap of the

pulses reduces it is plausible that the reduced interacting intensities result in a reduction of the QUIC effect. In this case it would be expected that the region of best agreement would be centred around zero delay, but the displacement seen could simply be an error in establishing the zero-delay using the GaAsP photodiode (c.f. Section 3.3.2). The given offset-centre would be explained by an error in establishing the zero delay of just 1  $\mu\text{m}$ . Note that on several of the curves the largest variations with delay also centre around the -1000 nm delay, suggesting the QUIC effect is strongest here and lending support to the latter theory.

The downward slant of the signal with delay is an unexpected effect and the origin of it is not necessarily clear. The most likely explanation would be some slight lateral displacement of the 1500 nm pulses as the retroreflector is translated, reducing the spatial overlap of the 1500 nm pulses with the 750 nm pulses and so having an adverse effect on the QUIC signal. One problem with this explanation though is that it would not be expected to affect the lowest value of the signal (when the QUIC signal should be absent, due to  $\sin(2\phi_\omega - \phi_{2\omega}) = 0$ ), unless the base signal is the result of some form of two-photon signal being detected.

It is clear then that there are a number of features in Figure 3.9 which are unexpected and while reasonable explanations exist for most of them further investigation would be required to establish the causes of these features. From the point of view of establishing the presence of QUIC however it can be said that these measurements have been a success, in that a predicted variation with phase was seen.

### 3.4.5 Discussion of charge accumulation results

A system was built, based around a femtosecond OPO, that produced femtosecond pulses at 1500 nm and 750 nm and enabled properties of these pulses, such as their relative delay, to be controlled. These pulses were focused onto the gap of an Au on LT-GaAs device where, it was predicted, a photocurrent should be injected in the absence of a voltage bias due to the QUIC effect of the one-photon and two-photon processes. On connecting a

lock-in amplifier to measure the voltage due to charge accumulation across the gap (modulating the 1500 nm pulses) a signal was indeed seen and so a series of experiments were carried out to confirm that this signal was in fact the result of a QUIC effect and not due to some other unknown mechanism. In each experiment a prediction was made from QUIC theory as to what would be expected before the actual measurements were carried out.

The effect of a variation in the polarisation of the 750 nm pulses was calculated and found to be a function that was not a simple circular (sine or cosine) function. Variations in a signal due to changes in reflectance, caused by changes in polarisation, from an optical surface as a mirror for example would be expected to follow such a simple circular function. Experiments carried out on separate occasions demonstrated a variation that followed the predicted QUIC-result, and only when the 1500 nm and 750 nm pulses were present, indicating that this was an effect caused by an interaction between the two.

Next a calculation of the expected variation with 750 nm power predicted a sharply rising gradient that slowly levelled off to a peak beyond which a slow decrease would be seen. Such a gradient was seen in the measured results, peaking at a power value close to that predicted theoretically. This peak was already close to the maximum available 750 nm power, but those few measurements that could be made at higher 750 nm powers found an expected decrease, unlike either the further increase that might be typically expected, or a levelling of the signal if saturation was occurring.

A calculation of the expected dependence on pump power predicted a square dependence, yet the power dependence found was closer to 1.6 than to 2. This however correlated with similar experiments made by another group who had confirmed the presence of a QUIC effect, and moreover eliminates the possibility of optical rectification.

Finally, a number of measurements were carried out to test for a variation in signal with relative delay (phase) between the two pulses, one of the most distinctive features of QUIC effects. These were all found to exhibit the expected 750 nm period over displacements of several times this amount in a remarkably consistent way. Other unexpected features were also seen,

such as an offset decreasing with delay, the seeming displacement of the zero delay and a reduction of the effect away from the zero delay point. However explanations could be proposed for each of these that were consistent with the existence of a QUIC signal.

In summary, a series of experiments were carried out in which the dependence of the signal on very different aspects of the system was measured. In each case the dependence found was that which corresponded to the signal being due to a QUIC process, suggesting strongly that QUIC was injecting photocurrent into the device. Having determined this, the next stage was to attempt to resolve this injected QUIC photocurrent in the time domain.

### 3.5 Time-resolved Measurements

Having established through charge accumulation measurements that QUIC current was being injected into the LT-GaAs device the interest now is to resolve the injected current in time and see if there really is, as surmised, the creation of electrical pulses far faster than those seen using photoconductive switch techniques.

The simplest method of measuring the QUIC electrical signal in time would be to use the coaxial connection to the device to connect it to an oscilloscope, yet doing this failed to reveal any sign of a signal. The reason for this is most likely due to the limited bandwidth of the measurement system. The QUIC injected current is expected to have a rise-time of the order of the optical pulses responsible, leading to  $< 1$  ps electrical pulses. Yet the rise-time of a typical high-bandwidth (50 GHz) oscilloscope used at NPL, such as the Tektronix SD32, is approximately 7 ps, and there would be additional broadening and attenuation of the electrical pulse due to the effects of cables and the transition from the coplanar waveguide to the coaxial connection. If the injected pulse is of low amplitude to begin with then the broadening, which disperses the energy of the pulse in time, can reduce it until it lies within the noise level of the oscilloscope preventing the signal being detected. Besides, the primary interest is the structure of such a fast pulse. What is needed is a system that has sufficient bandwidth to measure

these ultrafast pulses, and electro-optic sampling is such a system.

### 3.5.1 Electro-optic sampling

Electro-optic sampling (EOS) is a widely-used technique that enables electrical fields to be probed via optical means, often in a non-invasive way. The optical nature of the probe allows measurement systems to be built with much higher bandwidths than conventional electronic systems, and EOS also enables the mapping of electric fields in two or even three dimensions.

The basis of EOS is the Pockels effect (Section 1.2.2), where a dc electric field causes a linearly proportional change in the refractive index of a suitable material (commonly referred to as an electro-optic material). The resulting birefringence causes polarisation changes in light passing through the material, so by measuring the change in polarisation the magnitude of the electric field can be calculated. If pulsed light is used, only a small part in time of the electric field is measured, making what is essentially a sampling gate. Synchronising the arrival of the pulses to a periodic electrical signal enables a section of that signal to be repeatedly sampled, improving the signal-to-noise ratio. Introducing a delay between the timings of the optical pulses and the periodic signal results in a different section of the signal being sampled, and so varying the delay results in the build up of a waveform, effectively turning the system into an sampling oscilloscope. Because it is only the polarisation change that relates to the electric field some form of power measurement after a polarising beam-splitter or similar is sufficient, and so there is no need for a fast optical detector. As a result, the measurement bandwidth of the system is dependent only on the electro-optic material used (including its geometry) and the optical probe pulse.

Some systems use the substrate of the device itself as the electro-optic material (a technique known as internal EOS), but most commonly a separate piece of material forms the electro-optic probe and is brought into close proximity to the field to be measured (external EOS). The method of synchronising the EOS probe pulse to the electrical signal to be measured is also a matter of consideration. If the signal source is a photodiode this can often

be done optically by splitting the output from the laser source and using part of the output to pump the photodiode and the remainder to perform EOS (this technique is usually referred to as pump-probe EOS). If the signal source is purely electrical then some form of electrical oscillator can be used to drive it and the laser synchronised to the oscillator using a phase-locked loop or similar.

### 3.5.2 The NPL EOS system

To carry out the time-resolved measurements of the QUIC injected current the NPL Ultrafast Group's EOS system was employed after performing a number of adaptations as described in Section 3.5.3. The original EOS system is shown in Figure 3.10. All mirrors are dielectric-coated for high-reflection at 700–900 nm unless otherwise stated.

The NPL EOS system uses an external lithium-tantalate ( $\text{LiTaO}_3$ ) EOS probe ( $250 \mu\text{m} \times 250 \mu\text{m} \times 25 \mu\text{m}$ ) to make high-bandwidth electrical measurements in coplanar waveguide as part of the Ultrafast calibration service. Its primary use is to calibrate high-speed photodiodes (commercial or photoconductive switch-based) which are then used in turn to calibrate fast sampling oscilloscopes. The source of ultrafast pulses is a Coherent Mira providing  $< 150$  fs pulses centred at 800 nm or 850 nm (depending on the application) and a repetition rate of 80 MHz, with an average power of  $> 1$  W and horizontal polarisation. An optical glass wedge picks off two beams from the main beam, each about 4% of the power, and these are focused onto two separate photodiodes to provide electrical triggering synchronised to the laser when using an oscilloscope to record the electrical pulses generated from the pump pulses. The rest of the output can either be passed through a Coherent pulse-picker (not discussed in this thesis) or guided around via mirrors M1–M4. The optical path followed afterwards is the same in either case.

The output is raised by mirrors P1–P2 (the polarisation remaining unchanged) and another part of the output picked out by a second optical glass wedge to provide pump pulses as needed. These are used to generate the electrical pulse signal to be measured by either pumping a commercial

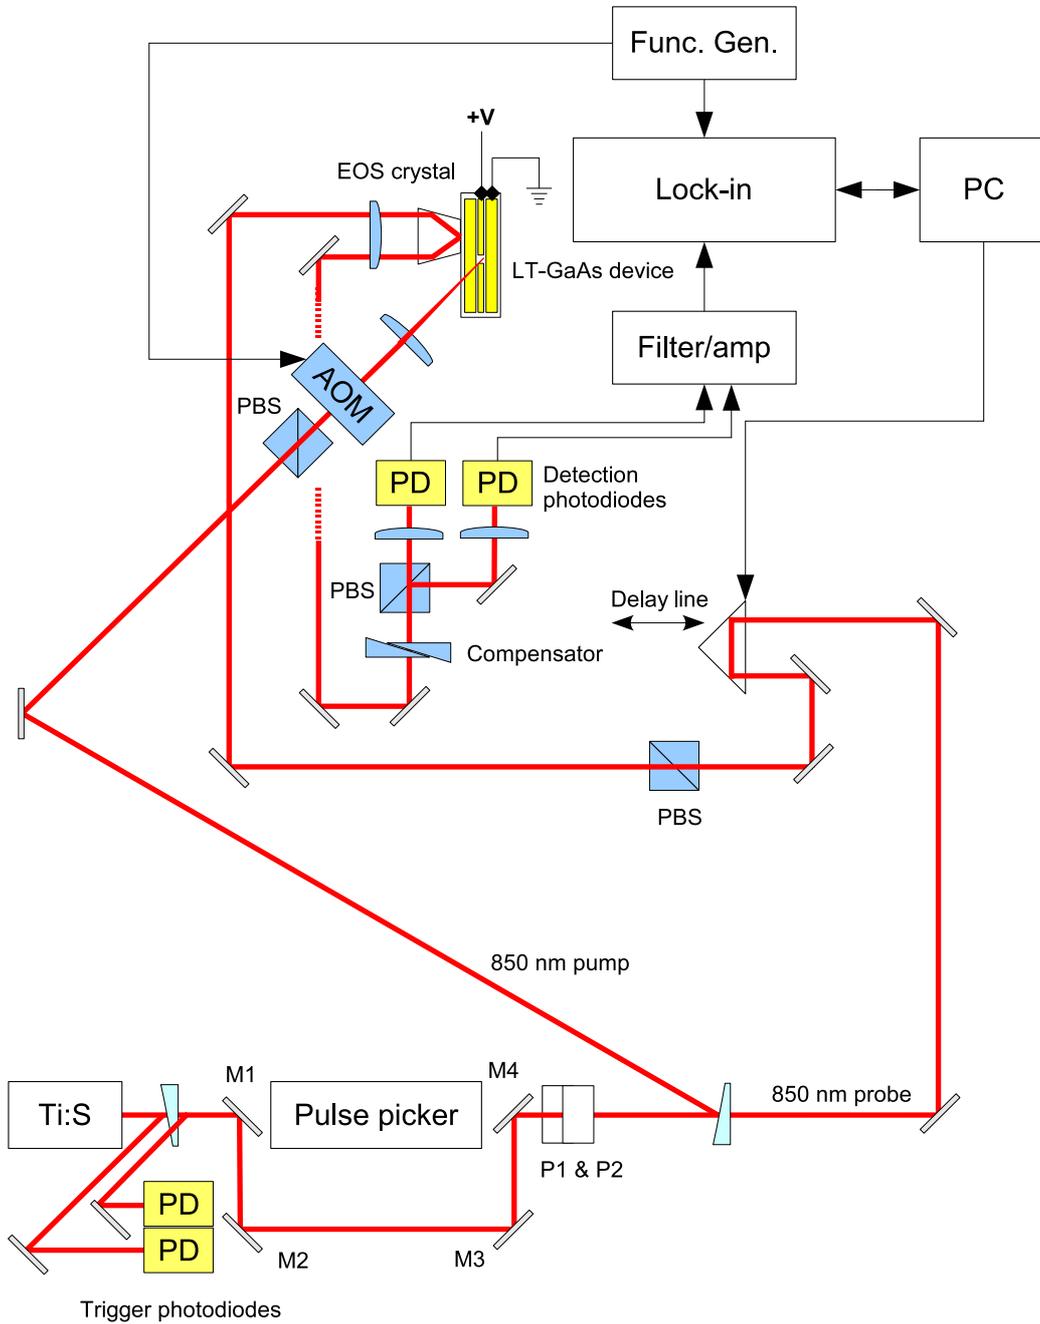


Figure 3.10: The NPL EOS system. Full details are given in the text. PBS denotes a polarising beamsplitter cube.

photodiode which is connected to a transmission line device or, as shown in the figure, pumping a photoconductive switch in the transmission line. The remainder of the laser output after the second wedge is guided to the EOS probe via a free-space corner-cube Au retroreflector, which is mounted on an delay stage of length 400 mm to provide the variable arrival times required for sampling.

The  $\text{LiTaO}_3$  crystal is mounted on a piece of quartz which is in turn attached to a quartz cylinder using optic cement, and this cylinder mounted on a thin flexible transparent piece of plastic in the same way. A microscope objective lens (removed from its casing due to space restrictions) focuses the input pulses onto the electro-optic crystal. A number of miniature translation stages enable the EOS crystal to be aligned correctly with this lens. The pulses undergo total internal reflection inside the electro-optic crystal and the returned beam is passed through a compensator (having been effectively collimated on the return journey through the objective lens) before being split by a polarising cube beamsplitter. Each polarisation is focused onto a multimode optic fibre and transported across the laboratory to a Si photodiode. (The reason for using optic fibre is to keep the detection system away from electro-optic crystal so as to prevent RF interference.) Using a photodiode for each polarisation permits differential detection, increasing the sensitivity of the system (since a polarisation change effectively causes twice the change in signal). To achieve this a VNDF is placed before each optic fibre so that the optical power in each fibre can be balanced when no electric field is present at the EOS point.

The output from each photodiode is passed through a bias voltage circuit and from there to a lock-in amplifier (SRS SR844) via a differential amplifier/filter (SRS SR560). The lock-in frequency for the amplifier is provided by a signal generator (SRS DS345) and the electrical pulse generated by the pump pulses modulated by passing the pump pulses through an acousto-optic modulator (AOM) driven by the same signal generator. A polariser is often placed before the AOM as different pulse polarisations travel at different speeds through the AOM crystal, leading to spurious results. (For a similar reason a polariser is also included in the probe path before the EOS

crystal.)

For typical measurements the circuit under test is mounted vertically (the input path of the EOS probe beam being horizontal). Translation stages allow the circuit to be moved in the xyz directions relative to the EOS probe. Two further xyz stages are used to bring Picoprobe coaxial-to-coplanar waveguide connectors into contact with the circuit to provide a transition between the two methods of electrical transmission. Three camera systems provide views as an aid to aligning the circuit under test, the Picoprobes and the EOS probe. One camera provides a 45-degree view and a second a view almost in the plane of the circuit. A third camera provides a view down the EOS crystal by means of a pellicle beamsplitter and a collimated light-source in the same way as that used for the charge accumulation measurements.

### 3.5.3 Alterations to the NPL EOS system

Although it had proved more than capable for the purposes of photodiode measurement[28], several alterations needed to be made to the existing EOS layout to make it suitable for the time resolution of the QUIC injected current. These can be seen in Figure 3.11 and are discussed below.

#### **Laser source**

The first issue is that of the laser source, the Mira. For EOS of QUIC signals the Mira is required to both pump the OPO and provide the probe pulses for EOS. This is readily achieved using a beam-splitter, but the needs of both systems must be balanced — too little power to the OPO will reduce the QUIC signal, yet if not enough is diverted to the electro-optic probe the sensitivity of the EOS system will be even more severely reduced. It was decided that the best compromise was an 80:20 (in favour of the OPO) beam-splitter coated for 850 nm, and this was installed in place of mirror M2 in Figure 3.10. Splitting the source in this way ensures that the pump and probe pulses are synchronised to each other and will have virtually no jitter between them (since the same pulse that does the probing is the one that



ultimately created the pump pulse, via pumping the OPO and generating the second harmonic for QUIC).

(In fact, the beam-splitter and mirror M2 were each mounted on separate relocatable magnetic base-plates to give the option of diverting the full power to the EOS system or not, according to the need of experimental work unrelated to this thesis. This mounting was done extremely carefully so as to keep the optical path to the EOS system the same no matter which reflector was used, which is no mean feat. For the purposes of this discussion, though, it can be assumed that the beam-splitter was in position permanently.)

### **Path lengths**

The second issue is that of matching path lengths so that the pump (i.e. the QUIC) pulses and the probe pulses arrive at the same point simultaneously. In this case the path lengths that require matching begin at the beam-splitter at M2 and end at the electro-optic probe. As the OPO path also necessarily contains the cavity length of the OPO as well as the QUIC apparatus the EOS path length needed to be extended to around 4 m in length, and to have the signals arriving within even a few picoseconds of each other requires the path lengths to be equal to 1 mm. Fine-tuning is possible using the delay line once the bulk of the path lengths have been equalised through the use of mirrors, but the short length (400 mm, equating to under 3 ns allowing for double-pass) can lead to frustration if the bulk lengths are different enough such that the delay line cannot match them. Therefore it was decided to incorporate a second delay line, previously procured by NPL but as yet unused, which was a full 2 m in length. This was capable of introducing over 13 ns of delay in double-pass — more than the time between successive pulses from the Mira. Ultimately this would replace the smaller delay line but because of issues in adapting the control software it was initially employed in conjunction with the smaller delay line, providing a large, easily adjustable delay to roughly match the path lengths while the smaller delay line was used to perform the EOS measurements.

### Pump-probe alignment

The third issue is that of delivering the pump (QUIC) pulses to the device together with the probe pulses. This is actually easier to achieve than in the conventional NPL set-up where part of the Mira output is used to pump a photoconductive switch on an LT-GaAs device. In the latter case, the pump pulses must be delivered from outside of the electro-optic crystal or they might produce spurious signals that interfere with the EOS. This requires tight and close alignment of optics, and the physical mounting of the EOS crystal can make it difficult to place the probe point as close to the source point as might be liked. For the QUIC pulses, however, matters are much simpler. The 1500 nm pulses are not detected by the Si photodiodes used in the EOS system and the 750 nm pulses are too weak to have any significant effect. Therefore the pulses can be delivered down through the EOS crystal, enabling sampling to take place very close to the generation point, simplifying the optics required and allowing the use of the probe-view camera to align the QUIC pulses with the LT-GaAs gap. The one difficulty is bringing the 1500 nm and 750 nm pulses parallel to the 850 nm pulses; however, it was found that a mirror coated for 1500 nm was sufficiently reflective at 750 nm (due to the harmonic) yet transparent to 850 nm, serving as a form of dichroic mirror (see Figure 3.11).

### 3.5.4 Experimental set-up

The experimental layout is as shown in Figure 3.11. The 1500 nm and 750 nm pulses are directed along a path that is parallel to, but non-collinear with, the path taken by the 850 nm pulses. In this way the 1500 nm and 750 nm pulses are passed down the centre of the electro-optic crystal to generate QUIC current within the device whereas the 850 nm pulses strike the side of the crystal mount, double-pass through the crystal itself and are reflected back along a separate path for electro-optic detection.

The optical chopping used during the charge accumulation measurements can only provide frequencies of perhaps a few hundred kHz, which has been found to raise too many signal-to-noise issues in the EOS system. Therefore

the reference signal for the lock-in amplifier needs to be supplied by an AOM placed in the arm of one of the QUIC beams, running at 1 MHz. Since an AOM works by diffracting the beam the power of the modulated beam will be, at best, 50% that of the unmodulated beam. Therefore modulating the 1500 nm will result in a lower QUIC signal. Modulating the 750 nm could feasibly improve the QUIC signal since there is a possible excess of the power in the 750 nm pulses. Unfortunately this would also introduce the considerable difficulty of trying to align the two pulse paths and find the zero delay with only a few microwatts of 750 nm. As a result it was decided to use a 1500 nm AOM.

### 3.5.5 EOS Measurement with 1500 nm modulator

The AOM was placed in the 1500 nm arm and the modulated signal maximised by appropriate adjustment of the position of the AOM crystal relative to the incident pulses. The diffracted pulses were aligned to be collinear with the 750 nm pulses after the second dichroic beamsplitter and the new zero delay found using a GaAsP photodiode. The two beams were then guided onto the gap of the LT-GaAs MSM device under test.

To match in time the arrival of the pump (1500 nm and 750 nm) and probe (850 nm) pulses the 2 m delay line was first adjusted to provide a coarse agreement as measured with a metal ruler. For finer adjustment the time delay was measured by using a Picoprobe to connect an oscilloscope and dc bias to the MSM device. (The other side of the device was connected to a 50  $\Omega$  termination using a second Picoprobe.) The application of the bias produces pulses of photocurrent from the 850 nm pulses and also the 750 nm pulses. By inspecting the voltage trace for the peaks indicating 750 nm and 850 nm absorption and the resulting photoconductive switch effect the delay between the pulses can be determined from the positions of the peaks along the timebase. The 2 m delay line was then adjusted until the peaks were < 100 ps apart, bringing them well within range of the smaller delay line.

Once aligned the delay between the pulses was often checked in the manner described above, and it was found during these checks that the

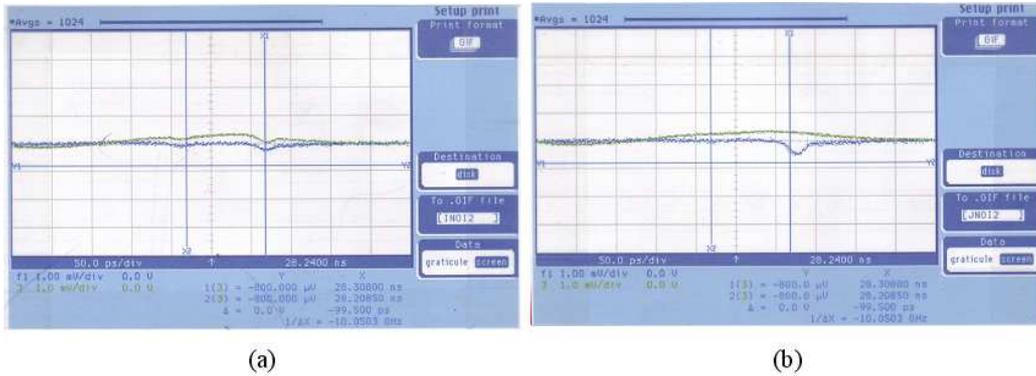


Figure 3.12: Screen grabs of the electrical signal produced by the 750 nm optical pulses with a bias voltage applied as measured on an HP54752A 50 GHz oscilloscope (blue trace). (a) shows the electrical pulse with the cursor line X1 aligned to the peak. (b) was taken some time later when a change in position had been spotted. The cursor, trigger settings etc. all remain unchanged. The cause for the occasional shifts such as this one has not been determined. Note that it was observed as a definite abrupt shift, not a slow drift over time.

750 nm peak would shift occasionally between oscilloscope measurements (Figure 3.12). These shifts appeared to be random and varied in magnitude by up to 10 ps. The cause was not identified; it may have been due to slight changes in alignment between measurements or errors in the triggering.

Once it was established that the probe and QUIC pulses were temporally close enough to be expected to fall within the range of the short delay line the dc bias was removed and both sides of the MSM device terminated with  $50 \Omega$  loads to prevent electrical reflections interfering with the measurement. A series of EOS scans were carried out over an interval of 30 ps (15 ps either side of the short-delay retroreflector position that corresponded to an overlap of the pump and probe pulses) at a resolution of 300 fs per point. Since it was surmised previously that use of the GaAsP photodiode might not provide the exact zero delay between the signals (Section 3.4.4) a series of scans were taken for various voltages applied to the PZT on the 1500 nm retroreflector to ensure the best chance of making a measurement when generating the maximum QUIC signal.

During the measurements there was a great deal of difficulty encountered in creating a good signal-to-noise ratio, for which there were several unavoidable causes. The loss of 50% or more of the 1500 nm power through the modulator is one. Modulation is required for successful lock-in amplification but results in the direct loss of 50% or more of the QUIC injected current. In addition, the amount of 850 nm available for EOS probing was around an order of magnitude lower than normal due to the requirement of splitting off a large proportion of the normal laser output to pump the OPO. Finally, it was not always possible to maintain a suitably high level of stability in the OPO over the duration of a scan, which could have caused variations in the generated electrical pulses that distorted the sampling measurement. The near-quadratic dependence of the QUIC injected current on the OPO output power meant the fluctuations of 5–20% which were sometimes seen would greatly reduce any signal levels and result in the scan having to be repeated. Despite these problems, however, in several scans a candidate signal peak emerged, positioned near the centre of the epoch where it would be expected. Four of the most promising are shown in Figure 3.13.

The peak is most clearly seen in (a) and (b); in (c) and (d) corroborating peaks are present at the same position but are too low to be conclusively distinguished from noise. Slight (0.5–2ps) variations in the positions of the peaks are assumed to be due to the same unknown factors which caused the shifts in the 750 nm dc bias-produced signal noted earlier. Adjusting the timebase to overlap the peaks and averaging (Figure 3.14) makes it easier to distinguish the peak from the noise. This may be considered artificial, although repeating the technique for randomly selected noise spikes fails to produce any such clear peaks. The structure of this peak is certainly similar to what might be intuitively expected, in that there is a fast initial rising edge followed by a somewhat slower, uneven falling edge.

### 3.5.6 EOS Measurement with 750 nm modulator

Due to the difficulties in obtaining significant signal encountered with the 1500 nm AOM it was decided to try the 750 nm AOM. This would maximise

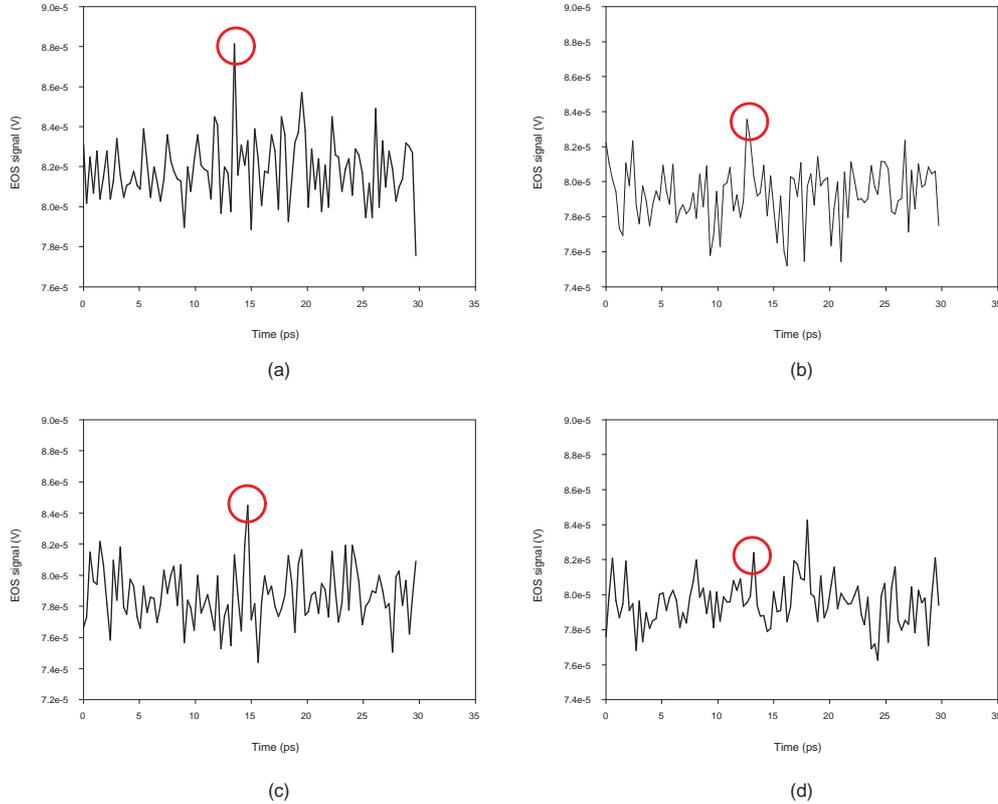


Figure 3.13: EOS measurements on the LT-GaAs device close ( $< 50 \mu\text{m}$ ) to the generation point of the QUIC signal. Note the candidate peak in each around the 13–15 ps point (circled).

the 1500 nm power available but of course reduce the 750 nm power. However, if successful there would be the option of attenuating the 1500 nm power to maximise the QUIC signal (c.f. Section 3.4.2).

The 1500 nm AOM was therefore removed from the 1500 nm arm of the QUIC system and the 750 nm AOM inserted in the 750 nm arm. The pulse paths were again realigned, first with each other (zero delay being found using the GaAsP photodiode as usual) and then realigned onto the photoconductive switch gap. Since the paths had changed as a result of the swap-over of AOMs the zero delay between the QUIC pulses and the probe pulses needed to be ascertained again, but as expected there was considerable

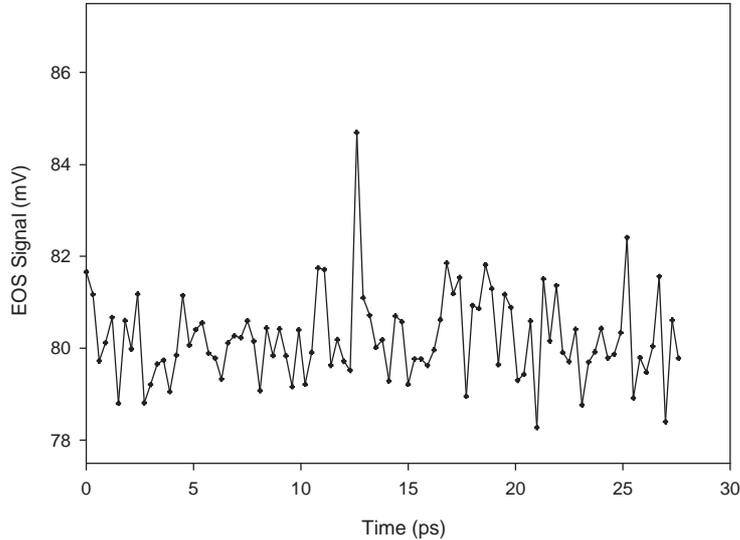


Figure 3.14: Average of results in Figure 3.13 a–d after shifting the timebases to overlap the candidate peaks.

difficulty in achieving this using an applied bias and an oscilloscope as the 750 nm signal was greatly reduced. Despite this, an approximation to the zero delay point was found and the EOS scan range increased to try and ensure the waveform of any QUIC electrical pulse was captured. The preliminary scans found no evidence of a signal, and further attempts were frustrated when it was found that the AOM had developed a fault and was failing to produce any modulated signal.

A second AOM was available and replaced the broken one in the 750 nm arm, requiring the beams to be realigned once more. Preliminary scans were carried out but the experimental work had to be stopped once again when the second AOM also failed. An investigation located the source of the fault to the power supply that supplied the 12 V required to power the AOM. It appeared that the power supply was occasionally spiking, sending excessive voltage to the connected AOM that damaged the internal circuitry. Unfortunately although the cause of the fault had been identified and the power supply could be replaced to prevent future occurrences, the lack of another

spare 750 nm modulator prevented this part of the experiment from being completed. The lack of conclusive results from the 750 nm measurements was frustrating, but this form of modulation should still be considered in any such future work.

### 3.6 Conclusions

A system was successfully designed and constructed with the purpose of generating QUIC currents within a transmission line device mounted on semiconductor. The creation of a QUIC current was demonstrated conclusively by measurements of the behaviour of charge accumulation across the device for several different kinds of variation in the generating optical pulses, each measurement agreeing with the theory. An existing EOS system was then adapted and integrated with the QUIC system to enable time-resolved measurements of the electrical signals generated and a number of promising measurements were made[29]. These experiments are highly important to this particular phenomenon and have, to current knowledge, never been carried out before. What has been highlighted is the difficulty of balancing the requirements of powering both a signal source and a probe using a single laser, as well as balancing the need for high frequency modulation with the consequent loss of power. Yet it must be realised that it is not as simple a matter as providing a separate laser source for the OPO and EOS system, since the problem is then introduced of needing to somehow synchronise both sources together and the inherent zero-jitter advantage of using a split source is inevitably lost.

The presence of a time-resolved QUIC signal is extremely encouraging, as is the size of the apparent pulse width. The issues are in the lack of power in the signal which are insufficient, it seems, to be of practical use in characterisation of high-speed instruments. It is important to realise that diverting power from the probing system is not enough, as a measurement of the pulse will still be required if the response of an instrument is to be characterised. The difficulties encountered when performing EOS suggests the best direction for further work is primarily one of improving the power of the OPO, its

stability, maximising the available power for EOS and investigating methods to enhance the EOS signal-to-noise ratio at low powers.

## Chapter 4

# Measurement of Ultrafast Pulses

Once there are sources capable of producing optical pulses with durations measured in femtoseconds, there is a requirement for means of exploring and measuring the characteristics of such pulses at femtosecond timescales. How is it known that the pulse really is as short as it is thought to be, or even the shape that it should be? How is it known if it is transform-limited or frequency-chirped, if it is a well-behaved Gaussian or a distorted or even doubled pulse? It is true that the spectrum provides some information about the pulse, but without knowing the spectral phase there is no way of knowing the temporal shape of the pulse — a pulse that has travelled along a kilometre of optical fibre can have the same spectrum as one that has not. The most the spectrum alone will tell the researcher is that the pulse cannot be any shorter than the value set by Equation 2.3, but if there is a need to know more (and often there is) then what is required is some scheme for comprehensive temporal measurement of the femtosecond pulses. Yet measurement of an event requires a shorter event itself, so how is it possible to measure the shortest event available? The answer lies in using the pulse to measure itself.

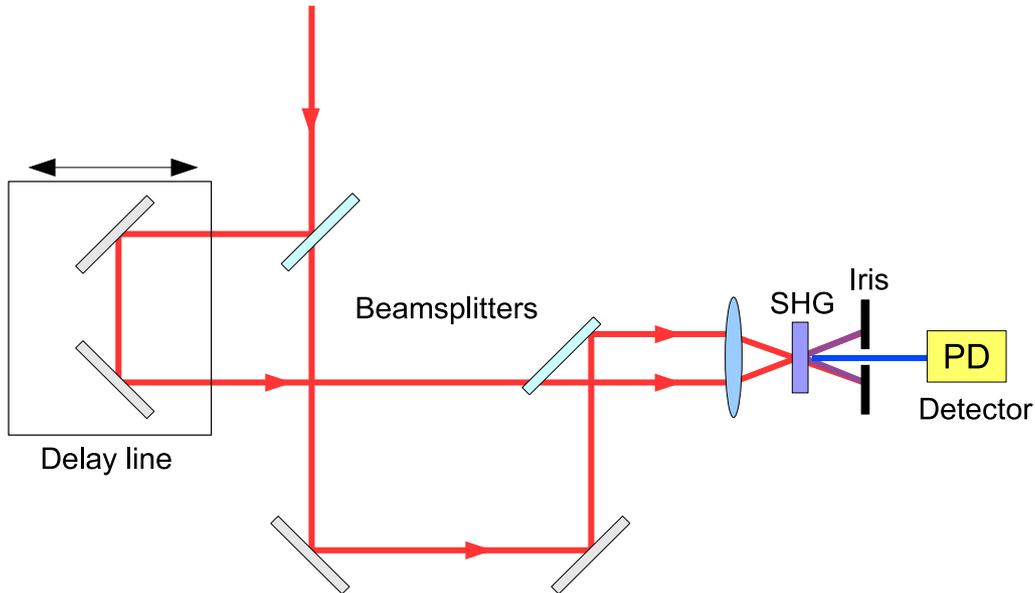


Figure 4.1: A schematic for an intensity autocorrelation. The incoming pulse is split into two replicas one of which is time-delayed with respect to the other before both replicas are focused non-collinearly onto a nonlinear crystal. The resulting second-harmonic signal is measured using a photodetector. A slit prevents any second harmonic produced by each individual pulse reaching the detector.

## 4.1 Autocorrelation

For a long time, the (unsatisfactory) answer to the measurement of ultrafast pulses was to deduce this from an autocorrelation of the pulse. A typical autocorrelation set-up is shown in Figure 4.1. The pulse is split into two intensity replicas which are then overlapped in some medium (in this case a nonlinear crystal) to produce a signal that corresponds to the mutual correlation of the overlapped pulses.

By varying the time-delay between the pulses one serves to gate the other and the autocorrelation is produced. In the case where the signal is the second-harmonic of the two pulses and the pulse paths to the crystal are non-collinear as shown in the diagram, and making the assumption that the nonlinear medium responds instantly in time, the polarisation signal

resultant towards the detector will be

$$P(\tau) \propto E(t)E(t - \tau) \quad (4.1)$$

where  $\tau$  is the delay between the two pulses. On making the reasonable assumption that the electric field of this SHG signal is proportional to the polarisation, the intensity of the signal will be

$$I(\tau) \propto |A(t)A(t - \tau)| \quad (4.2)$$

where  $A(t)$  is the complex amplitude of the pulse. It is usually the case that the detector used to measure the resulting signal has a response much slower than the duration of the pulse being measured and so the signal measured will be an integration over time, this result being the *intensity autocorrelation*:

$$I_{AC}(\tau) \propto \int_{-\infty}^{+\infty} I(t)I(t - \tau)dt \quad (4.3)$$

(This is also known as the *background-free* intensity autocorrelation, due to the lack of background signal from SHG of the individual pulses reaching the detector.) Since the intensity  $I(t) = |A(t)|^2$ , it is immediately apparent that the phase information of the pulse is lost. Also, on making the substitution  $t' = t - \tau$  it is seen that

$$I_{AC}(\tau) = \int_{-\infty}^{+\infty} I(t' + \tau)I(t')dt' = I_{AC}(-\tau) \quad (4.4)$$

and so the intensity autocorrelation is symmetric in time. Therefore, it follows that symmetric and asymmetric pulses can yield the same autocorrelation, and more information about the pulse has been lost. Yet the intensity autocorrelation still provides some information, since if the pulse shape is already known then deconvolution can be used to calculate the pulse duration. For some basic pulse shapes it is enough to divide the  $\Delta\tau$  of the autocorrelation by a simple factor — by  $\sqrt{2}$  for a Gaussian pulse, for example, or by 1.54 for a  $\text{sech}^2$  pulse.

Even though the intensity autocorrelation can provide a useful diagnos-

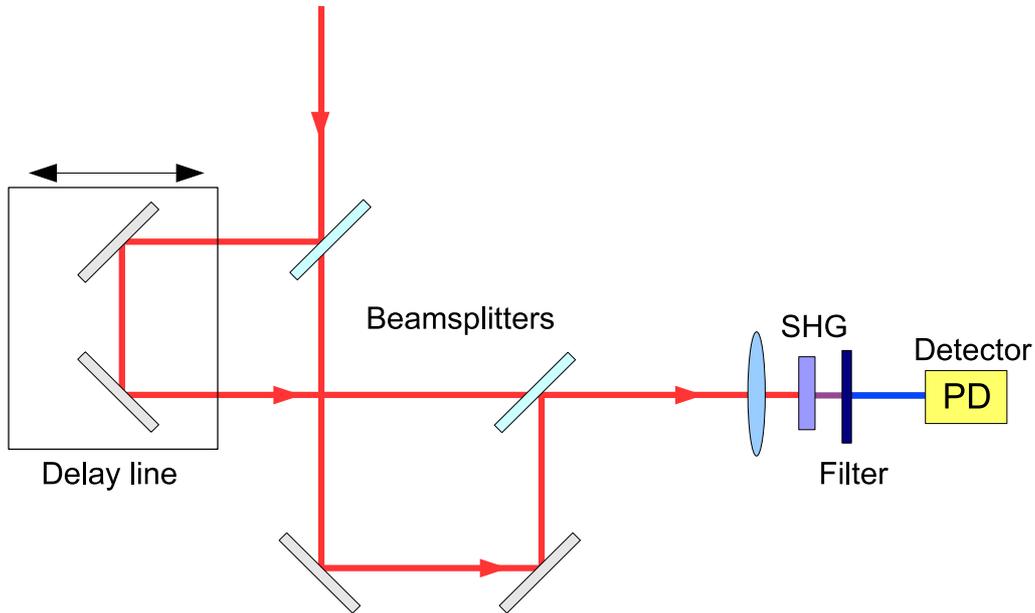


Figure 4.2: A schematic for an interferometric autocorrelation. The incoming pulse is split into two replicas one of which is time-delayed with respect to the other before both replicas are focused collinearly onto a nonlinear crystal. A filter blocks the fundamental signal and the second-harmonic is measured using a photodetector. This second harmonic will be a combination of the product of the two pulses and their individual second harmonics.

tic, the lack of phase and temporal shape information limits its utility and so other methods have been investigated to try to obtain more pulse information. One of the most important developments was the *interferometric autocorrelation* [30] (also known as the fringe-resolved autocorrelation, or FRAC). This differs from the normal intensity autocorrelation in that the two replica pulses overlap collinearly within the crystal, with the same polarisation. The autocorrelation is then resolved interferometrically. This set-up is shown in Figure 4.2.

In this situation the polarisation signal incident on the detector will be given by

$$P(t) \propto (E(t) + E(t - \tau))^2 \quad (4.5)$$

$$\propto (A(t)e^{i\omega t}e^{i\phi} + A(t-\tau)e^{i\omega(t-\tau)}e^{i\phi})^2 \quad (4.6)$$

$$(4.7)$$

where  $\phi$  is the phase of the carrier wave relative to the amplitude envelope and  $\omega$  is the carrier frequency. As with the intensity autocorrelation, the slower detector integrates over time and this produces the interferometric autocorrelation equivalent to (4.3)

$$I(\tau)_{FRAC} \propto \int_{-\infty}^{+\infty} \left| A^2(t)e^{i2\omega t} + A^2(t-\tau)e^{i2\omega(t-\tau)} + 2A(t)A(t-\tau)e^{i\omega t}e^{i\omega(t-\tau)} \right|^2 dt \quad (4.8)$$

Expansion of (4.8) leads to four terms in the integral, which can be separated out into four integrals that combine to make the interferometric autocorrelation

$$I(\tau)_{FRAC} = I_{back}(\tau) + I_{int}(\tau) + I_{\omega}(\tau) + I_{2\omega}(\tau) \quad (4.9)$$

where

$$\begin{aligned} I_{back}(\tau) &= \int_{-\infty}^{+\infty} |A(t)|^4 + |A(t-\tau)|^4 dt \\ &= 2 \int_{-\infty}^{+\infty} I^2(t) dt \end{aligned} \quad (4.10)$$

$$\begin{aligned} I_{int}(\tau) &= 4 \int_{-\infty}^{+\infty} |A(t)|^2 |A(t-\tau)|^2 dt \\ &= 4 \int_{-\infty}^{+\infty} I(t)I(t-\tau) dt \end{aligned} \quad (4.11)$$

$$I_{\omega}(\tau) = 4 \int_{-\infty}^{+\infty} \Re \left( (I(t) + I(t-\tau))A^*(t)A(t-\tau)e^{i\omega\tau} \right) dt \quad (4.12)$$

$$I_{2\omega}(\tau) = 2 \int_{-\infty}^{+\infty} \Re \left( A^2(t)(A^*(t-\tau))^2 e^{i2\omega\tau} \right) dt \quad (4.13)$$

The first term (4.10) is clearly a constant value, and is simply the background signal of the autocorrelation. The second term (4.11) is the intensity autocorrelation (4.3) multiplied by a factor of 4. The last two terms (4.12) and (4.13) are the two coherence terms that oscillate with  $\omega$  and  $2\omega$ , respectively. A typical interferometric autocorrelation is shown in Figure 4.3. Note again the same time-symmetry as for the intensity autocorrelation:

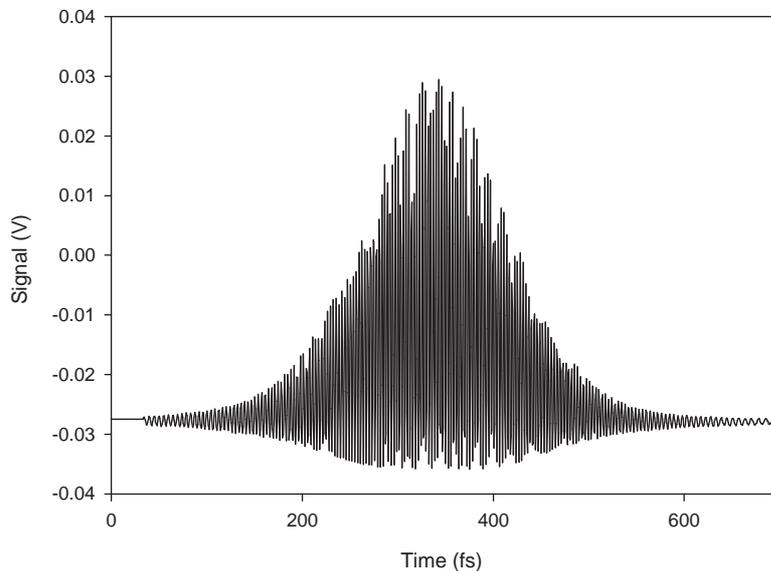


Figure 4.3: Interferometric autocorrelation of a 'Mira' Ti:S laser. Note the 8:1 ratio between the top-to-bottom distance and the bottom-to-baseline distance. This ratio provides a useful check that the autocorrelation has been made correctly.

$$I_{FRAC}(\tau) = I_{FRAC}(-\tau).$$

An important feature of the interferometric autocorrelation arises from the integral values at  $\tau = 0$

$$I_{int}(0) = I_{\omega}(0) = I_{2\omega}(0) = I_{back}(0) \quad (4.14)$$

$$I_{back}(0) = 2 \int_{-\infty}^{+\infty} |A(t)|^2 dt \quad (4.15)$$

Therefore  $I_{FRAC}(0) = 8 \int |A(t)|^4 dt$ . As  $\tau \rightarrow \pm\infty$  we also see that  $I_{FRAC} \rightarrow 0$  and so there is an 8:1 ratio between the peak of the interferometric autocorrelation and the background level. This feature proves extremely useful for checking the alignment of an autocorrelation system, since an absence of this 8:1 ratio indicates an error somewhere.

Another useful feature of the FRAC is the ability to detect chirp (Sec-

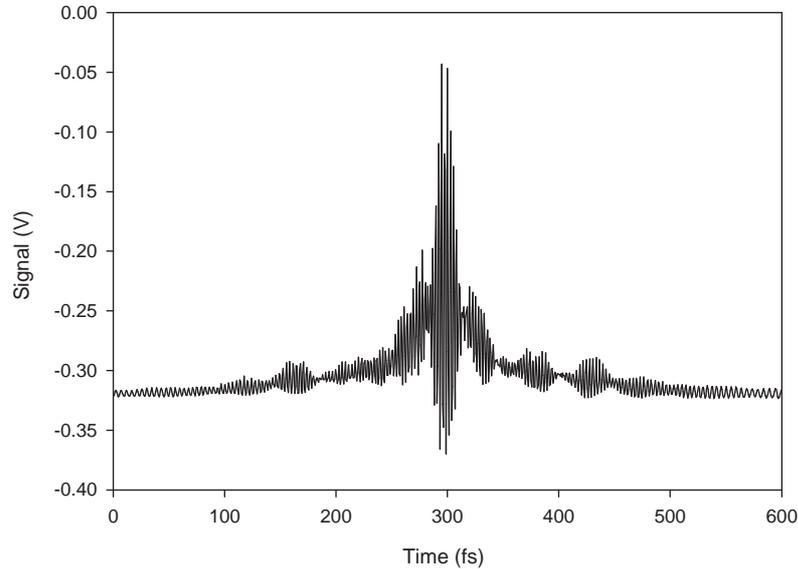


Figure 4.4: Interferometric autocorrelation of the 'Femtolaser' Ti:S <12 fs laser. The extra-cavity dispersion compensation required for this laser is absent, resulting in highly chirped pulses. The effect of the chirp can be seen in the wings of the autocorrelation where the interference fringes have reduced, to the point of loss in places, and the envelope is approaching that of the intensity autocorrelation.

tion 1.1), since for a chirped pulse it can be shown [30] that destructive interference in the wings of the autocorrelation will cause the FRAC envelope to merge with the intensity autocorrelation envelope. Figure 4.4 shows an example measurement of a chirped pulse.

## 4.2 Frequency Resolved Optical Gating

*Frequency-resolved optical gating*[31] (FROG) is a correlation-based technique that is capable of extracting both amplitude *and* phase information about the pulse. The core measurement at the heart of FROG is the *spectrogram*

$$S(\omega, \tau) = \left| \int_{-\infty}^{+\infty} E(t)g(t - \tau)e^{-i\omega t} dt \right|^2 \quad (4.16)$$

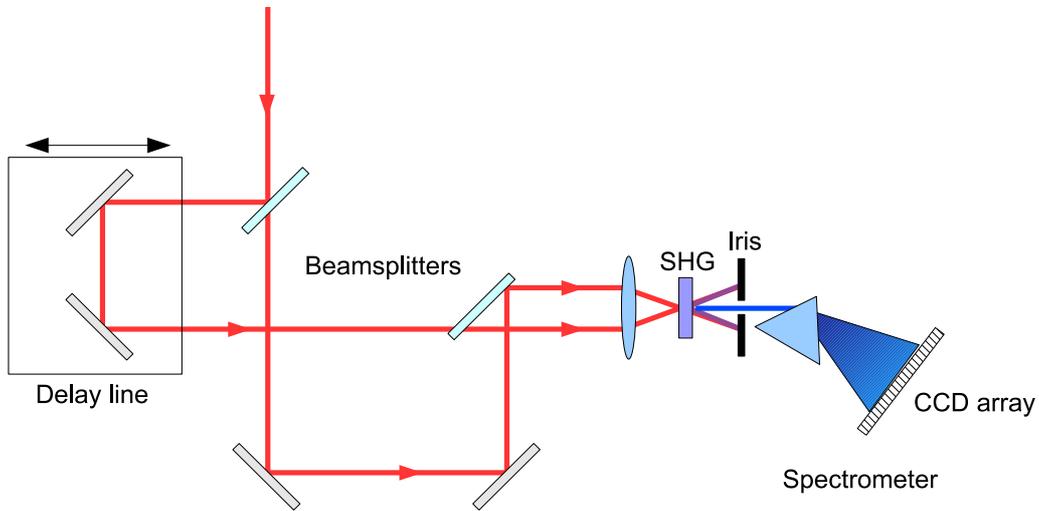


Figure 4.5: Schematic of a SHG frequency-resolved optical gating set-up. Two pulse replicas are created by a beamsplitter and one delayed with respect to the other. A second harmonic signal is generated in the manner of an autocorrelation but instead of measuring the intensity of the signal it is spectrally resolved. The resulting plot of wavelength (or frequency) against delay is the spectrogram.

where  $g(t - \tau)$  is a gate pulse delayed by time  $\tau$  relative to the optical pulse  $E(t)e^{-i\omega t}$ . When the gate pulse is a replica of the electrical pulse to be measured this is, in essence, a measurement of the spectrum of the SHG intensity autocorrelation. A schematic of an SHG FROG measurement is shown in Figure 4.5

(Other forms of FROG measurement exist such as Polarisation-Gating FROG (PG-FROG), Third-Harmonic-Generation FROG (THG-FROG) and more besides, but they all share the basic theme of spectrally measuring a nonlinear correlation process. The interested reader is pointed towards Trebino's book [32] which details many of these, but within this thesis only SHG-FROG is considered.)

From Equation 4.16 it can be seen that the spectrogram will be a three-dimensional plot of intensity against frequency and time, but the question is how can the  $E(t)$  information be retrieved from it? The key step here is to consider the spectrogram from the viewpoint of the Fourier domain. Taking

$U(t, \tau) = E(t)E(t - \tau)$  and considering the following representation in the Fourier domain

$$U(t, \tau) = \int_{-\infty}^{+\infty} \tilde{U}(t, \varpi) e^{-i\varpi\tau} d\varpi \quad (4.17)$$

(i.e. the Fourier transform of the field with respect to  $\tau$  and not  $t$ , hence the use of  $\varpi$ ) means that, if  $\tilde{U}(t, \varpi)$  is known, a Fourier transform can be used to produce  $U(t, \tau)$ . The substitution  $t = \tau$  can then be made to produce  $U(t, t) = E(t)E(0)$ . Since  $E(0)$  is a constant factor it can be scaled to unity to give  $U(t, t) \equiv E(t)$ , so finding  $\tilde{U}(t, \varpi)$  will give  $E(t)$ . Substituting Equation 4.17 into Equation 4.16 based on this gives

$$S(\omega, \tau) = \left| \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} \tilde{U}(t, \varpi) e^{-i\omega t - i\varpi\tau} dt d\varpi \right|^2 \quad (4.18)$$

This shows that the spectrogram provides the magnitude of the 2D Fourier transform of  $\tilde{U}(t, \varpi)$ . Thus if the phase could also be obtained it would be possible to determine  $\tilde{U}(t, \varpi)$  through inversion and from there find  $E(t)$ . Determining the phase is known as a two-dimensional phase-retrieval problem [31], which means that, somewhat surprisingly, a near-unique solution exists and can be found provided there is finite support, i.e.  $\tilde{U}(t, \varpi)$  is zero outside a given region of  $t$  and  $\varpi$ . The description ‘near-unique’ is used as there are ambiguities, solutions that differ by a common absolute phase factor, or are inverted, for example. There also exists the possibility of non-trivial ambiguities (i.e. other solutions that do not correctly represent  $E(t)$ ) but these are rare to the point of improbability.

Retrieval of the pulse from the spectrogram through the two-dimensional phase-retrieval problem can be achieved by an iterative algorithm[33] (Figure 4.6). A guess pulse is used to generate a pseudo-spectrogram and the actual measured spectrogram used to suitably modify the pseudo-spectrogram by applying some constraint, such as replacing the magnitude of the guess with the square-root of the intensity of the measured.  $E(t)$  is derived from the modified spectrogram to produce a new guess pulse, and the procedure reiterates until it converges towards a solution. Differing implementations or variations of the basic algorithm may provide additional advantages and

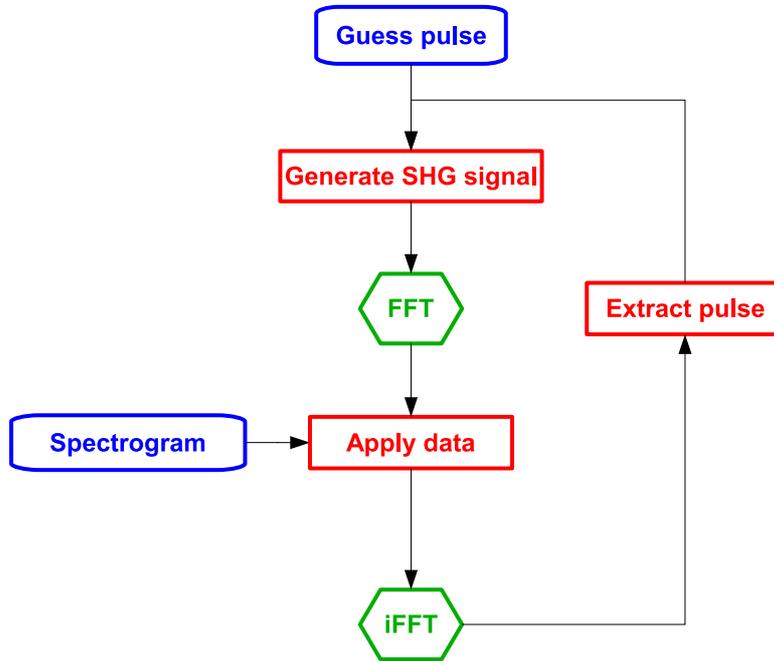


Figure 4.6: Flowchart of the FROG retrieval algorithm.

disadvantages, but the overall idea remains the same — iteration to an amplitude and phase solution by applying constraints to a guess pulse in the frequency domain. FROG retrieval is now approaching maturity as a tool, and has been used to measure pulses of varying complexity as far down as 4.5 fs FWHM[34]. For these reasons — the wealth of existing research, the relative simplicity of measurement and retrieval and the scope for further refinements — it was decided to design, construct and develop an in-house version of a SHG-FROG system at NPL.

### 4.3 The NPL FROG system

The FROG system for NPL was based on SHG-FROG and primarily intended for use with Ti:S laser pulses of between 10 fs and 100 fs in duration. It would also need to be compact enough to be portable, allowing for measurement of different lasers. The basic design (Figure 4.7) is simple enough, being similar

to an autocorrelator. A beamsplitter divides the input pulses along two arms, one fixed in length and the other variable through both a manual translation stage (for coarse adjustment) and a translation stage controlled by a PZT<sup>1</sup> (further details are given in Section 4.3.2). Both pulses are then recombined by a second identical beamsplitter to be parallel but non-collinear. A mirror is used to focus the pulses onto a SHG crystal (a parabolic is used as opposed to a lens to minimise dispersion) which produces the second harmonic of each pulse plus (where the pulses temporally overlap) the second-harmonic FROG signal. A matching parabolic mirror collimates the second harmonics and directs them to a detection system for measuring the spectrum of the FROG signal. An iris prior to the detector eliminates the basic second harmonics to reduce the background level of signal.

Silver mirrors were used for all planar mirrors in the FROG system. The focusing parabolic was gold and the collimating parabolic and all subsequent mirrors were aluminium. The software for setting of the delay line, measurement of the FROG signal and storage of the data was written in LabVIEW (Appendix A). This software allowed setting of the integration time of the detection system; the ability to make a background signal measurement and remove said background from subsequent measurements; user control of the delay line; and automatic measurement of spectrograms. The software also allowed the storage of spectrograms and individual measurements at the detector (e.g. the signal at zero delay).

### 4.3.1 Beamsplitters

The arrangement of the beamsplitters within the FROG system requires more careful thought than may be apparent at first. The first consideration is that ideally each pulse replica that arrives at the SHG crystal should have travelled identical paths. This means the arrangement should be such that each pulse has undergone one reflection from a beamsplitter and one transmission, and the transmission should be from the same face in each case (Figure 4.8). (This is especially important in the case of the shortest (10 fs) pulses, which

---

<sup>1</sup>Newport AD-100

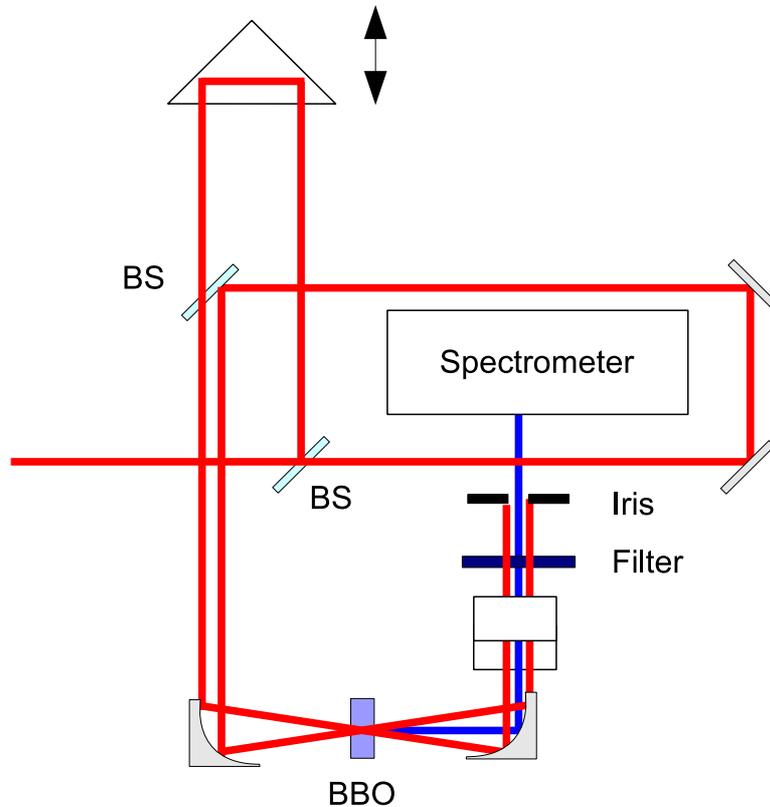


Figure 4.7: Schematic of the NPL FROG system. Beamsplitters divide and recombine the input pulses with the minimum of dispersion. The use of a focusing parabolic mirror also avoids unwanted contributions to dispersion. An identical parabolic mirror is used for collimation after the SHG crystal purely for ease of alignment. A persicope is used to enable the spectrometer be included on the breadboard by locating it below the path of one FROG arm.

are more susceptible to dispersion.) If possible the same beamsplitter can even be used, but often this is difficult to arrange correctly and it should be enough that the beamsplitters are identical.

The second consideration is the polarisation of the pulses. The FROG signal will depend on the intensity of the pulses when they arrive at the SHG crystal and beamsplitters generally reflect S (parallel to the plane of incidence) and P (perpendicular to the plane of incidence) polarisations in

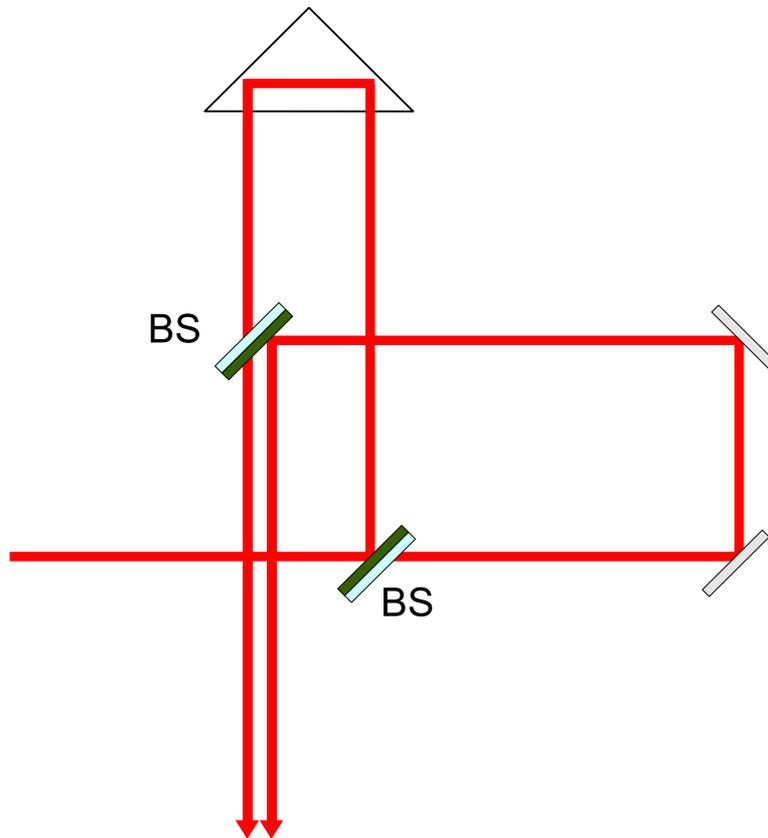


Figure 4.8: Correct arrangement of the (identical) beamsplitters to present equal paths to each pulse replica. The darkened sides of the beamsplitters represent the coated sides. Arranging them in this manner ensures each replica undergoes one transmission and one reflection.

different proportions. Nor will beamsplitters divide pulses perfectly 50:50, but may transmit more or less than 50% (with the corresponding difference in reflectance). Therefore it makes sense, from the perspective of maximising the FROG signal, to calculate the optimum arrangement of the beamsplitters according to polarisation.

The first set of beamsplitters received for the FROG system were found to be unsuitable due to manufacturing errors. Specifically, one beamsplitter was visibly thicker in substrate than the other by some margin, and also reflected only 15% of the incident light at any polarisation. A new order was placed

with a different company and the replacements found to be far more identical, each splitting the incoming pulses 63:37 in favour of the transmission in S polarisation and 25:75 for P polarisation.

In the case where P polarisation is used, it is clear that the power reaching the crystal from each arm will be  $0.25 \times 0.75 = 0.19$  of the incident power compared to  $0.63 \times 0.37 = 0.23$ , indicating that S polarisation is preferable for the input pulses. Given the drastic difference from 50:50 for P polarisation there is, however, a further option. By using a periscope arrangement in one arm to rotate the polarisation the greater reflectivity of the beamsplitters for P polarisation can be exploited (Figure 4.9) to produce levels at the crystal of 0.47 and 0.23 of the input power. The FROG signal varies with the product of the square of the power in each arm, as calculated (SHG from two combining pulses) and confirmed by measurement (Figure 4.10), giving greater weight to the advantages of effecting the polarisations this way. However, as the BBO crystal has been designed for Type-I phase matching SHG the polarisations of the incoming pulses need to be aligned with the e axis of the crystal for correct SHG. If the pulses are cross-polarised then SHG will require that the e axis is aligned at 45 degrees to each and so the power contribution from each arm will be reduced by  $\frac{1}{\sqrt{2}}$ . Working through this in comparison with S polarised input and no rotation within the FROG gives

$$0.23^2 \times 0.23^2 = 0.0028 \quad (4.19)$$

with no rotation and

$$\left(0.23 \times \frac{1}{\sqrt{2}}\right)^2 \times \left(0.47 \times \frac{1}{\sqrt{2}}\right)^2 = 0.0029 \quad (4.20)$$

when incorporating rotation, i.e. there is very little difference between the two. However, this is dependent on the beamsplitter ratios used for each polarisation being correct.

To confirm this, the FROG was first constructed with periscopes in place and the powers measured in each arm at the crystal. For an input power of 337 mW from the Mira Ti:S laser, powers of 74 mW were measured in the

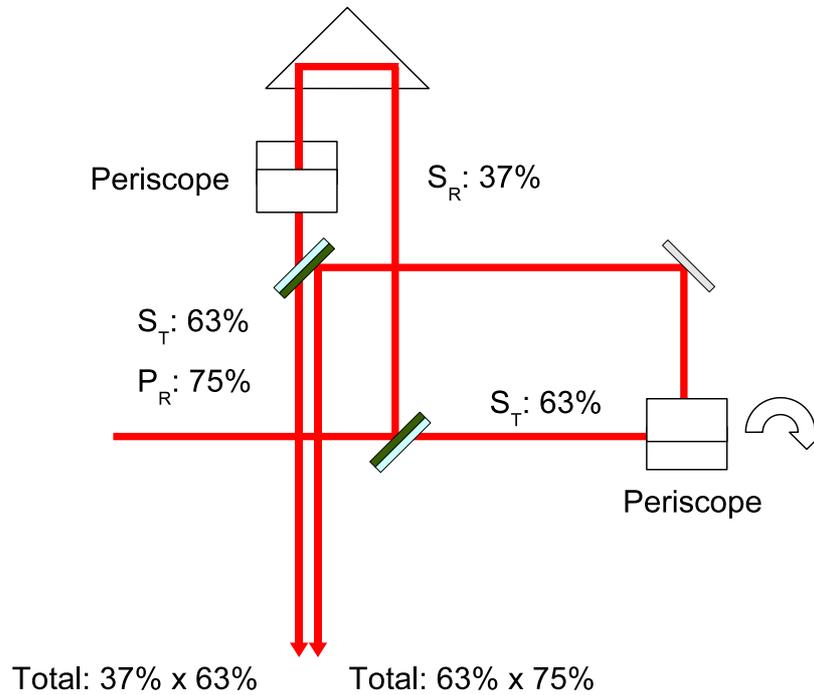


Figure 4.9: Incorporation of periscopes to maximise the power of each pulse replica. The periscope in one arm is used to rotate the polarisation of one replica from S to P to take advantage of the high reflectivity of the second beamsplitter to P polarisation. The second periscope is required to match the beam heights. Without the rotation the replica power would be  $63\% \times 37\%$  instead of  $63\% \times 75\%$ .

arm with no polarisation rotation and 138 mW in the arm with rotation. This corresponds to 22% and 41%, which can be compared to the predicted powers of 23% and 47% — close but slightly lower than expected. Removing the periscopes and realigning produced (for a slightly changed input of 339 mW) powers of 77 mW and 79 mW, corresponding to 23% each, more in line with what was correctly expected. Comparing the zero-delay signal measured in each configuration (which after all is the incentive behind the whole exercise) in Figure 4.11 shows, remarkably<sup>2</sup>, virtually identical performance in each case. (A slight offset in central wavelength was noted, but this was found to be due to a slight difference in the central wavelength of the source laser as

<sup>2</sup>in that the theory and the experimental match so closely!

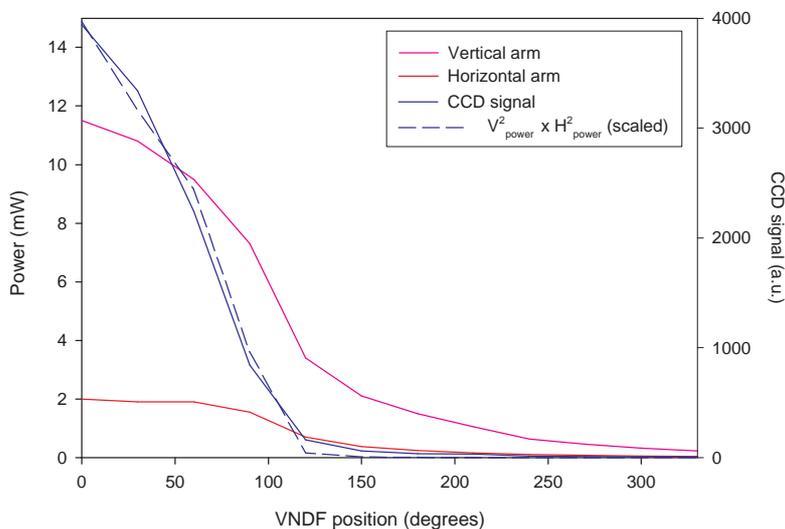


Figure 4.10: Plot showing how the peak signal level at the spectrometer varies with the power in each arm of the FROG. The red and the pink curves plot the arm powers as measured with a power meter at the location of the SHG crystal with a single VNDF used to attenuate both beams. The blue curve is the peak FROG signal measured at the CCD array (plotted with respect to the right-hand axis) for the same positions of the VNDF and the dashed blue curve is the predicted signal level based on  $\alpha V^2 H^2$ , where  $\alpha$  is some scaling value to convert from mW to CCD counts. It can be seen that the peak CCD signal follows the predicted values.

measured with a spectrometer.)

A further check on the correctness of the beamsplitters was carried out by measuring the spectrum of the input pulses, both before the FROG and after. The possibility of measuring the pulses after the FROG arrangement arises due to the fact that the final beamsplitter transmits *and* reflects parts of the pulses, meaning that some of the pulse energy is directed away from the FROG as well as towards the crystal (Figure 4.12). (The reason that the pulses cannot simply be measured at the position of the SHG is a practical one — even with the crystal removed there is no room to squeeze the appropriate spectrometer into the compact set-up. This method also permits the laser spectrum to be measured simultaneously with the FROG spectrogram.)

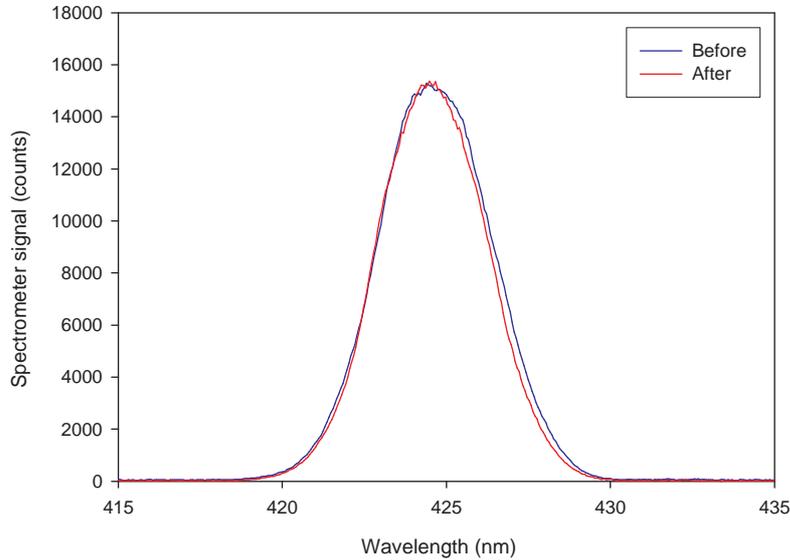


Figure 4.11: Zero delay signal measured before and after removal of the periscopes in the FROG arrangement. An offset of 1 nm was required due to a change in the central wavelength of the laser between measurements. The CCD integration times were identical in each case and no scaling has been applied. The agreement between the two measurements indicates the use of periscopes provided neither advantage nor disadvantage.

This measurement, shown in Figure 4.13, shows that the spectral amplitude is not being distorted by the beamsplitters, indicating that the pulses arriving at the crystal are the same as those entering the FROG (neglecting additional phase dispersion introduced by the beamsplitters and assuming no spectral amplitude or phase changes are introduced by the Au-coated parabolic mirror). This ‘wasted’ signal output was also found to have another use for calibrating the delay line, as explained in Section 4.3.2.

To conclude, it was found that although a configuration that takes advantage of the greater reflection ratio of a beamsplitter for P polarisation might be able to give greater FROG signal level, for the beamsplitters provided no such advantage is offered. As a result, the configuration without periscopes was decided on due to the simpler alignment and reduced number

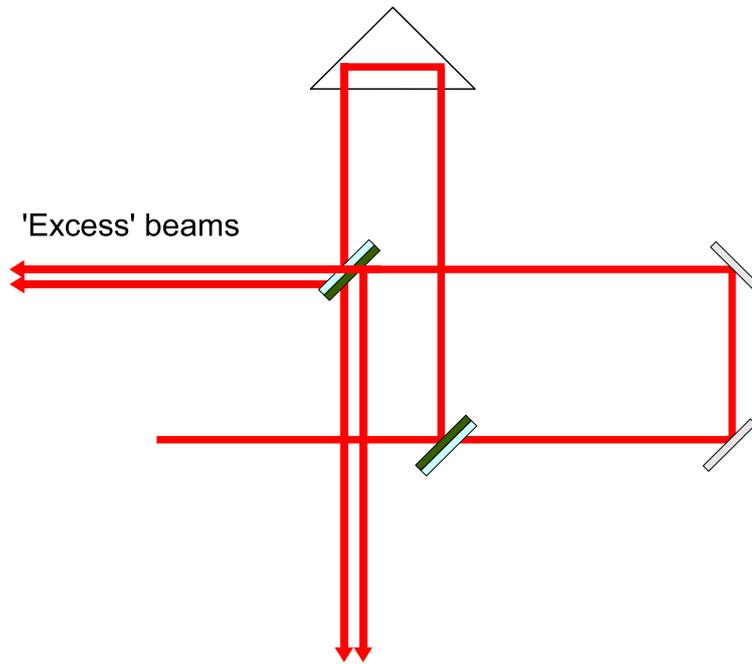


Figure 4.12: Schematic diagram showing how some of the pulse energy exits the FROG from the second beamsplitter. Normally this energy is discarded but it can be used to monitor the spectrum of the pulses, prior to arrival at the SHG crystal, for any errors introduced by the FROG components.

of reflections within the system.

It should be apparent that typical beamsplitters such as those used here consist of coated substrate and therefore as the pulses pass through the substrate glass they will undergo dispersion. Although arranging the beamsplitters as described in this section ensures that both pulse replicas undergo identical dispersion this still means that the pulses that arrive at the SHG crystal will not be the same as those that arrived at the input to the FROG. The ideal case would be to have no dispersion at all, and the closest realisation of this is the use of pellicle beamsplitters. A pellicle beamsplitter consists of thin nitrocellulose film stretched to tension over a supporting ring. Their extreme ( $5 \mu\text{m}$ ) thinness all but eliminates the secondary reflections that occur with traditional beamsplitters and, more importantly from the point of view of the FROG, greatly reduces dispersion in the pulses trans-

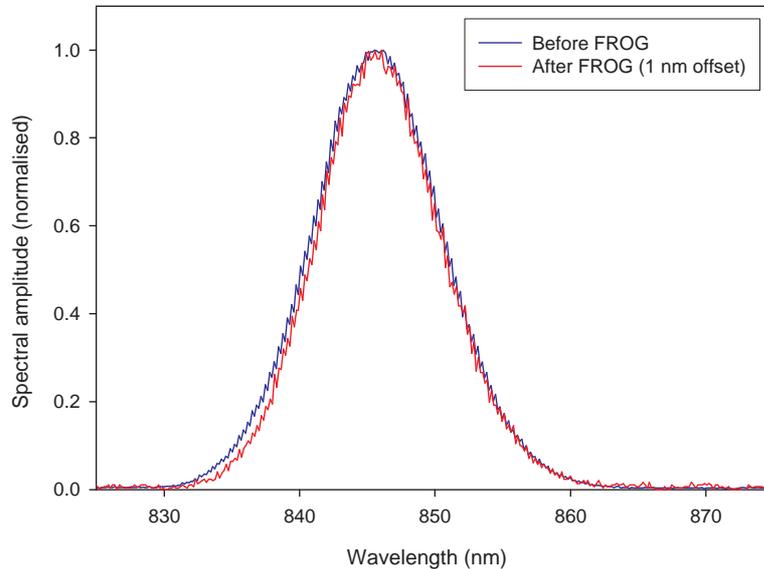


Figure 4.13: Spectrum of the laser pulses measured at the entrance to the FROG and at the additional exit. The two spectra have been scaled to account for the differing power levels. The agreement between the two confirms the suitability of the FROG components for their purpose.

mitted through the beamsplitter. The penalty is an extremely fragile optical component that cannot even be touched, and moreover the tension of the film means it can be sensitive to vibration and resonance effects. Even so, a pair of such beamsplitters, suitably coated for pulses centred in the 800–850 nm range, has been installed in the NPL FROG system as can be seen in Figure 4.22. Although their fragility might be thought a disadvantage for a system where portability is the aim it can be said that this FROG system is robust because it has been moved across optical benches and between laboratories without incident so far!

### 4.3.2 The Delay Line

The delay line used was a Newport AD-100 PZT (with associated Newport ESA-C motion controller). Voltages applied to the PZT, ranging from 20 V

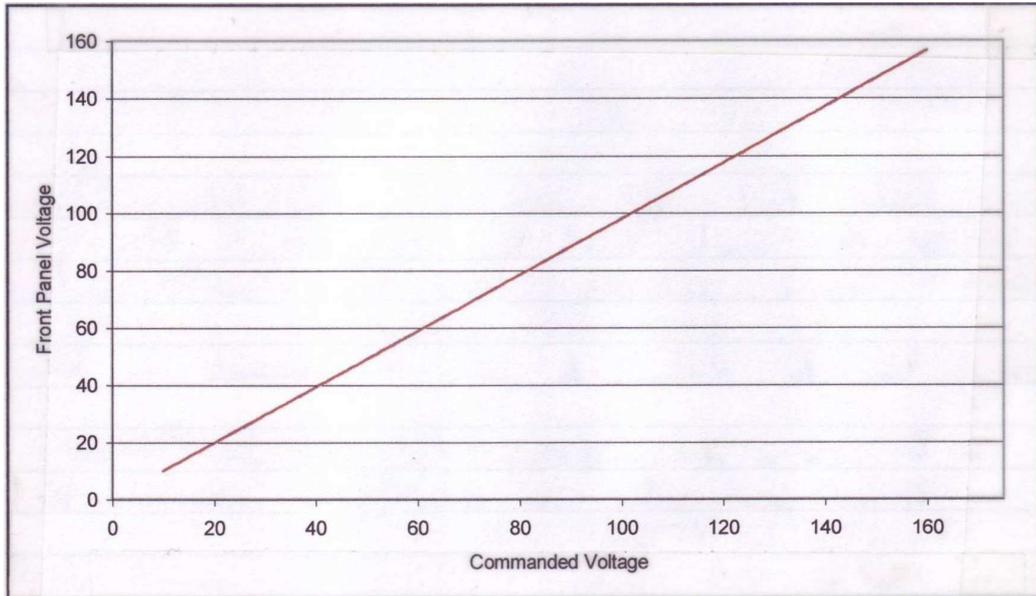


Figure 4.14: Comparison of the voltage issued to the Newport ESA-C via GPIB command with the voltage displayed on the front panel after the command has been processed. A large range of measurements has been plotted, but the close overlap indicates that the discrepancy between the two voltages is consistent. No reason for the discrepancy could be found.

to 150 V, adjust the delay line over a range of approximately 150  $\mu\text{m}$ . The resolution of the voltage supplied by the controller was 0.1 V digitally displayed and could be set either manually, via a dial on the front panel, or remotely using the GPIB interface.

It was noted that when setting the voltage remotely the front panel displayed a value different to that requested. Investigation of this demonstrated that the difference was consistent (Figure 4.14). No cause could be found, but there was at least confidence that, even if the absolute voltage at the PZT was unknown (there being no reason to believe the front panel voltage over the commanded voltage or vice versa), it was at least being set consistently. Since all control was being carried out remotely, any voltage values referred to in this thesis are those that were set remotely.

An approximate voltage/displacement relation was provided by the manufacturer in the form of a plot in the user manual (Figure 4.15). As precise

knowledge of the actual delay line displacement is required for correct FROG retrieval an interferometer setup was used to calibrate the delay line (Figure 4.16). This was based around a HP5519A laser head and an HP10887A interface card. As the interface card was based on the ISA interface which was now obsolete an ISA-to-USB converter box (Agilent 55292A) designed for this specific card had to be used. The setup is typical of that for interferometry, using a polarising cube beamsplitter and a fixed reference arm. With this system the displacement was repeatedly measured for a number of different voltages. These measurements were made for voltages in sequence, both increasing and decreasing, and also for random values, to limit the effects of hysteresis. The relationship found between the voltage and the displacement is shown in Figure 4.17.

Ideally the interferometer would be a permanent part of the system, used by the software to correctly determine the delay during spectrogram acquisition. Two issues, however, prevented this. The first and main issue was that the interferometer could not be integrated with the FROG software. Although a DLL<sup>3</sup> code module was provided, and despite lengthy correspondence with one of the chief software engineers from the manufacturer[35], all attempts to produce code capable of reading the interferometer within LabVIEW proved fruitless. The exact cause was difficult to pinpoint but seemed related to an inability of the DLL to set the interferometer internal timing. The second issue was that the FROG system should ideally be portable, and including the laser head for the interferometer would have reduced significantly the portability of the existing system. (As well as requiring a new, larger, optical breadboard the laser head weighed as much again as the existing system.) Instead, the measurements shown in Figure 4.17 were used to form a look-up table for the FROG software so that requested delays could be converted to correct voltages. The software used linear interpolation for delays corresponding to values between those present in the look-up table.

As an additional check of the delay line calibration, the ‘excess’ beams leaving the FROG system (c.f. Section 4.3.1) were found to have a novel use. When monitoring the spectrum by guiding one of the beams free-field

---

<sup>3</sup>Dynamic-link library. A Windows file that provides software functionality.

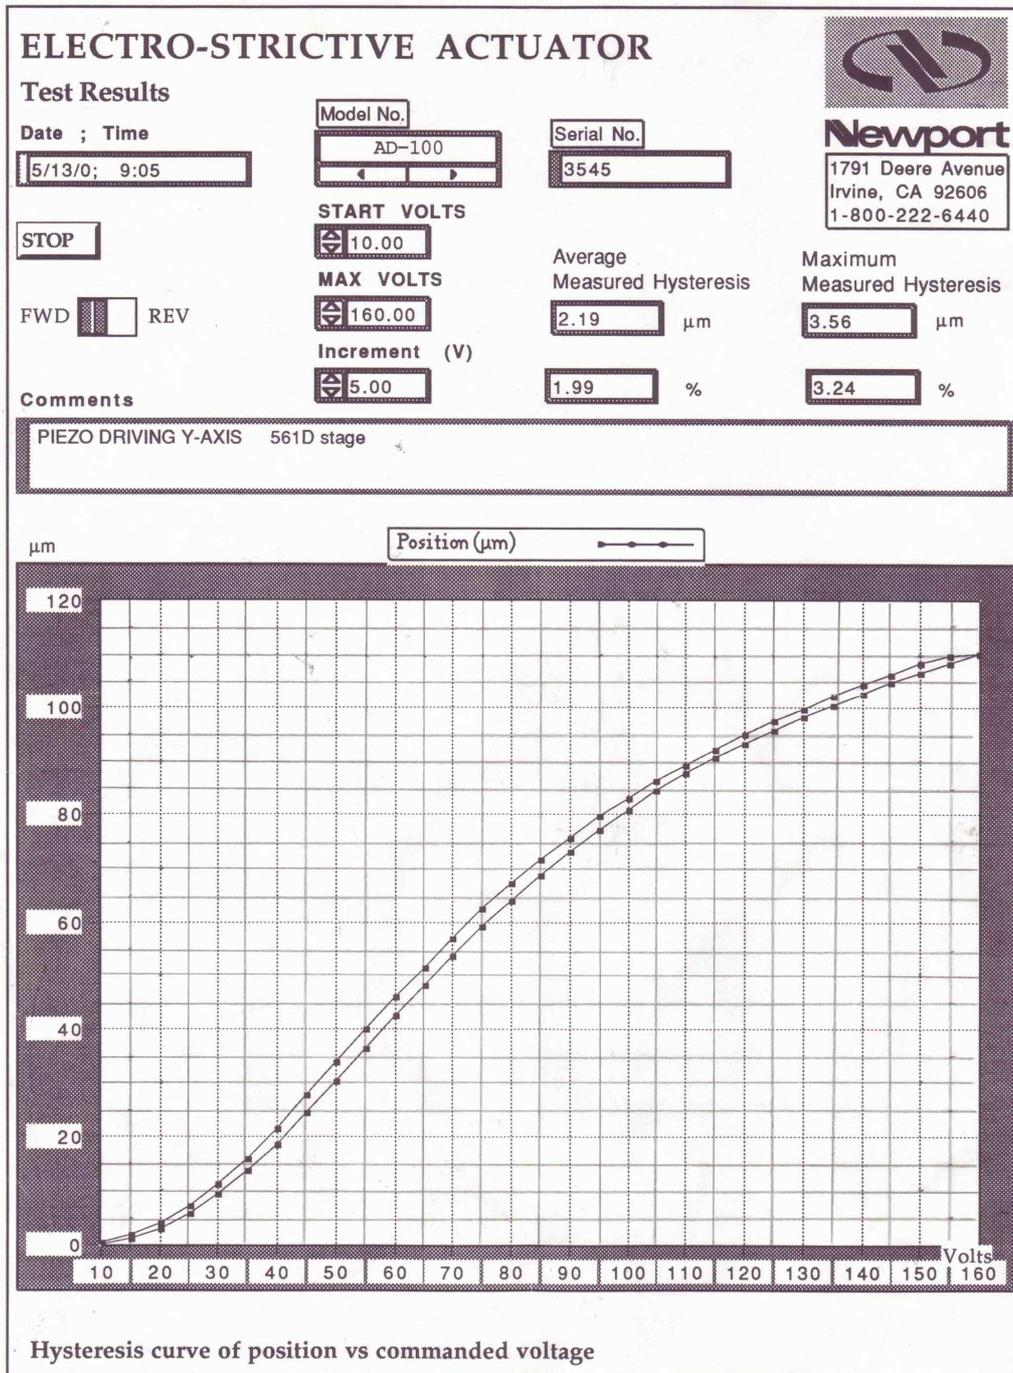


Figure 4.15: Voltage/displacement relationship for the AD-100 actuator as provided by Newport.

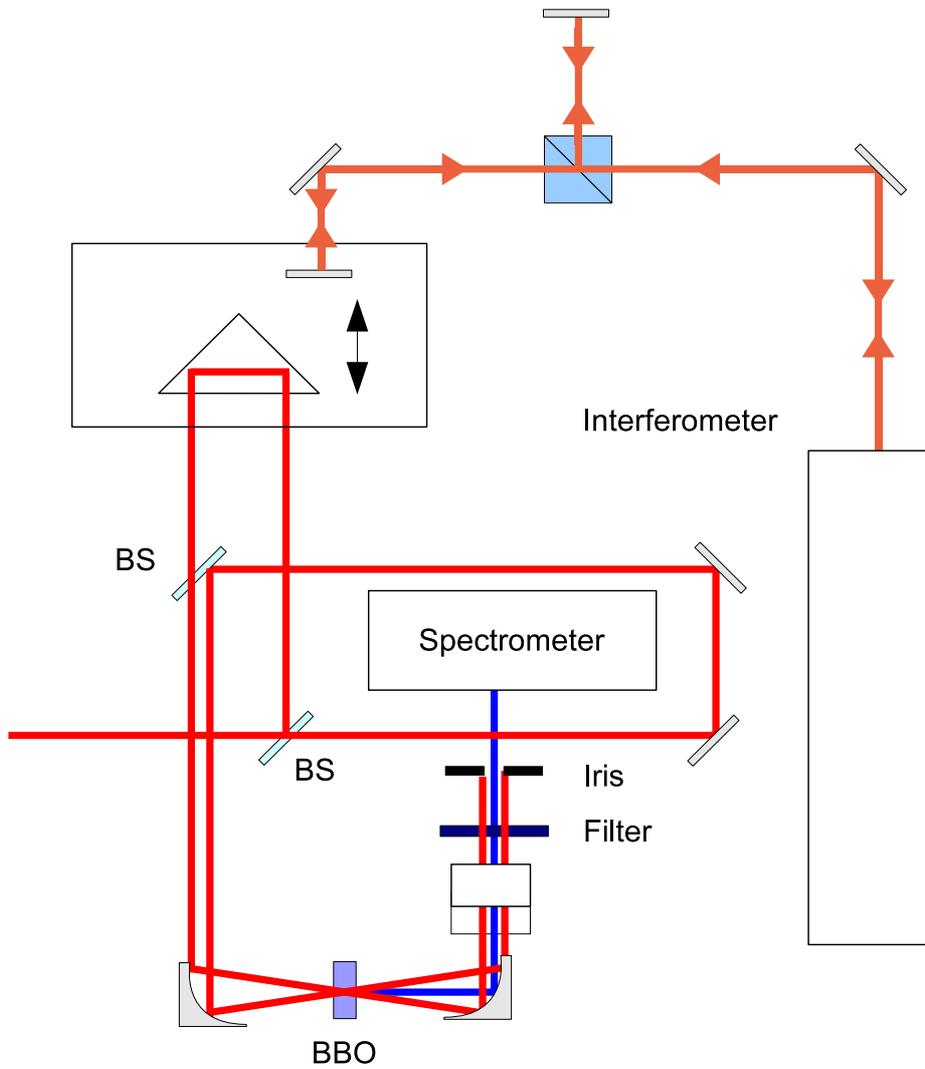


Figure 4.16: Schematic of the interferometer arrangement for calibration of the FROG delay. The arrangement is typical for this kind of measurement, using a polarising cube beamsplitter and a fixed reference arm. The mirror for the moving arm is installed on the translation stage bearing the delay line.

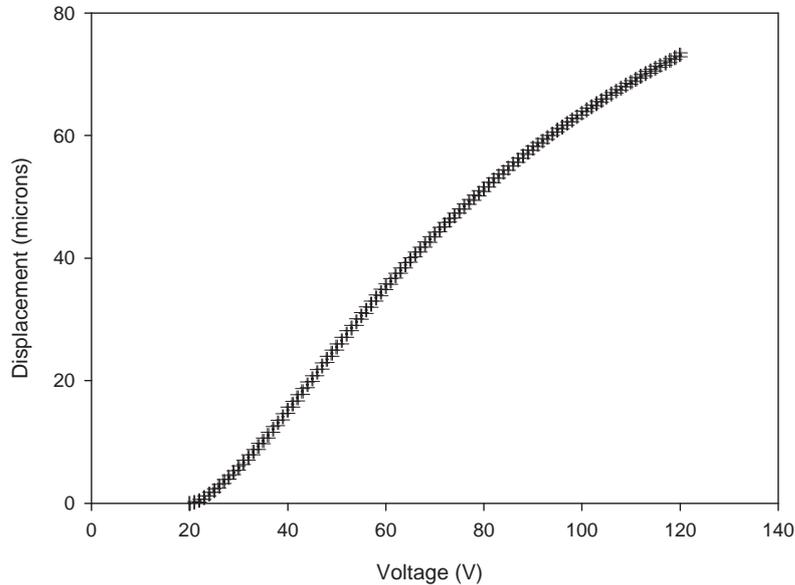


Figure 4.17: Plot showing the displacement of the FROG delay determined by the interferometer as a function of the voltage. The error bars account for the error in the voltage setting ( $\pm 0.1$  V and the standard deviations of the displacements.)

into the spectrometer some form of fringing effect was noticed at positions away from the zero delay (Figure 4.18). On investigating it was found that there was some overlap in the (unfocused) beams at the spectrometer causing an interference effect. Blocking the unwanted beam removed the fringing, allowing the spectrum to be measured correctly, but it was realised that the changes in the fringes as the delay line was moved could be used to determine the displacement. Note that a true interferometric measurement is not possible using this method because there is a large spectral bandwidth within the pulses. Even so, it provides a useful cross-check of the look-up table generated using the interferometer measurements by checking whether the delays produced corresponded to those requested by the software.

A colleague carried out the cross-check<sup>4</sup> by locating the PZT at zero

---

<sup>4</sup>EOS190 page 102

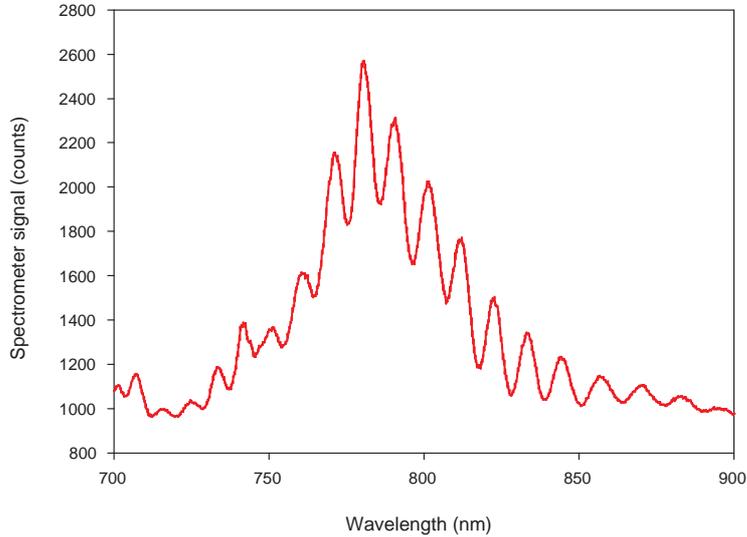


Figure 4.18: Spectrometer measurement of the Femtolaser Ti:S pulses exiting the FROG after the second beamsplitter. There are two sets of pulses that overlap slightly at the spectrometer due to divergence. Here the delay line has been moved away from zero, producing interference fringes in the combined spectral output.

delay and then changing the path length using the manual translation stage by a fixed number of fringe movements (from constructive to destructive interference to constructive again). A wavelength of 800 nm was used to determine the change from the zero delay and compared to the corresponding change in the PZT required to return the path lengths to zero difference. It was established that the two propagation times agreed to within 5 fs ( $1.5 \mu\text{m}$ ) of each other.

### 4.3.3 Crystal thickness

For correct operation of the FROG it is necessary that the SHG crystal is suitably phase-matched (c.f. Section 2.2.1) for the pulses being measured otherwise the spectrogram will lack the full spectral information. For the

BBO crystal the phase-matching angle was calculated to be 29.1 degrees using the SLNO software package. The crystal thicknesses as calculated according to the Sellmeier coefficients for a BBO crystal were 300  $\mu\text{m}$  for the 100 fs pulses and 100  $\mu\text{m}$  for the 10 fs pulses.

As a confirmation, the expected zero delay signal was calculated from the measured spectrum of the laser according to

$$S_{\text{zerodelay}}(\lambda) = \sqrt{S_{\text{laser}}(\lambda)} \quad (4.21)$$

as an approximation. A comparison of this to a measurement of the zero delay taken using the 300  $\mu\text{m}$  crystal is shown in Figure 4.19. The  $\Delta\nu$  of the measured and the predicted traces agree to 0.1 nm, follow the same form and match in wavelength. This suggests that the SHG crystal is correctly phase-matched for the pulses from the laser.

Additionally the zero delay signal for the Mira Ti:S laser (100 fs, 850 nm) was measured using both the 300  $\mu\text{m}$  and the 100  $\mu\text{m}$  crystals. Although the 100  $\mu\text{m}$  crystal is designed for a different central wavelength (800 nm) the difference is slight. If there are any phasematching bandwidth issues due to crystal thickness such a measurement should indicate this. The results are shown in Figure 4.20. It was found that it was necessary to shift the zero delay spectra by 0.8 nm relative to each other to produce an overlap. Subsequent investigation (Section 4.3.4) determined that this was due to differences in coupling of the signal into the optical fibre used with the spectrometer. Offset aside, the clear agreement in shape and width between the two measurements (normalised to account for the lower signal level from the thinner crystal) again demonstrates that the SHG crystal is correctly phase matched to the pulses from the Mira laser.

#### 4.3.4 Spectrometer

The final part of the FROG system is the spectrometer used to measure the FROG signal. The original set-up used a grating-based monochromator with

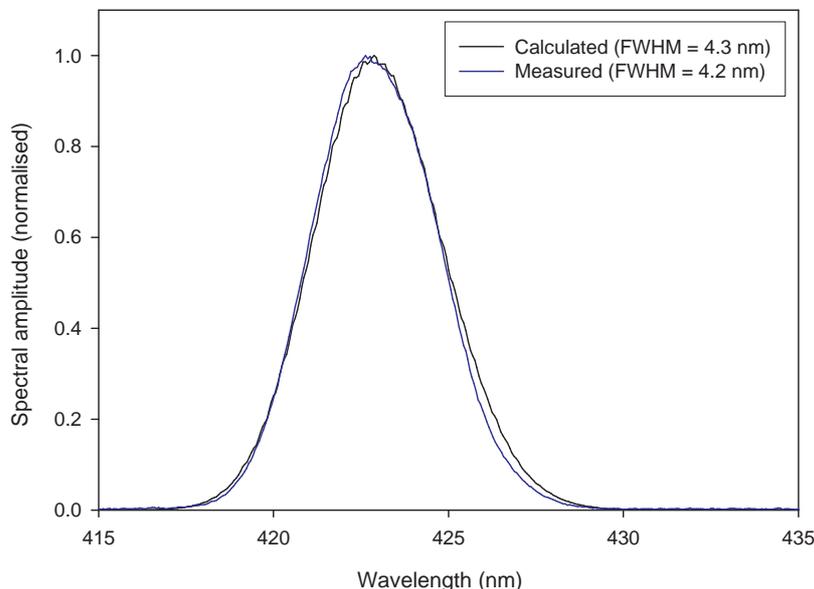


Figure 4.19: Theoretical calculation of the zero-delay signal from the Mira spectrum compared to the signal measured by the FROG system, indicating that the crystal used provides suitable phasematching bandwidth.

the final slit removed and replaced by a CCD array<sup>5</sup> to make the monochromator into a spectrometer. This array consisted of 2048 pixels and gave a total spectral span of roughly 140 nm. Calibration of the pixel/wavelength values was carried out using two He-Ne lasers at 633 nm and 543 nm to provide wavelength references. Different areas of the spectrum are accessed by rotating the monochromator gratings and the accuracy of the spectrometer for different rotations confirmed by checking the returned peak wavelengths for each He-Ne for different grating positions. Figure 4.21 shows an example of such a measurement.

The CCD/monochromator arrangement was used for a number of the early FROG measurements. LabVIEW was used for the control software and communication with the CCD array was achieved using a DLL and the supplied PCI DAQ (data acquisition) card with associated National Instru-

<sup>5</sup>CVI AD150 — not related to the AD100 actuator.

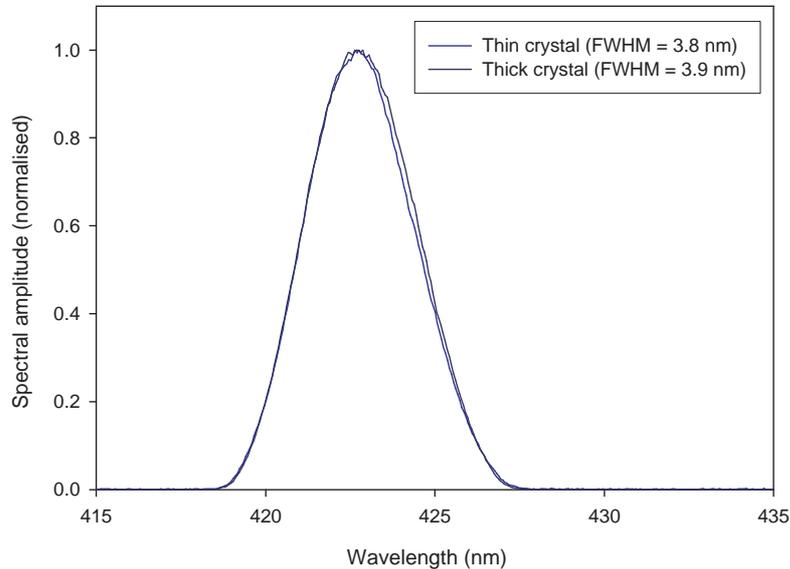


Figure 4.20: The zero-delay signal measured by the FROG system for the Mira using a  $300\ \mu\text{m}$  and a  $100\ \mu\text{m}$  SHG crystal. The  $300\ \mu\text{m}$  crystal was calculated as having suitable phasematching bandwidth for the Mira pulses. This is supported by the absence of any additional bandwidth when using the thinner crystal.

ment drivers. However, during a LabVIEW version upgrade communication with the CCD array was lost and after investigation this was found to be due to the upgraded drivers no longer supporting the model of card. The driver installation turned out to be non-reversible, and the resulting interface issues led to the purchase of a replacement spectrometer. This in fact aided the portability of the FROG, as the existing monochromator had to be mounted outside the breadboard of the FROG system.

The new spectrometer was an Ocean Optics HR4000 and is shown installed in Figure 4.22. The spectrometer interfaced with the PC via USB and contained 4096 pixels. The input to the spectrometer was designed to be by optical fibre, and one was purchased with the spectrometer for this purpose. The FROG signal was aligned via the fibre but it was found that, especially for the broader spectrum of the sub-12 fs laser, the signal could

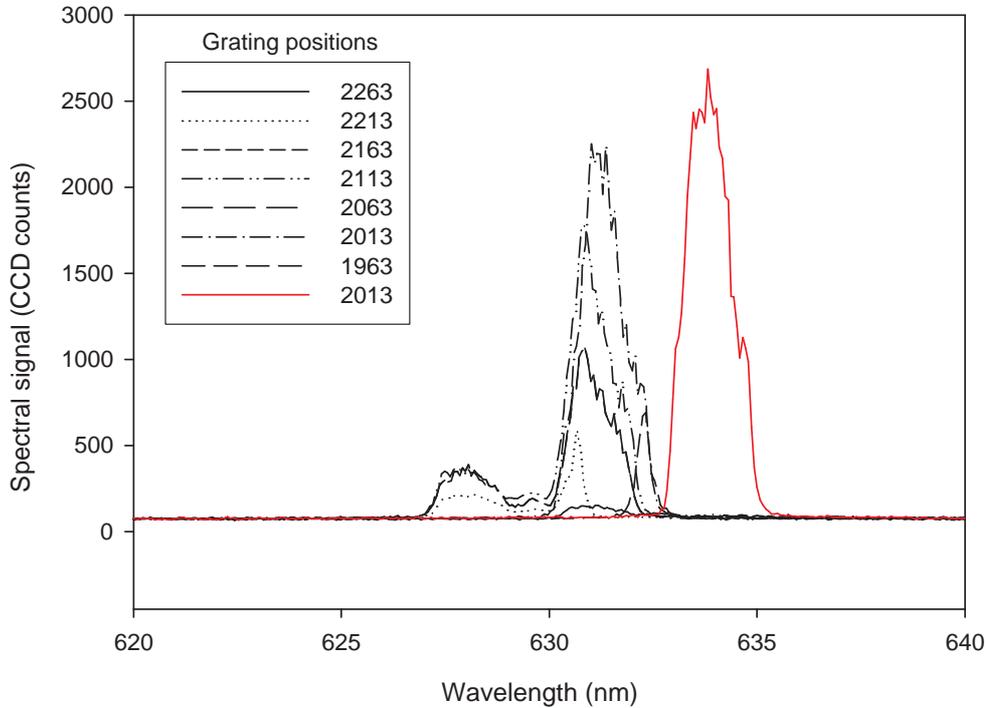


Figure 4.21: Checking the wavelength of the AD150 and monochromator system using a He-Ne. The traces represent different grating rotations (indicated in the legend). It is seen that the wavelength measurement is consistent for different rotations but that there is an error in the absolute wavelength. A screw on the monochromator allows this to be adjusted, and the resulting correction can be seen in the figure in red.

be made to vary according to the fibre alignment. Mindful of the errors this could introduce, free field alignment was tried instead and found to be far less error-prone, only the amplitude of the signal varying with alignment. Furthermore the free-field arrangement was found to be more sensitive to low signal levels, leading to improved measurements when using thinner SHG crystals.

Pellicle beamsplitters are used to minimise dispersion in the pulses (c.f. Section 4.3.1) and a gimbal-type mount on the fixed arm enables adjustment of the beam overlap at the crystal to maximise the signal while avoiding de-

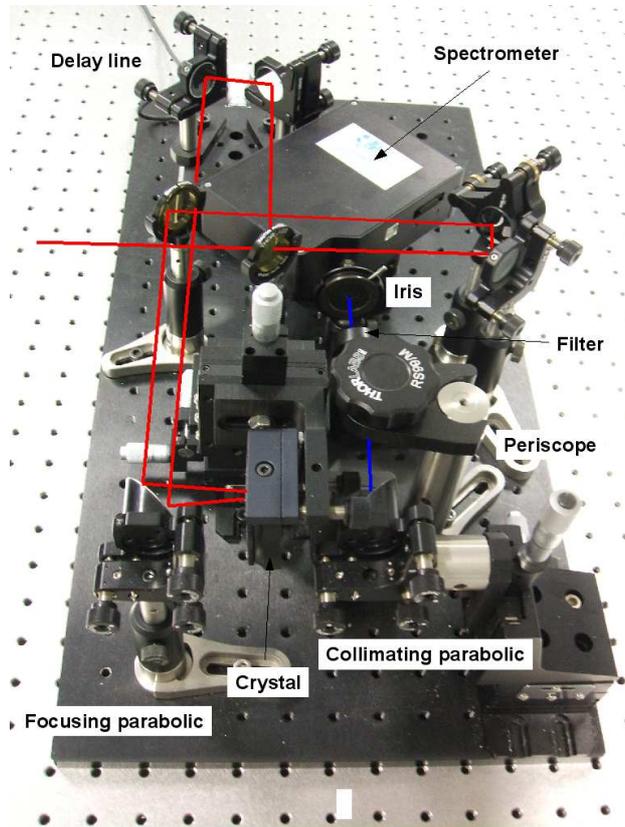


Figure 4.22: The portable FROG system mounted on a 300 mm  $\times$  600 mm breadboard.

lay line changes that could effect that signal level. The crystal itself is held in a rotational mount to allow optimal alignment of the crystal axes with the polarisation of the beams. Additionally the mount allows adjustment of the normal of the crystal relative to the beam paths to optimise the z axis. xyz-translation stages enable fine-positioning of the crystal and each parabolic mirror to ensure best focus of the pulses and collimation of the SHG generated. A periscope lowers the SHG height — this is purely to enable the spectrometer to be incorporated onto the breadboard beneath the paths of the input pulses. A filter prevents any of the input pulses reaching the spectrometer and an iris blocks the base SHG from the input pulses.

Note that it was found that if the input beam diameter to the FROG is

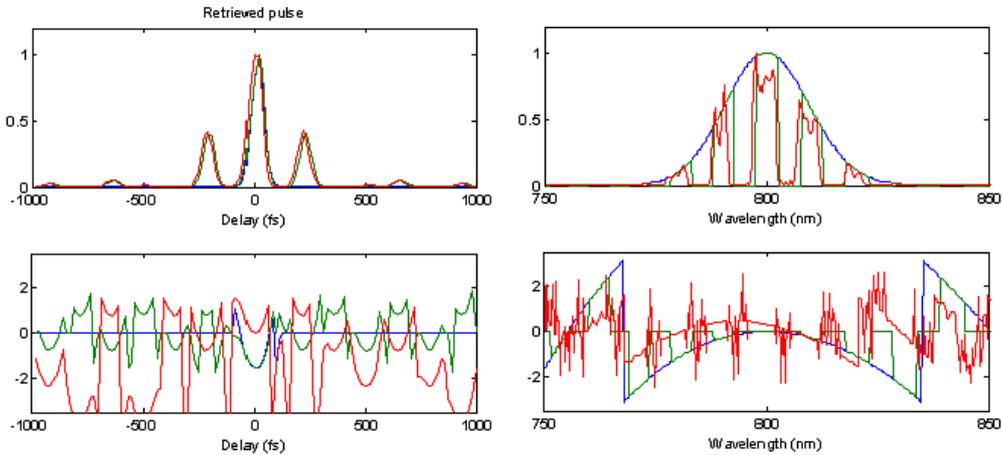


Figure 4.23: Retrieval of pulses from modelled spectrograms. The intensity and phase in the temporal domain are on the left and in the spectral domain on the right. Blue indicates the unshaped pulse, red the shaped pulse and green the retrieved pulse.

too large there will be some overlap of the SHG from the individual pulses at the spectrometer as well as the FROG signal, despite the use of an iris before the spectrometer, leading to an interferometric pattern in the spectrogram. This effect will give rise to interferometric fringes in the spectrogram itself.

### 4.3.5 Retrieval Algorithm

As described in Section 4.2, pulse amplitude and phase information is retrieved from a spectrogram by means of a suitable algorithm. Such an algorithm for SHG-FROG was supplied<sup>6</sup> in MATLAB code for use with this system. The bulk of the (modelled) testing of the algorithm and the processing of the results in MATLAB was carried out by an NPL colleague<sup>7</sup>. Initial testing of the algorithm using modelled pulses (Figure 4.23) demonstrated its suitability for pulse retrieval where more complex pulses are involved, as anticipated for the pulse shaper (c.f. Section 2.4.2).

<sup>6</sup>D. T. Reid, Heriot-Watt

<sup>7</sup>S. N. Lea

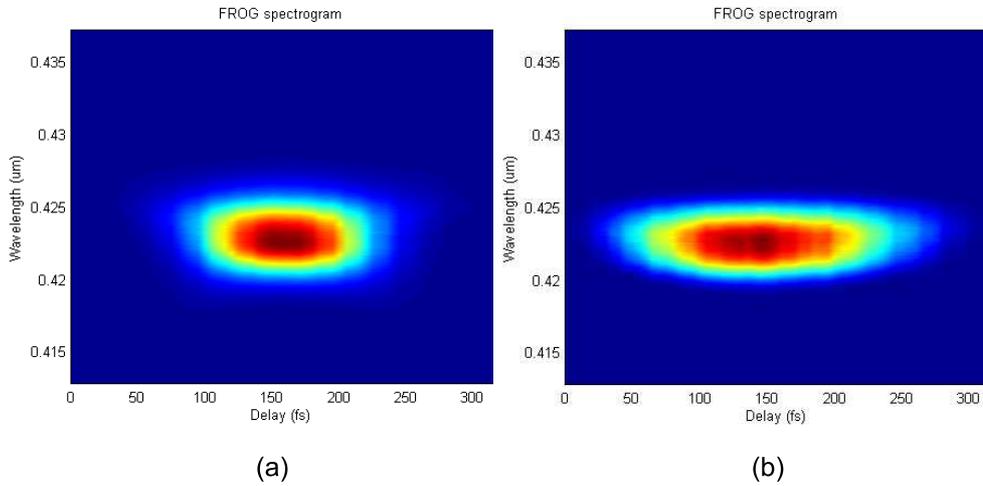


Figure 4.24: (a) Spectrogram of the output from a Ti:S laser. (b) Spectrogram of the same output after passing through a 25 mm block of SF57 glass. The increased pulse width from the resulting dispersion is obvious.

## 4.4 Test FROG measurements

### 4.4.1 Chirped pulses

For an initial test of the FROG system a measurement was made using a deliberately chirped pulse. Firstly a spectrogram was taken of the Mira output as passed directly to the FROG, to provide a reference. 64 delay steps were used, each of 5 fs. The resulting spectrogram is shown in Figure 4.24 (a). Then, in order to produce significant chirp in the pulses, a 25 mm thick block of SF57 glass was placed in the beam path. This glass was mounted on a rotational stage to present a normal face to the beam path in order to prevent deviation. (Even so some slight change to the Au and Al parabolic mirrors was required, indicating that the two faces of the glass were not parallel.) The spectrogram taken with the glass block in place is shown in Figure 4.24 (b).

It is immediately apparent on inspection of the spectrogram that there has been temporal broadening of the pulse, as seen from the increased extent along the delay axis. Processing of the data using the algorithm produces a retrieved pulse for each as shown in Figure 4.25. The broadening of the pulse

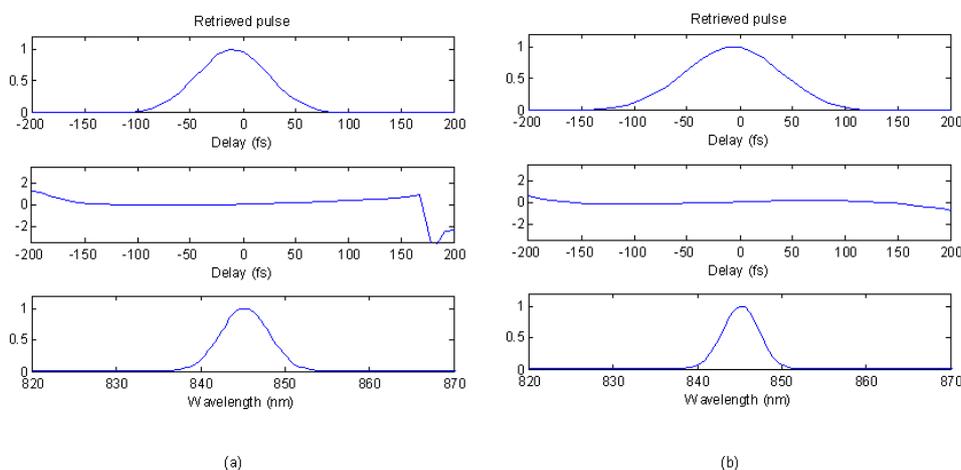


Figure 4.25: Retrieved pulses from the spectrograms (a) and (b) in Figure 4.24. The pulse broadening is confirmed. There also appears to be some truncation of the spectrum.

is confirmed, as is the change in phase due to frequency chirp (less apparent on the diagram). There also appears to be some narrowing of the spectrum.

#### 4.4.2 Sub-12 fs pulses

A further requirement of the FROG system was the ability to measure pulses as short as 10 fs. For this purpose the output from a Femtolaser Scientific Ti:S source was used. This laser has a spectral range from 700–900 nm (Figure 4.26) and a specified pulse  $\Delta\tau$  of  $< 12$  fs when using an external cavity dispersion compensation (ECDC) component consisting of a pair of chirped mirrors. Note that this should be considered a best-case rather than a typical pulse duration.

The output from the laser was directed into the FROG using three external Ag mirrors and the ECDC with two reflections per chirped mirror. The resulting spectrogram is shown in Figure 4.27 and the retrieved pulses in Figure 4.28<sup>8</sup>.  $\Delta\tau$  of the retrieved pulse is 20 fs, in the region of what would be expected from the laser. Comparison of the retrieved spectra with a mea-

<sup>8</sup>Retrieval by S. N. Lea.

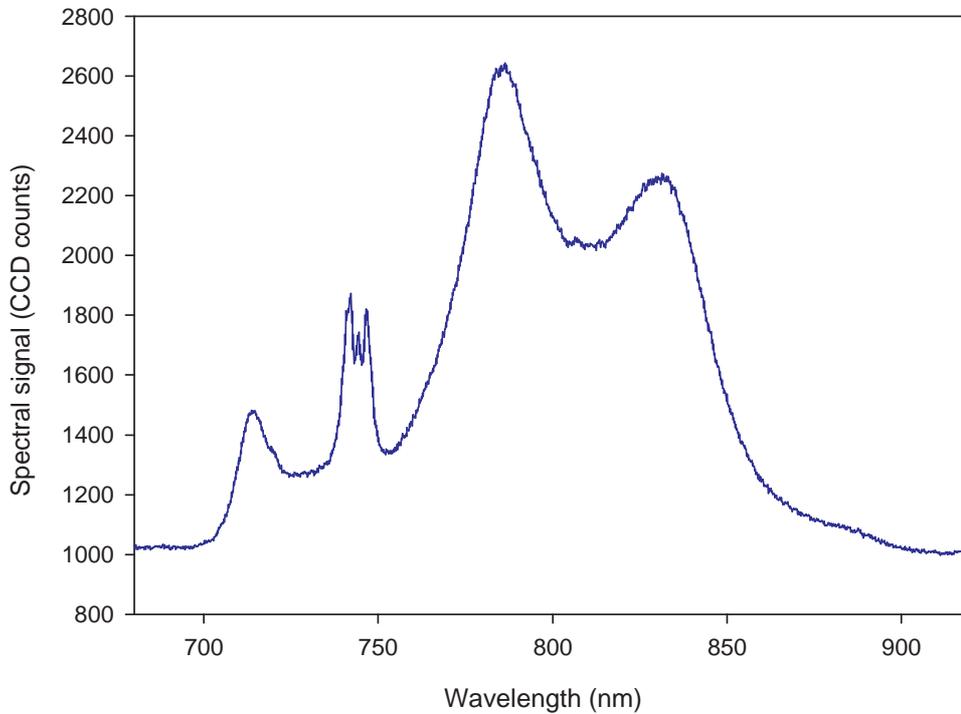


Figure 4.26: Example spectrum of the Femtolaser. The feature at 740 nm is not believed to be part of the true spectral output of the laser (since it is present even when no modelocking is occurring) but some kind of reflective error within the spectrometer.

sured spectrum taken at the time shows broad similarities but is not a close enough retrieval to give full confidence in the result. The lack of matching detail in the spectra could be due to loss of information outside the range of the PZT (see Section 4.4.3). It was also noted that there was an error in the symmetry of the spectrogram (recalling that SHG-FROG is symmetric in time). Repeated measurement of spectrograms for different adjustments of the crystal, and verification that the asymmetry was identical and reversed for a negative-sweep spectrogram, suggested this was due to some alignment issue at the crystal as opposed to any error at the detector. Even so, the indications are good for the use of this system for pulse measurement at

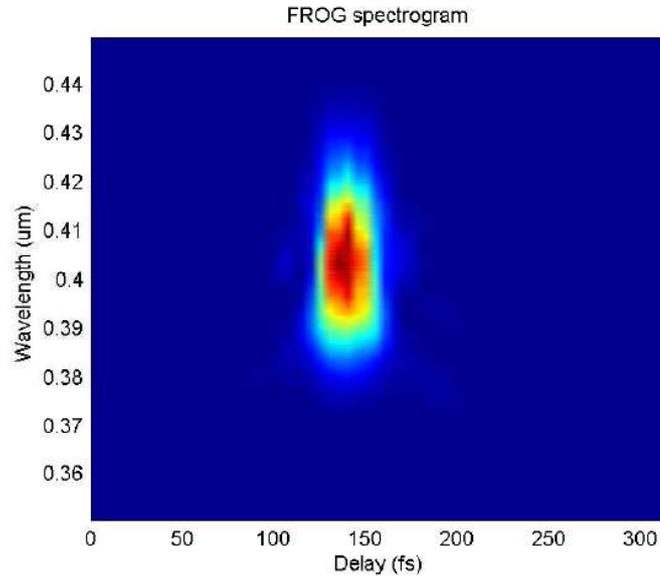


Figure 4.27: Spectrogram taken of the Femtolaser Ti:S output. Some asymmetry is noted, most likely due to an error in the alignment.

the order of magnitude of 10 fs pulses once the sources of errors have been identified and eliminated.

### 4.4.3 Shaped pulses

To test how the FROG worked with structured pulses, the output from the laser was passed through a  $4f$  pulse shaper (c.f. Section 2.4.2) prior to measurement by the FROG. Various semi-arbitrary masks (based around ‘picket fence’ arrangements where pixels are raised to voltages above a common level at various intervals) were applied to the LCM and measurements taken of the resulting spectra of the shaped pulses and the spectrograms that they produced in the FROG system. These measurements are shown in Figures 4.29 and 4.30. The introduction of structure into the FROG signal is readily apparent.

Processing of the measured spectrograms (Figure 4.31) initially produced a disappointing result. Despite the clear presence of structure in the mea-

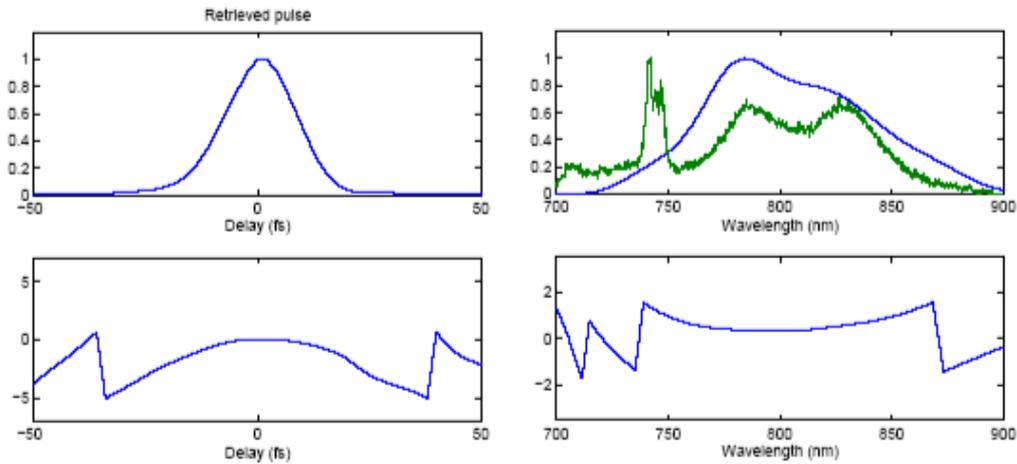


Figure 4.28: Retrieved temporal and spectral intensity from the spectrogram in Figure 4.27. The approximate FWHM of the pulse is 20 fs and there appears to be little chirp. The retrieved spectrum is compared to the measured spectrum (in green) and is broadly similar. It is believed that there is a lack of spectral resolution due to the limited length of the delay line. The feature at 740 nm in the measured spectrum is an artifact.

sured spectrograms no structure was seen in the spectrograms formed from the retrieved pulses, nor was there any structure in the spectra of the retrieved pulses. The clear cause is that the algorithm is not retrieving the correct pulses, yet testing with complex pulses produced satisfactory results (Section 4.3.5) so the reasons behind this were not immediately clear. Further examination of the spectrogram measurements within the FROG system were carried out and it was observed that beyond the perceived limit of the spectrogram (where the signal had reduced to background levels), and beyond the limit of the PZT, further signal would emerge almost as an ‘echo’ of the main spectrogram (Figure 4.32). (Inspection beyond the limit of the PZT was carried out by moving the manual translation stage.) A return to the study of the behaviour of the algorithm with modelled pulses found that, for a modelled pulse with a spectral modulation applied similar to that produced by the pulse shaper, the spectrogram contains extended ‘wings’ as seen in the physical system. Retrieval of the pulses for spectrograms where

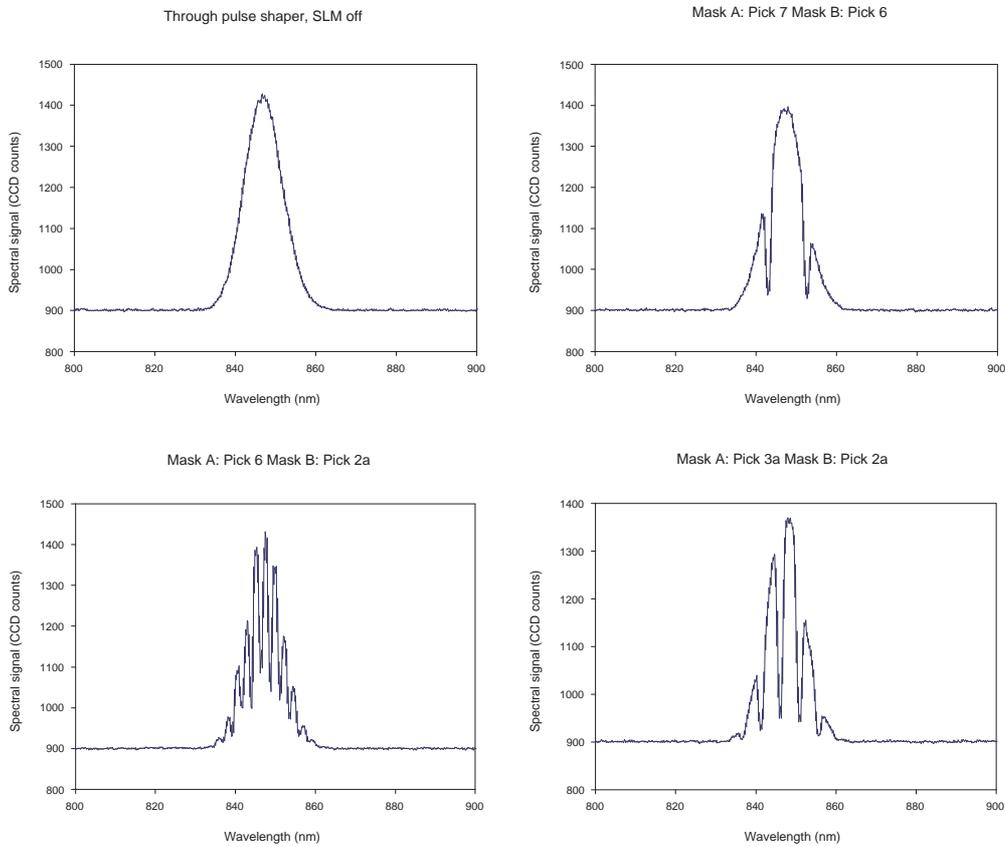


Figure 4.29: Spectra of the Ti:S laser pulses after travelling through the pulse shaper with various masks applied. The top left figure shows the spectrum with no shaping.

this information in the wings has been lost (in the case of the model, through a filtering function) demonstrated the loss of spectral resolution in the retrieved pulses. Figure 4.33 shows how the retrieved spectrum approaches the unfiltered Gaussian pulse spectrum as the information is lost. The conclusion, therefore, was that the failure to retrieve spectrally correct pulses from the spectrograms taken for shaped pulses was due to the loss of information caused by the restricted delay range of the PZT.

Since there was no scope for increasing the delay using the current PZT (the measured spectrograms were already being taken at the limits of the PZT) it was necessary to build up a composite spectrogram by measuring

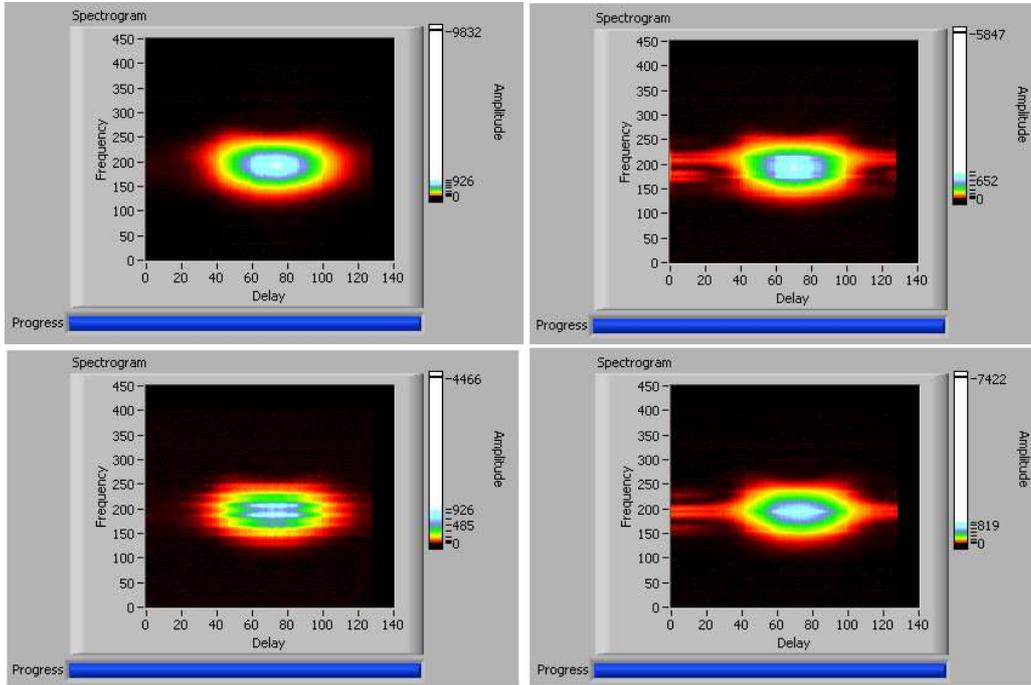


Figure 4.30: Spectrograms of the Ti:S laser pulses after travelling through the pulse shaper with various masks applied, corresponding to the spectra in Figure 4.29.

the spectrogram for various positions of the manual translation stage. Ten spectrograms were taken, each a  $50\ \mu\text{m}$  step of the manual delay stage apart (Figure 4.34). (Note that a  $50\ \mu\text{m}$  step corresponds to a total  $100\ \mu\text{m}$  change in the delay path length.) These individual spectrograms were then combined to produce a single composite spectrogram as seen in Figure 4.35. A cross-check of the delay of the PZT and the micrometer could also be carried out by comparing the delay marginals of several of the spectrograms. (A marginal is the summation of the spectrogram elements along one axis, in this case the delay (time) axis.) The result of this cross-check is shown in Figure 4.36, indicating agreement between the two.<sup>9</sup>

Processing of the composite spectrogram with the algorithm led to a retrieved spectrum that bore more resemblance to the measured spectrum

<sup>9</sup>Composition of spectrogram and marginal check by S. N. Lea.

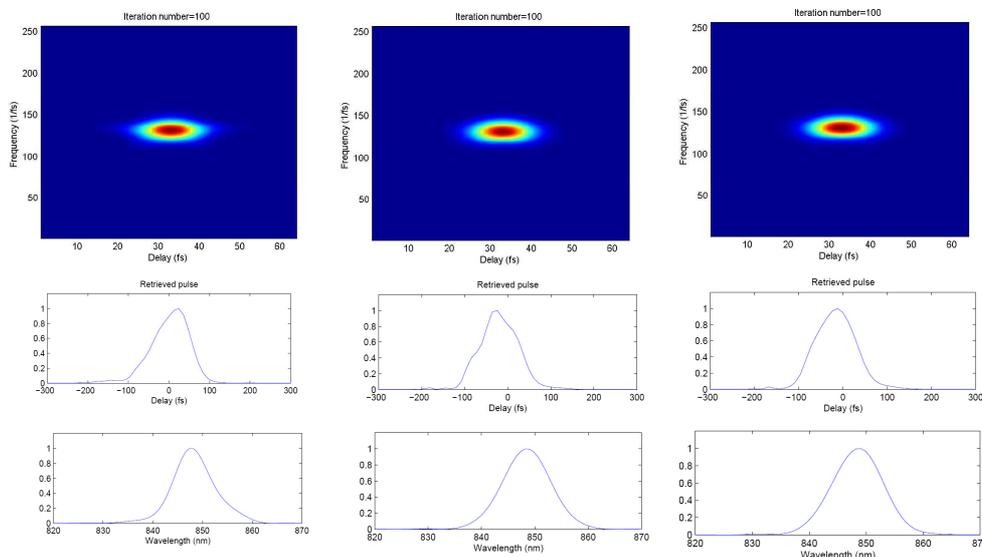


Figure 4.31: Retrieved spectrograms and pulse intensities for the shaped spectrograms in Figure 4.30. Despite the clear differences in the measured data the retrieved pulses are all simple in shape and the spectrograms formed from the retrieved pulses are all remarkably identical and lacking the structure seen in the measured ones. This was found to be due to missing information in the wings of the spectrograms.

(Figure 4.37). The retrieval indicates a train of pulses has been produced by the pulse shaper, decreasing in power, each with  $\Delta\tau$  comparable to the original pulse and spaced approximately 500 fs apart. (Note that there is a time ambiguity in SHG FROG so it could be that the lower intensity pulses lead the main pulse. Intuition would suggest that the displayed retrieval is the correct one, but then again intuition has been known to be wrong!)

The retrieved spectrum does not correspond ideally with the measured spectrum so this cannot be considered a conclusive determination of the pulse shape, but it does suggest a capability of the algorithm for dealing with real, complex pulses and moreover demonstrates the need for a large delay range within the FROG, beyond the existing range of the PZT for measurement of shaped pulses. Even so, given the current need for a manual adjustment to acquire the spectrogram for a shaped pulse, and the likely introduction of errors during the composition of the full spectrogram, it was decided to limit

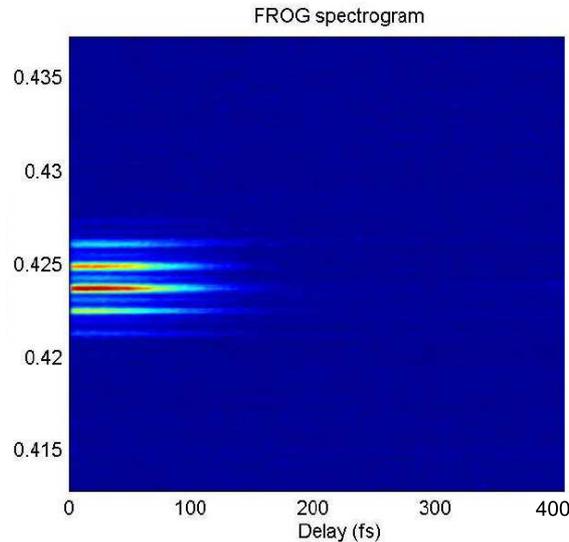


Figure 4.32: Spectrogram measured beyond the range of the PZT relative to the zero-delay point. The extra signal in this region is clearly apparent.

the use of the FROG to unshaped pulses until either a larger-range actuator was used or the application of the manual delay line was fully understood.

#### 4.4.4 Uncertainties in FROG

Retrieved pulses produced by FROG measurements are often quoted to significant levels of accuracy but the uncertainties associated are rarely provided or even discussed. Uncertainties are crucial when measurements are intended for more than basic indication and it is intended that a full and complete uncertainty budget is produced to generate uncertainties for measurements taken using the NPL FROG system.

This work is currently still in its early stages, but the intention is to assign uncertainties to the spectrogram measurement first and then, through simulation, modelling and/or theoretical calculation, determine how those uncertainties contribute to the uncertainties of the retrieved temporal intensity (and phase, if necessary) of the pulse. The method will be to determine how each aspect of the FROG system affects the three parameters of the

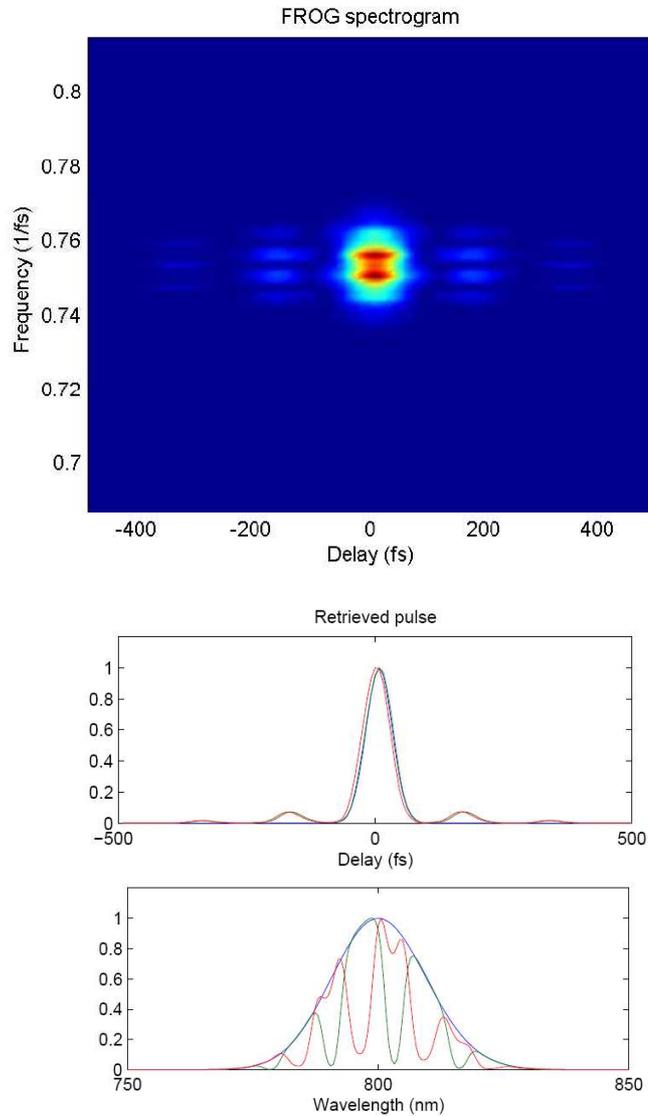


Figure 4.33: Retrieval of a pulse from a spectrogram (top) formed from a modelled shaped pulse. The blue curves in the plots of the retrieved pulse show the unshaped pulse and the green curve the retrieved pulse with all data present in the spectrogram. When filtering is applied to the spectrogram before retrieval to simulate the loss of data at the limits of the delay (red curve) it is seen that the retrieved pulse starts to resemble the unshaped pulse.

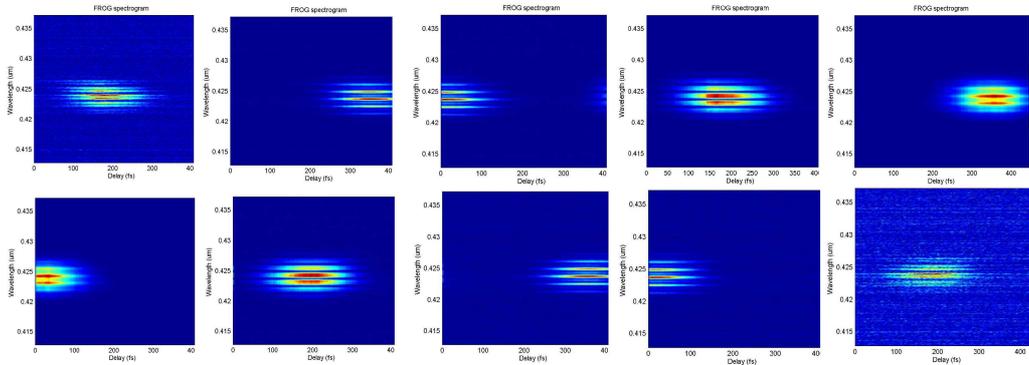


Figure 4.34: Sequence of spectrograms taken for differing positions of the micrometer on the delay line. The order is from top left across to bottom right, each spectrogram taken at a delay  $100 \mu\text{m}$  greater than the one before. Note that there is overlap between the spectrograms since the total delay used for each spectrogram is  $> 100 \mu\text{m}$ ). It can be seen that there are three ‘echoes’ visible on each side of the central image.

spectrogram points: The delay, the wavelength and the intensity. There may also be additional contributions due to aspects such as the retrieval process itself. A preliminary list is given below.

- **Delay:** The clearest contribution here will be the uncertainty in the delay setting at each point in the spectrogram. This will be a direct consequence of the error in the PZT, and also there will be a quantisation error in the setting. Repeatability may also be an issue, separate from the absolute accuracy of the displacement.
- **Wavelength:** The main contribution will be the wavelength calibration of the spectrometer used. There may also be phasematching errors due to any issues with the crystal or alignment, though these could be considered to be intensity errors (see below).
- **Intensity:** Background noise, and the signal to noise ratio, will be a factor in the intensity. Fluctuations in the pulse power may also contribute a random error. (Fluctuations in the pulse shape will require careful consideration.) The wavelength response of the spectrometer

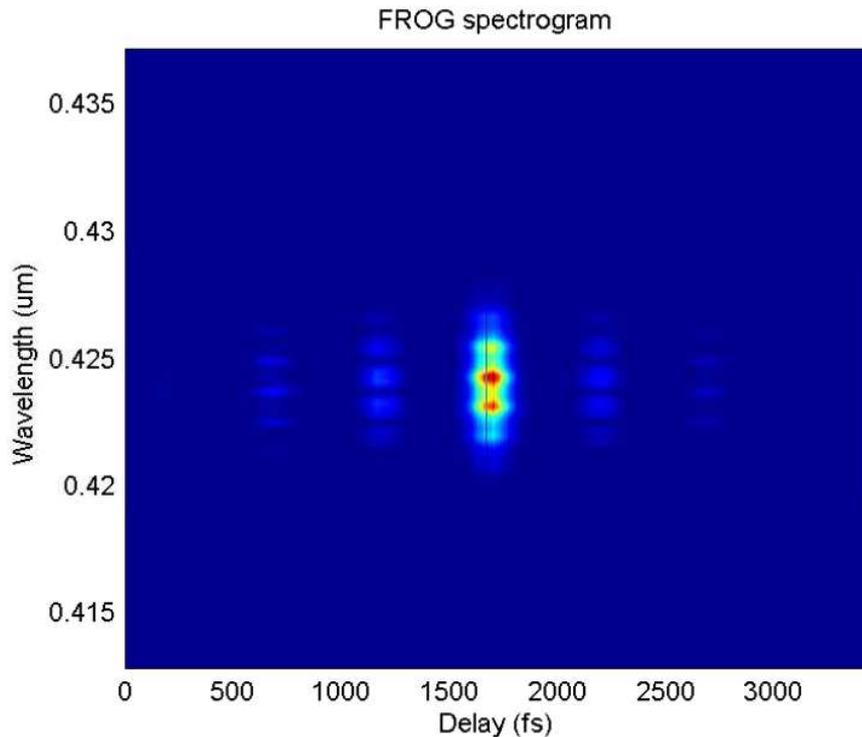


Figure 4.35: Composite spectrogram constructed from the multiple measurements of Figure 4.34

will also contribute as will any alignment errors (changing overlap as the delay line is translated, for instance.) Additional SHG from the individual arms may leak through to the spectrometer despite the iris and this will also contribute to the error, if only as an uncertainty in the subtracted background signal. Finally there will be an uncertainty associated with the digitisation of the spectrometer, although this is likely to be small.

- **Retrieval:** The effect of the retrieval process will be difficult to quantify. The nature of it should be such that the same spectrogram results in the same pulse, meaning only systematic errors remain, but this will need to be confirmed. Likely sources of error may lie in the interpolation (such as when converting from the wavelength domain to the

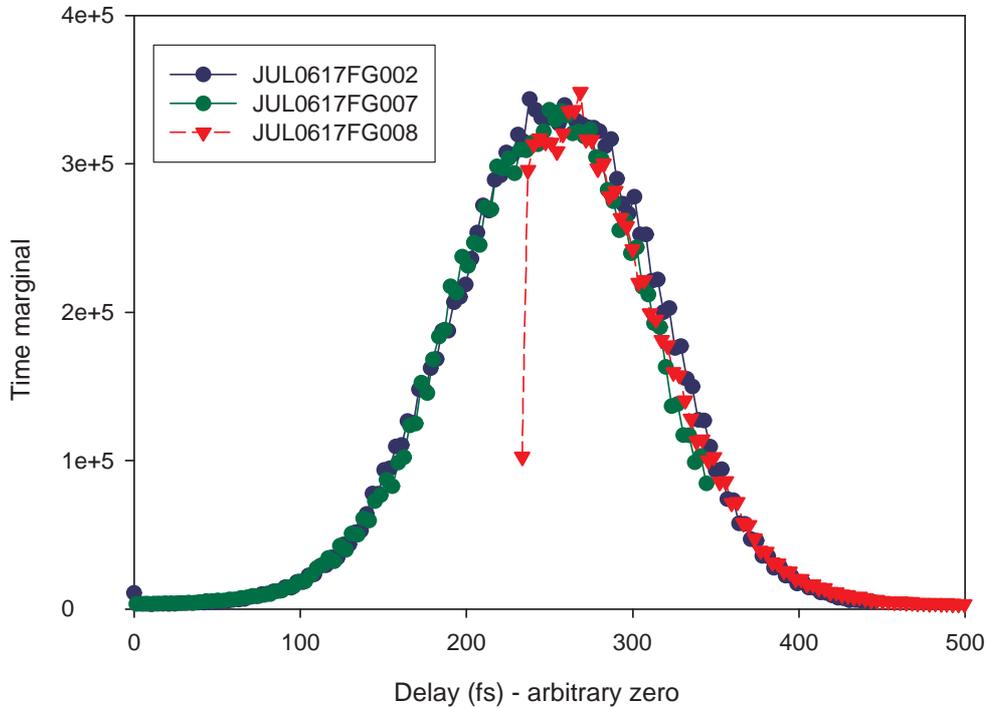


Figure 4.36: Plot of the marginals (the sums of the spectrograms along their delay axes) for spectrograms taken for differing positions of the micrometer on the delay line. The overlap confirms that the displacements read from the micrometer are correct. (The erroneous point was found to be due to a poor reading at the start of the given spectrogram.)

frequency domain) or the FFT conversions within the algorithm.

By first establishing how the three basic uncertainties, in delay, wavelength and intensity, contribute to the overall uncertainty of the measurement (plus any addition systematics in such aspects as the retrieval) and then in turn determining how the individual aspects of the FROG contribute to those three uncertainties, it should be possible to produce a complete uncertainty budget for pulse measurements made using the portable FROG system.

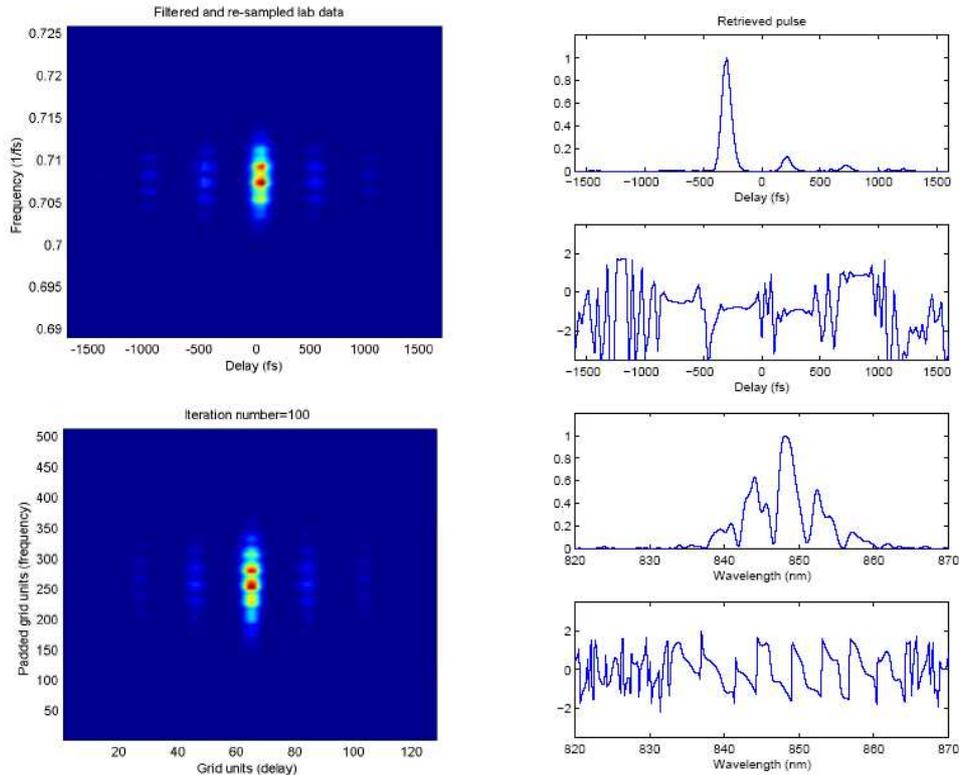


Figure 4.37: The retrieval result from the composite spectrogram of Figure 4.35 (top left). The spectrogram at bottom left is the one formed from the retrieved pulse shape. The plots on the right show the temporal and spectral intensities and phases.

## 4.5 Conclusions

In summary, it can be concluded that a compact and portable FROG system has been designed and constructed successfully for the purpose of the measurement of shaped and unshaped optical pulses in the 700–900 nm regime for pulse durations as short as 10 fs. Dispersion within the system has been kept to a minimum by the use of silver-coated mirrors, pellicle beamsplitters and an Au-reflecting parabolic mirror. The SHG crystal can be easily exchanged as appropriate for the expected bandwidth of the pulses to balance phasematching requirements with signal levels. Software written in LabVIEW enables real-time monitoring of the FROG signal for alignment

purposes and the automatic acquisition of spectrograms. A MATLAB algorithm is used to perform the retrieval of the amplitude and phase information (both temporally and spectrally) from the spectrogram. This system has been demonstrated to produce satisfactory results for unshaped pulses with 25 fs pulse durations but measurements of shaped pulses have determined a lack of resolution due to the limited range of the PZT. Extension of this range with a suitable actuator will enable more correct measurement of spectrograms from shaped pulses and allow the resolution limits of the system to be tested more fully. It is anticipated that this, along with an improved understanding of the assignation of uncertainties to the retrieved result, will produce a valuable tool for ultrafast studies typical of those carried out at NPL.

# Chapter 5

## Conclusions and future work

The application of control of ultrafast pulses in pump-probe arrangements has been investigated for two different metrological applications. The first application concerned the use of optical pulse shaping with a genetic algorithm to control the spectra of terahertz radiation as a means of developing wavelength standards. The second application investigated the use of quantum interference control for ultrafast electrical pulse generation as a means of providing rise-time calibration standards. Furthermore, a compact frequency-resolved optical gating system for characterisation of ultrafast pulses in pump-probe systems has been developed. The conclusions of these investigations are presented here.

### 5.1 QUIC for electrical pulses

#### 5.1.1 Generation and EOS of QUIC signal

The experimental scheme for generating a  $< 100$  fs electrical pulse by means of quantum interference control (QUIC) in a semiconductor device has been designed and constructed. This is significantly shorter in duration than electrical pulses generated by traditional bias-voltage methods (650–800 fs). The QUIC process exploited is one-photon and two-photon absorption in LT-GaAs using a built-for-purpose femtosecond optical parametric oscillator (OPO) for the fundamental frequency and a SHG crystal for the harmonic.

Measurements of the behaviour of charge build-up across the device with changes in polarisation, power and relative delay of the harmonic all agreed with the predicted outcomes for a QUIC-induced electrical signal. Electro-optic sampling (EOS) was used to make a temporal measurement of the QUIC-injected electrical pulses as would be required for metrological applications (e.g. determination of the impulse response of an oscilloscope). A trace signal was found with the expected characteristics but the sensitivity of the EOS was limited due to lack of probe power (the same laser was required for pumping the OPO) and an unexplained random change in the pump-probe relative timings. Further development was hampered by failures of the pump source for the Ti:S laser and the acousto-optic modulators (AOMs). Therefore the creation of ultrafast electrical pulses through QUIC has been demonstrated but they are not yet suitable for metrological applications.

### 5.1.2 Further work on QUIC

The main requirements if these pulses are to be used for metrology are improved sensitivity of the EOS measurement and increased signal level of the QUIC-induced electrical pulse. In the first instance, the cause of the shifts in timing should be investigated and eliminated to enable better averaging. Furthermore, any methods of maximising the power of the probe pulses (such as eliminating all unnecessary optics, retuning the laser, and so on) without depleting the power to the OPO should be investigated. To increase the QUIC signal level the use of a 750 nm AOM instead of a 1500 nm AOM should be investigated further because there is evidence that the loss of 750 nm power will have reduced effect compared to the current 1500 nm power loss. There may, however, be new problems due to the possible detection of any spurious signals related to one-photon absorption.

The OPO itself should be developed to improve its stability (which will aid the EOS sensitivity) and power output (which will aid the QUIC signal level). Improvements such as the use of lockable mounts have already been incorporated, and the introduction of prisms should also be revisited.

Should the EOS measurement of the QUIC signal improve to the stage

that the signal becomes of practical use then various voltage signals can be applied to the PZT (used to generate the phase delay between the fundamental and harmonic pulses) under the control of the genetic algorithm. This can then be used to investigate the ability to influence and tailor the sub-picosecond electrical pulse. The application of the pulse shaper should also be considered, although its inherent loss of power may prove a hindrance. It may, however, be used to produce phase variations within the 750 nm pulses that control the electrical QUIC pulse with a finer resolution than that achievable with the PZT.

## 5.2 Shaping of ultrafast pulses

### 5.2.1 Pulse shaper and genetic algorithm

A  $4f$  pulse shaper based around a dual-mask 640 pixel LCM has been designed and constructed to simultaneously control the spectral amplitude and phase of ultrafast laser pulses, therefore shaping the temporal amplitude and phase. The ability to alter a laser spectrum has been confirmed and a characterisation measurement demonstrated the conversion of a single pulse into multiple pulses spaced 500 fs apart. The inclusion of a genetic algorithm allows optimisation of ultrafast pulse-based systems based only on knowledge of the desired outcome, removing the need for any knowledge of the required pulse shape. This has been demonstrated experimentally with a pump-probe THz system where shaping of the pump pulses has altered the spectra of the THz radiation generated to emulate physical THz filters for wavelength standards. By carefully defining the evaluations of the algorithm outputs different LCM settings were produced that suppressed either the 0.7 THz, the 0.75 THz or the 1.5 THz frequencies of the generated THz pulses. The novelty of, and interest in, this result were demonstrated when the recent conference poster submission became an invited keynote presentation<sup>1</sup>.

The pulse shaper can be used with pulses ranging from 10 fs to a few hundred femtoseconds and optimised for a given bandwidth by careful selection

---

<sup>1</sup>IRMMW-THz conference, Cardiff, UK, September 2007

of gratings. The genetic algorithm is flexible enough that it can be used to find a solution in any situation where the suitability of the pulse shape can be defined in terms of some experimental parameter. The only restriction in finding a solution is one of time. Modelling of the genetic algorithm has been carried out and it has been determined that to minimise the number of evaluations carried out (and hence the time taken) to reach a solution small ( $< 20$ ) populations should be used with moderate (10–30%) numbers of new individuals introduced each generation. Together the pulse shaper and genetic algorithm provide a powerful tool for ultrafast pulse characterisation studies.

### 5.2.2 Further work

Extension of the THz filter emulation should be decided upon after the level of interest has been gauged in response to the presentation. The main aspect against its continuation is the long evolution times (60+ hour periods for reasonable results). However, given that the THz spectral manipulation has demonstrated that the pulse shaper and genetic algorithm work together successfully the next intention must be to apply this system to other applications where the evaluation can be made far more quickly.

One such proposal is to reduce the  $\Delta\tau$  of the electrical pulse produced by a photoconductive switch under voltage bias, as is currently used for typical oscilloscope calibrations at NPL. Normally these are measured using the EOS system but to reduce evaluation time a high-bandwidth (70+ GHz) oscilloscope should be used for the base measurement of the waveform (with minimal averaging). The oscilloscope firmware can return the  $\Delta\tau$  of the pulse waveform measured, removing the need for a time-consuming transferral of waveform data across a GPIB connection. Although the oscilloscope will return an erroneous  $\Delta\tau$  due to its insufficient response it should still detect noticeable reductions. Once an adequate solution has been reached by the algorithm a more accurate measurement can be made using EOS.

The pulse shaper should be rendered compact and portable by reconstructing it on a dedicated optical breadboard. Parabolic mirrors in place of

lenses will fold the beam paths and be more robust compared to the current lens mounts. Alignment of the pulse shaper should be such that alternative gratings can be interchanged as easily as possible according to the laser bandwidth.

Maintenance and further investigations into the genetic algorithm are already planned, and there will be further dissemination of the results throughout NPL to see if there are other areas that can benefit from an evolutionary algorithm. The emphasis will be on improving the speed of the algorithm in reaching a solution through methods such as variable rates of mutation or breeding. Refinements of the algorithm outputs, especially in relation to the pulse shaper, are also being examined (Appendix B).

## 5.3 Compact pulse measurement system

### 5.3.1 The NPL FROG system

A compact (600 mm  $\times$  300 mm) automated pulse characterisation system based on SHG frequency-resolved optical gating (FROG) has been designed and constructed for measurement of ultrafast optical pulses in the Ti:S range (700–900 nm). The emphasis has been on portability and the ability to measure down to 10 fs pulse durations. Careful and original design kept the whole system (excluding control electronics) within the optical breadboard. The potential to measure down to 10 fs pulse duration was achieved through the use of suitable mirrors and beamsplitters to prevent unwanted contributions to the dispersion of the pulses. The SHG crystal used to generate the FROG signal can be interchanged to balance signal level with bandwidth requirements.

The delay line has been calibrated using an interferometer and the calibration confirmed using a novel method based on interference between the unused portions of the replica beams. LabVIEW software has been written to perform automatic spectrogram measurements and the FROG has been shown to be capable of retrieving the amplitude and phase information of pulses down to 20 fs (the shortest available to the system). Retrieval of

shaped pulses has also been demonstrated but with reduced accuracy due to the limited range of travel of the PZT, necessitating manual alterations to the delay using the coarser micrometer and the combination of multiple spectrograms. The portability of the system has been demonstrated by movement between laboratories.

### 5.3.2 Further work

The next stage for the NPL FROG system is to develop a full uncertainty budget that can be used to provide full confidence in measurements. First, theory and modelling should be used to determine how the uncertainties in the spectrogram propagate through to the uncertainties in the retrieved pulse. Then the physical contributions to each parameter (delay, wavelength/frequency and intensity) of the spectrogram can be determined to produce an uncertainty budget spreadsheet for the FROG system as a whole. This will enable uncertainties to be ascribed to future ultrafast pulse measurements as well as identify those aspects of the system that require improvement.

The range of the delay line of the FROG must be extended by replacing the actuator whilst maintaining the resolution deemed necessary to enable the automatic measurement of complex shaped pulses. To allow measurement and retrieval to take place within the laboratory the current MATLAB algorithm will be translated into LabVIEW for full integration with the acquisition software. This will also enable future integration of the whole FROG system with the genetic algorithm and/or pulse shaper if required. Development of a LabVIEW version of the algorithm has begun but there is currently an issue that arises when the two algorithms differ greatly from each other within a single iteration, possibly due to their FFT implementations.

(Note that LabVIEW does allow the calling of MATLAB from within a LabVIEW program but this requires an installed licensed version on the same PC and prevents the compilation of the program into an executable file. There is also evidence that a LabVIEW implementation may run faster.)

Once the FROG is more suitable for shaped pulse measurements it will be

used to measure the ‘optimum’ pulses returned when the genetic algorithm is applied to other experiments. By knowing and understanding the pulse shape that has produced the desired experimental result it should be possible to gain new insights into the physics of the process being investigated.

## 5.4 Polarisation shaping project

A new three-year project has been successfully proposed and accepted for the current Quantum (now Pathfinder) research programme at NPL. Its purpose is to investigate QUIC for the manipulation of chemical processes along preferable reaction pathways. To achieve this the output from an amplified femtosecond laser system will be split into pump and probe pulses. The pump pulses will excite a molecule to an intermediate, evolving state that is studied at some later time by the photoionisation of the state by a probe pulse. By shaping the pump pulse and employing the genetic algorithm it will be possible to evolve toward the selection of preferred pathways through QUIC. Such control of molecular dynamics has a number of long-term potential benefits in areas such as chemical synthesis and drug design. FROG measurement of the shaped pulses that give the preferred reaction will then enable new understanding of the QUIC processes happening within the molecule. There will also be further development of the pulse shaper to incorporate polarisation pulse shaping.

It can be seen that in this project are various research aspects. These include the pulse shaper, the genetic algorithm, QUIC and FROG, and the whole approach is taken into yet another new, exciting direction in ultrafast pump-probe measurement. This project builds on what has been produced so far and offers the welcome potential to generate new knowledge and a much clearer understanding of the physics.

And that, after all, is what scientific research is about.

# Bibliography

- [1] P. M. W. French. The generation of ultrashort laser pulses. *Reports on Progress in Physics*, 58:169–262, 1995.
- [2] Hermann A. Haus. A theory of forced modelocking. *IEEE J. Quant. Elec.*, QE-11:323–330, 1975.
- [3] D. E. Spence, P. N. Kean, and W. Sibbett. 60-fsec pulse generation from a self-mode-locked Ti:sapphire laser. *Opt. Lett.*, 16:42–44, 1991.
- [4] P. F. Moulton. Spectroscopic and laser characteristics of Ti:Al<sub>2</sub>O<sub>3</sub>. *J. Opt. Soc. Am. B*, 3:125–132, 1986.
- [5] R. L. Fork, O. E. Martinez, and J. P. Gordon. Negative dispersion using pairs of prisms. *Opt. Lett.*, 9:150–152, 1984.
- [6] Q. Fu, G. Mak, and H. M. van Driel. High-power, 62-fs infrared optical parametric oscillator synchronously pumped by a 76-MHz Ti:sapphire laser. *Opt. Lett.*, 17:1006–1008, 1992.
- [7] David A. Humphreys, Matthew R. Harper, Andrew J. A. Smith, and Ian M. Smith. Vector calibration of optical reference receivers using a frequency-domain method. *IEEE Trans. Instrum. Meas*, 54:894–897, 2005.
- [8] D. T. Reid, M. Ebrahimzadeh, and W. Sibbett. Design criteria and comparison of femtosecond optical parametric oscillators based on KTiOPO<sub>4</sub> and RbTiOAsO<sub>4</sub>. *J. Opt. Soc. Am. B*, 12:2168–2179, 1995.

- [9] K. Finsterbusch, R. Urschel, and H. Zacharias. Fourier-transform-limited, high-power picosecond optical parametric oscillator based on periodically poled lithium niobate. *Appl. Phys. B*, 70:741–746, 2000.
- [10] M. R. Harper. *Operation of parametric oscillator*, dem/29/desk008 edition.
- [11] A. M. Weiner, J. P. Heritage, and E. M. Kirschner. High-resolution femtosecond pulse shaping. *J. Opt. Soc. Am. B*, 5:1563–1572, 1988.
- [12] A. M. Weiner, D. E. Leaird, J. S. Patel, and J. R. Wullert. Programmable femtosecond pulse shaping by use of a multielement liquid-crystal phase modulator. *Opt. Lett.*, 15:326–328, 1990.
- [13] Marc M. Wefers and Keith A. Nelson. Programmable phase and amplitude femtosecond pulse shaping. *Opt. Lett.*, 18:2032–2034, 1993.
- [14] John H. Holland. Genetic algorithms. *Scientific American*, 267:66–72, 1992.
- [15] Yongqian Liu, Sang-Gyu Park, and A. M. Weiner. Enhancement of narrow-band terahertz radiation from photoconducting antennas by optical pulse shaping. *Opt. Lett.*, 21:1762–1764, 1996.
- [16] H. S. Margolis, M. R. Harper, and S. N. Lea. Metrological applications of femtosecond optical pulse shaping. *NPL report, unrestricted*, DEM-EM-010:19–21, 2006.
- [17] M. R. Harper, R. A. Dudley, S. N. Lea, and H. S. Margolis. Control of terahertz pulse generation by optical pulse shaping. *Joint 32nd International Conference on Infrared and Millimetre Waves and 15th International Conference on Terahertz Electronics*, Invited Keynote Paper, 2007.
- [18] E. A. Manykin and A. M. Afanas'ev. On one possibility of making a medium transparent by multiquantum resonance. *Sov. Phys. JETP*, 25:828–830, 1967.

- [19] Yi-Yian Yin, Ce Chen, D. S. Elliott, and A. V. Smith. Asymmetric photoelectron angular distributions from interfering photoionization processes. *Phys. Rev. Lett.*, 69:2353–2356, 1992.
- [20] W. Pötz. Coherent control of terahertz radiation from semiconductor nanostructures. *Appl. Phys. Lett.*, 72:3002–3004, 1998.
- [21] W. C. Hurlbut, B. J. Norton, N. Amer, and Y. S. Lee. Manipulation of terahertz waveforms in nonlinear optical crystals by shaped optical pulses. *J. Opt. Soc. Am. B*, 23:90–93, 2006.
- [22] G. Kurizki, M. Shapiro, and P. Brumer. Phase-coherent control of photocurrent directionality in semiconductors. *Phys. Rev. B*, 39:3435–3437, 1989.
- [23] E. Dupont, P. B. Corkum, H. C. Liu, M. Buchanan, and Z. R. Wasilewski. Phase-controlled currents in semiconductors. *Phys. Rev. Lett.*, 74:3596–3599, 1995.
- [24] R. Atanasov, A. Haché, J. L. P. Hughes, H. M. van Driel, and J. E. Sipe. Coherent control of photocurrent generation in bulk semiconductors. *Phys. Rev. Lett.*, 76:1703–1706, 1996.
- [25] A. Haché, J. E. Sipe, and H. M. van Driel. Quantum interference control of electrical currents in GaAs. *IEEE J. Quant. Elec.*, 34:1144–1154, 1998.
- [26] A. J. A. Smith, A. G. Roddie, and D. Henderson. Electrooptic sampling of low temperature gas pulse generators for oscilloscope calibration. *Optical and Quantum Electronics*, 28:933–943, 1996.
- [27] A. Haché, Y. Kostoulas, R. Atanasov, J. L. P. Hughes, J. E. Sipe, and H. M. van Driel. Observation of coherently controlled photocurrent in unbiased bulk GaAs. *Phys. Rev. Lett.*, 78:306–309, 1997.
- [28] A. J. A. Smith, M. R. Harper, and M. A. Basu. Progress in high-speed time domain metrology at NPL. In *ARFTG Microwave Measurements Conference*, pages 35–43, 2003.

- [29] M. R. Harper and A. J. A. Smith. Coherent control of electrical currents. Technical Report NPL Report – CETM S132, National Physical Laboratory, 2003. Restricted – Commercial.
- [30] J. M. Diels, J. J. Fontaine, I. C. McMichael, and F. Simoni. Control and measurement of ultrafast pulse shapes (in amplitude and phase) with femtosecond accuracy. *Appl. Opt.*, 24:1270–1282, 1985.
- [31] D. J. Kane and R. Trebino. Characterization of arbitrary femtosecond pulses using frequency-resolved optical gating. *IEEE J. Quant. Elec.*, 29:571–579, 1993.
- [32] Rick Trebino. *Frequency-resolved optical gating: The measurement of ultrashort laser pulses*. Kluwer Academic Publishers, PO Box 322, 3300 AH Dordrecht, The Netherlands, 2000.
- [33] R. Trebino and D. J. Kane. Using phase retrieval to measure the intensity and phase of ultrashort pulses: frequency-resolved optical gating. *J. Opt. Soc. Am. A*, 1993.
- [34] A. Baltuška, M. S. Pshenichnikov, and D. A. Wiersma. Amplitude and phase characterisation of 4.5-fs pulses by frequency-resolved optical gating. *Opt. Lett.*, 23:1474–1476, 1998.
- [35] M. R. Harper and A. Butterworth. Private correspondence.

# Appendix A

## LabVIEW Code

### LabVIEW

Much of the work covered in this thesis required original computer programs to be written, both for control of the instruments used and for acquisition or processing of measurements. The majority of these programs were coded in LabVIEW, a National Instruments software product. LabVIEW differs from normal programming languages in that it is not textual but instead graphically based. Commands are represented by symbols; for instance, a triangle containing a + is used to denote an addition operator. Variables are passed between the symbols by ‘wires’. Flow-of-control statements such as For...Next loops or If...Then conditionals are represented by boxes.

Each LabVIEW program or subroutine is known as a VI (Virtual Instrument) and consists of two parts: a *diagram* which contains all the symbols and wires that make up the program, and a *front panel* which contains *controls* and *indicators*, which represent those variables that can be passed into or out of the VI by a calling VI. The front panel can also act as a graphical user interface for program operation. LabVIEW is chiefly designed for interfacing with laboratory equipment but is flexible enough for purely programmatical uses as well.

## The programs

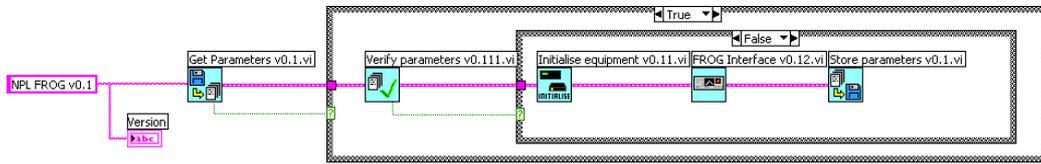
This appendix presents the code for the two main programs written and used in the course of this thesis and some of their associated subroutines. It does not present the code for every single VI, nor display the code for every conditional case within a VI. To do this would greatly increase the size of this thesis for little benefit. The intention is not to list the code to enable replication but to give the reader some insight into the software developed and used throughout this research.

The two main programs are the software used for spectrogram acquisition with the FROG, and the genetic algorithm software used for the THz work. In most cases just the VI diagram is presented. The front panel is shown where it is used as a user interface. Also shown for each is the *hierarchy*, which shows how the different VIs interlink with the main calling routine.

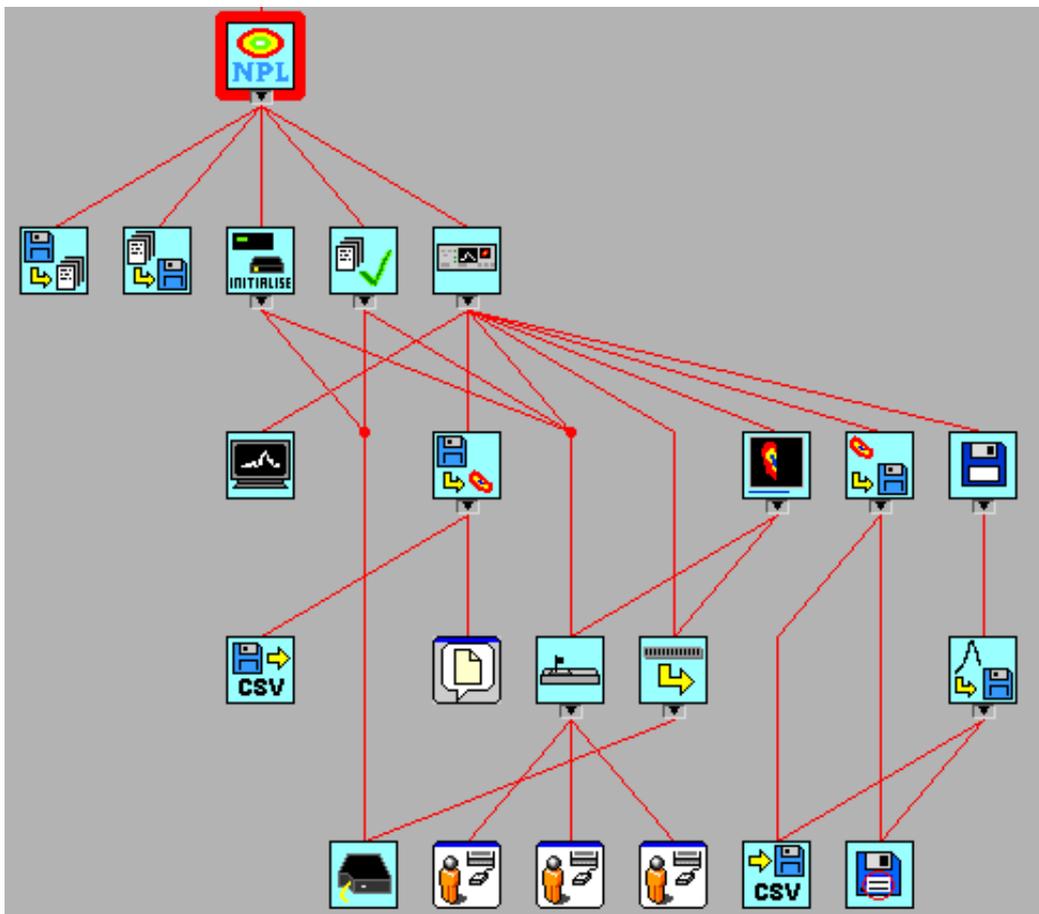
VIs written by the author can be identified by their blue or green backgrounds. VIs written by other parties are identified in the accompanying captions.

## A.1 The FROG program

The FROG software consists of a user-interface that allows control over the FROG delay line, the spectrometer integration time and the acquisition of spectrograms. The spectrometer output can be continuously monitored and saved to file. Background measurements can be made for subtraction purposes. Information such as previous settings and the last user ID are recalled from a ASCII file and saved on exiting the program for user friendliness.

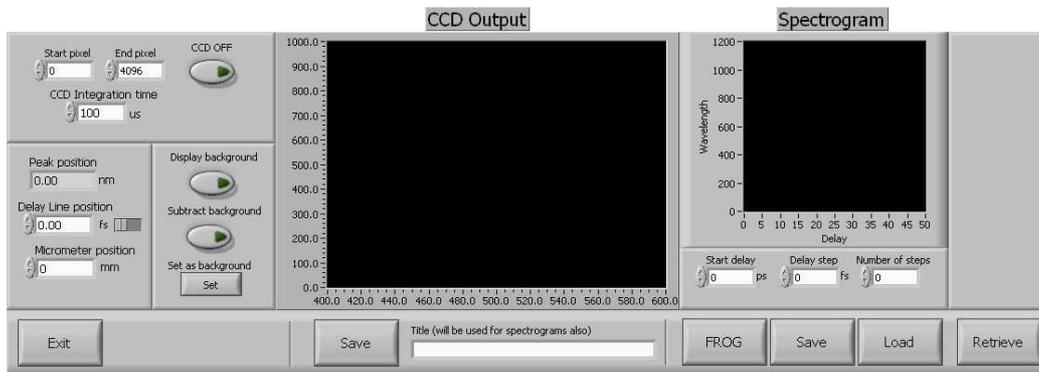


Code A.1: The main calling routine for the FROG program.

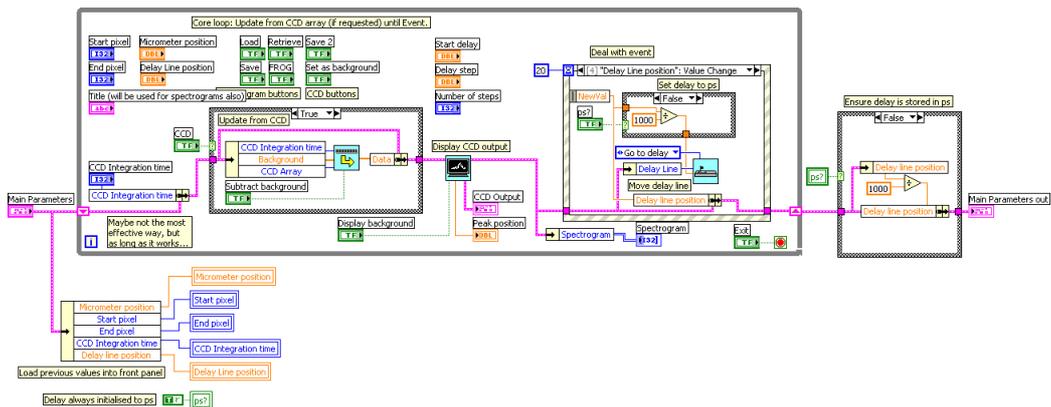


Code A.2: The hierarchy for the FROG program. The VIs with dark-blue top edges are LabVIEW's own user-interface VIs.

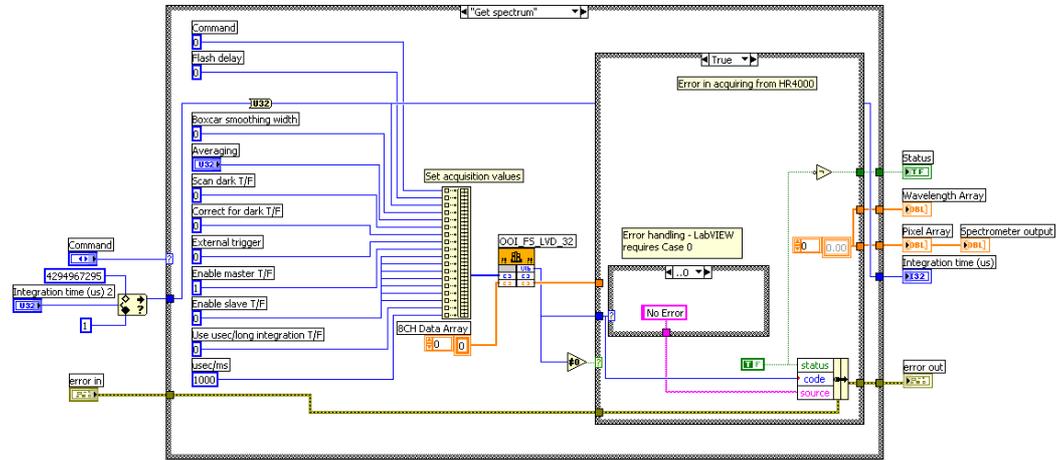
## APPENDIX A. LABVIEW CODE



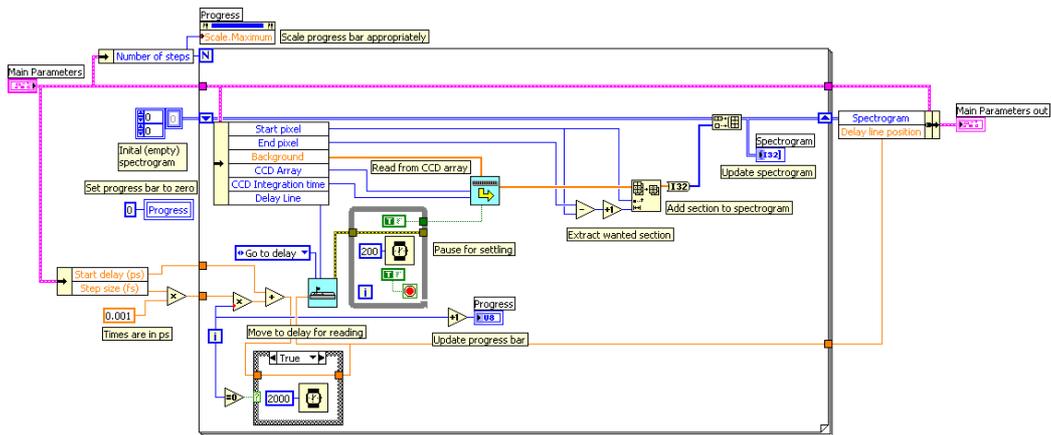
Code A.3: The user interface for the FROG software. To the left, the user is able to set the delay line position (with a switch to change between ps and fs delays), the integration time of the spectrometer and which pixels of the spectrometer are used during measurement. There are also controls to allow the subtraction of background signals. On the right, various spectrogram parameters can be set.



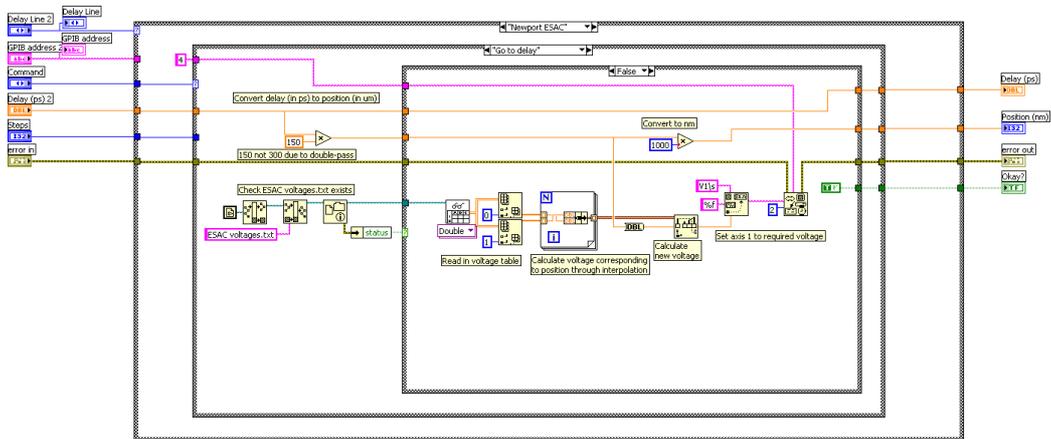
Code A.4: The user interface routine. This is the core of the software during execution.



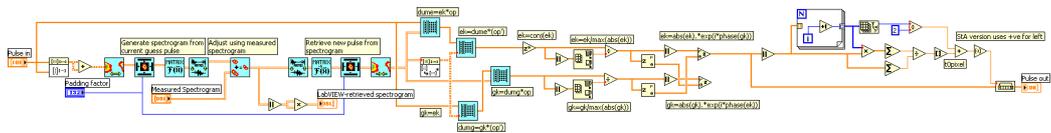
Code A.5: Driver routine for control of the Ocean Optics HR4000 spectrometer. A DLL (dynamic link library) routine is accessed using a 'code' module. Multiple commands and requests can be handled by this driver. Shown is a read of the CCD array.



Code A.6: Routine to automatically acquire a spectrogram according to the supplied parameters.



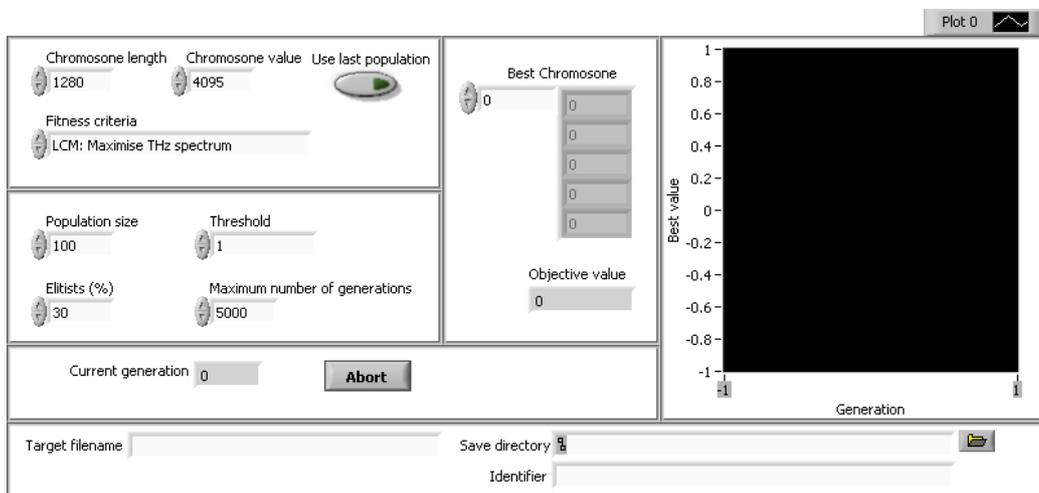
Code A.7: Driver routine for control of the FROG delay line. Commands are issued via GPIB. Shown is the move to a given delay using the look-up table stored in a separate file.



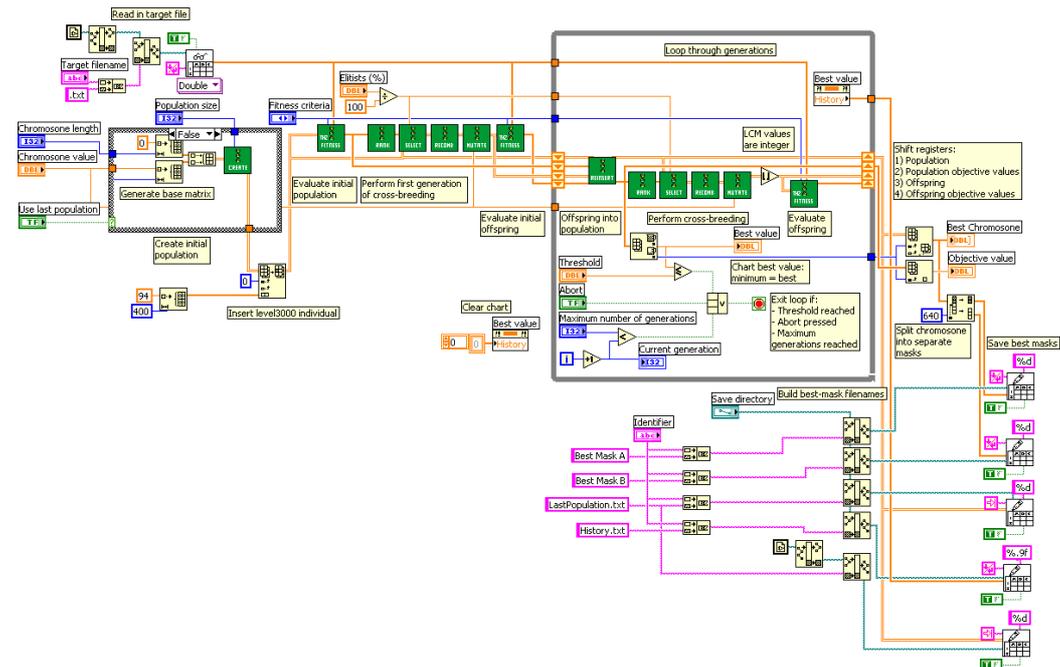
Code A.8: LabVIEW implementation of the MATLAB retrieval algorithm. This is work-in-progress. Currently the two versions of the algorithm diverge within a single iteration — differences in the FFT routines are a possible cause.

## A.2 The GA program

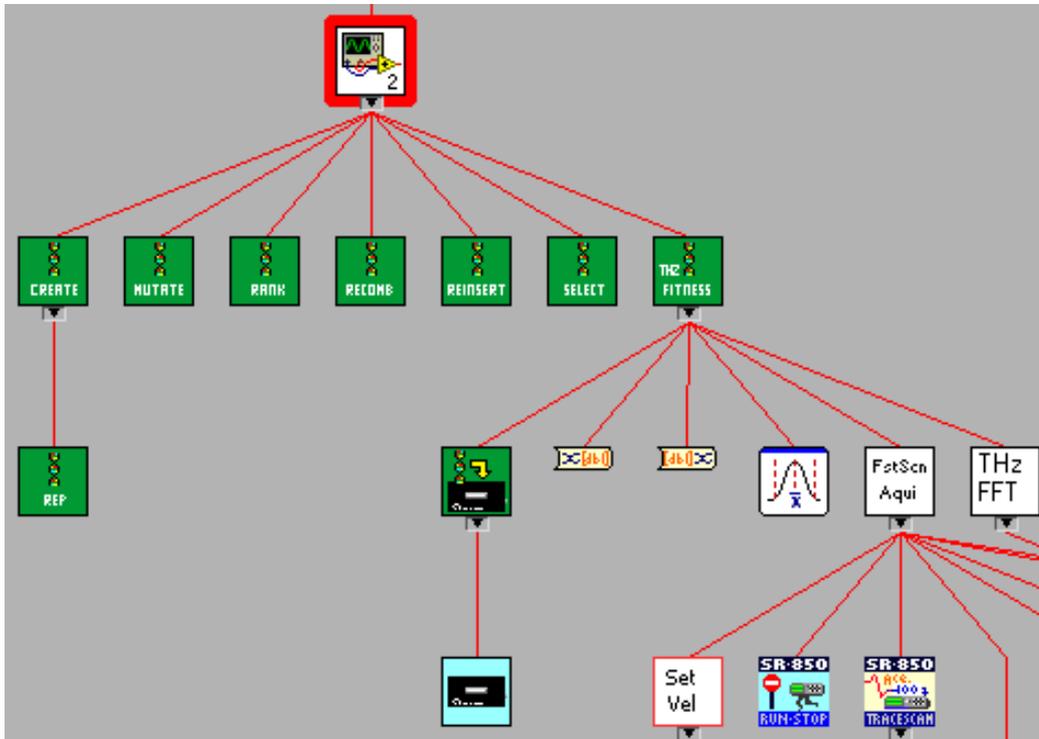
The genetic algorithm software is currently less polished than the FROG software. The user must set the program parameters, including the fitness requirement, in the front panel before running the code. During execution, the fitness history is plotted and the user may abort the program to ‘force’ a solution, i.e. the program writes the current best solution to a file, whether or not the threshold value has been reached, before exiting. This version of the software is designed exclusively for use with the THz and pulse shaper system, but the genetic algorithm code (in green) can be used in any suitable program. This genetic algorithm code is based on the MATLAB code provided by Reading University, with minor modifications and corrections.



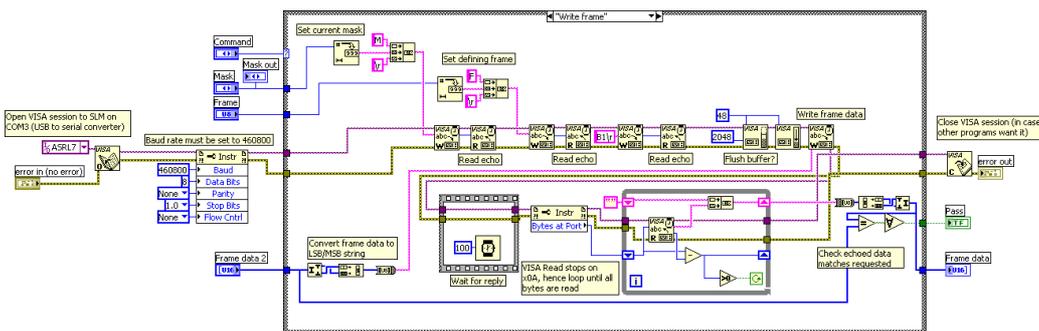
Code A.9: The front panel for the genetic algorithm. The user enters the various parameters here before running the code.



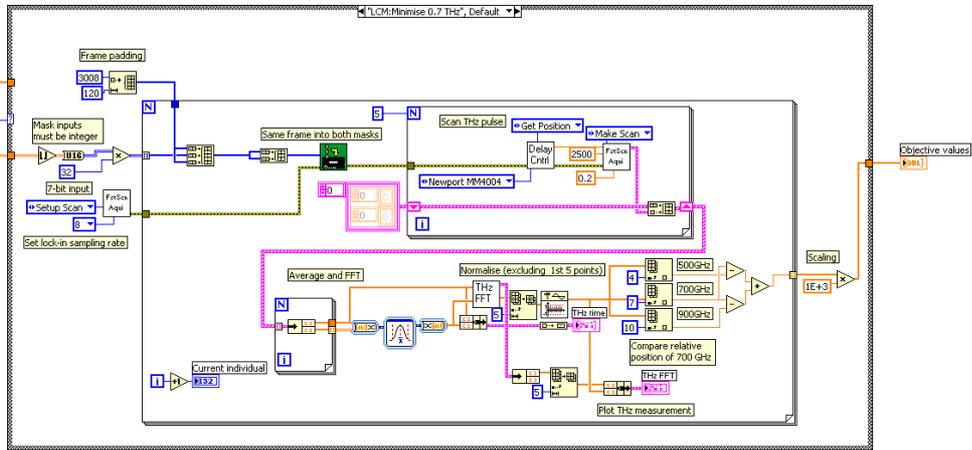
Code A.10: The main code for this implementation of the genetic algorithm. The core is the grey box which represents a Do...While loop. The code to the left of this prepares and evaluates the starting population. An option is included to use a previous population saved to a file. On completion the program outputs four files to a user-defined directory: The best values found for Mask A within the LCM; the best values found for Mask B; the complete population at the end of the run; and the history of the algorithm progress. A fifth file, the complete population, is stored in the same directory as the VI so that it can be used as a starting point on the next run, if desired.



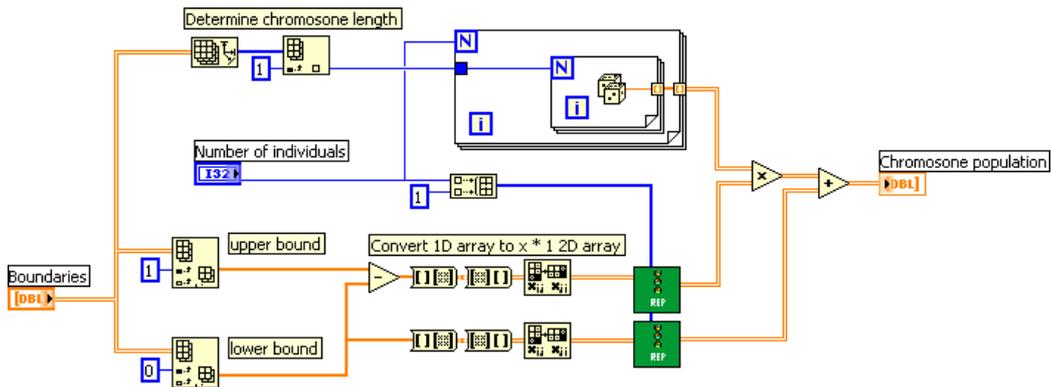
Code A.11: The hierarchy for the GA program. The VIs without blue or green backgrounds were supplied by R. A. Dudley (NPL) for acquisition of THz waveforms using EOS.



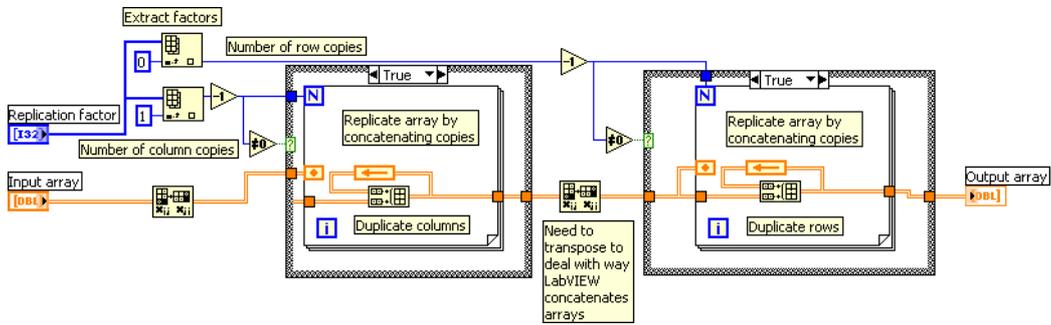
Code A.12: Driver routine for the CRi LCM. A DLL (dynamic link library) routine is accessed using a 'code' module. Multiple commands and requests can be handled by this driver. The routine shown here sets a mask within the LCM according to the data supplied to the routine.



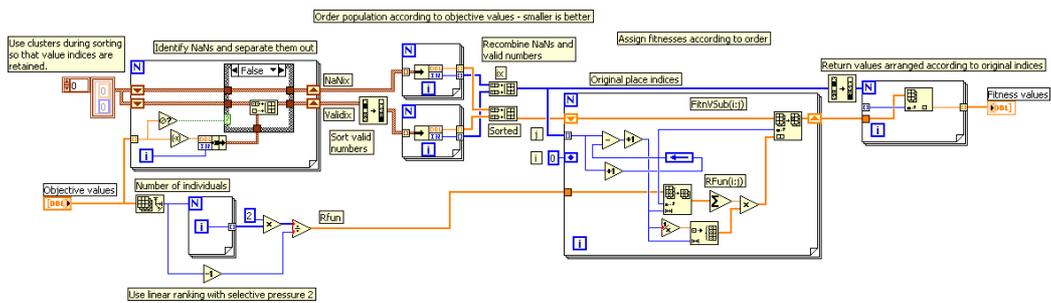
Code A.13: Fitness evaluation routine for minimising the 0.7 THz signal. The pulse shaper LCM is set and then five EOS measurements carried out. These are then averaged and Fourier transformed. The fitness is evaluated by comparing the value at 0.7 THz with those at 0.5 THz and 0.7 THz. The VIs FstScn Acqui, Delay Cntrl and THz FFT were supplied by R. A. Dudley.



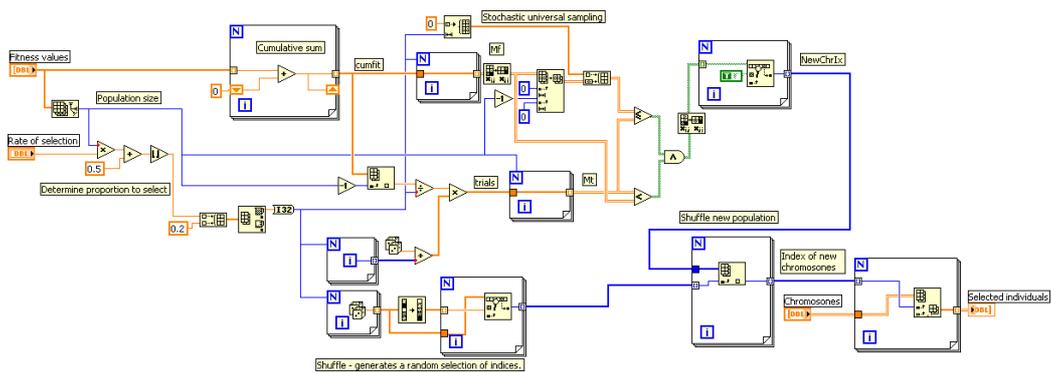
Code A.14: Create routine to generate a starting population for the genetic algorithm.



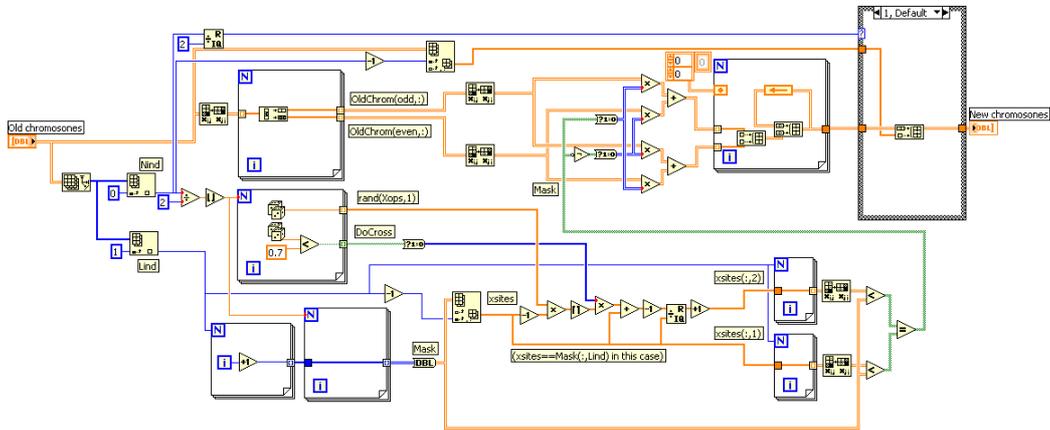
Code A.15: Replicate routine to set the bounds of the individuals to be used in the genetic algorithm.



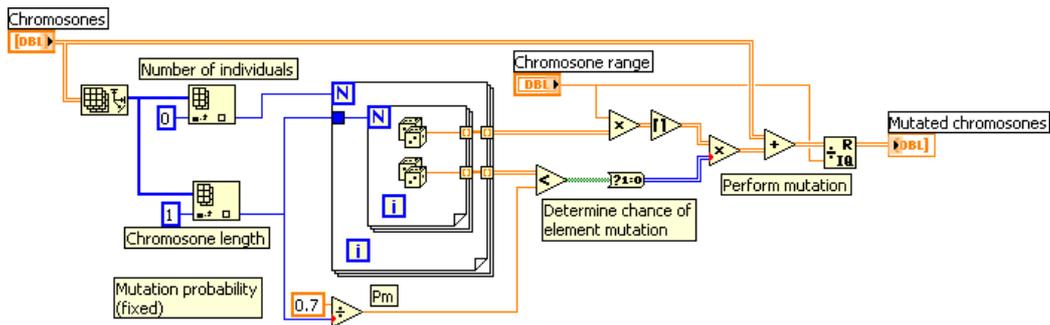
Code A.16: Ranking routine to order the current population of individuals according to their fitnesses.



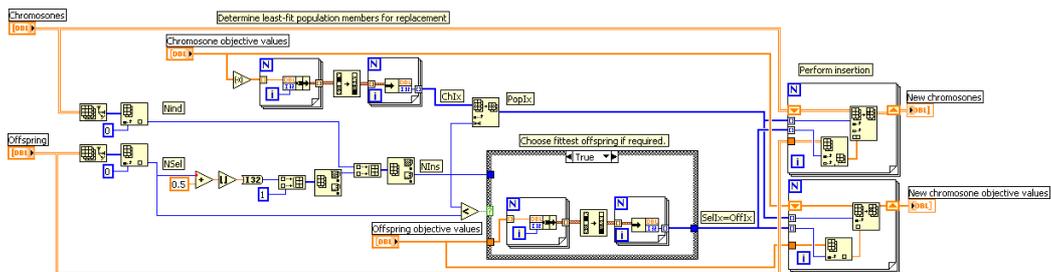
Code A.17: Select routine to select individuals from the current population for breeding. The selection is made using a stochastic sampling method.



Code A.18: Recombine routine to create offspring from those individuals selected for breeding. This is done by ‘crossover’, where for each breeding pair a random point is selected along their length, the two individuals split and matching portions swapped.



Code A.19: Mutate routine to introduce mutations into a population of newly-created individuals.



Code A.20: Routine to replace the least-fit members of the population with the new individuals. Note that older individuals are replaced by newer individuals even if the newer individuals have lower fitness.

# Appendix B

## Reducing the universe

The CRi LCM has two 640 pixel CCD arrays for a total of 1280 pixels, each of which can be independently set to one of 4096 voltage values. For the genetic algorithm, this gives a total of  $4096^{1280}$  possible individuals in the universe that represent all possible pulse-shaper settings. However, this is based on the initial assumption that all possible mask settings represent possible valid solutions to the problem at hand, yet in reality this is not the case. There are a number of reasons why certain individuals can be deemed invalid, and by identifying these before a time-consuming fitness test is carried out or, better still, by preventing the algorithm generating such individuals in the first place, the universe can be reduced and the speed of the algorithm increased as a result.

The most obvious requirements of the correct solution in the case of the LCM are as follows:

- **Phase-only control:** As mentioned in Section 2.6, if it is known that only the spectral phase of the pulse is to be changed, both masks in the LCM must be set to the same frame and so only individuals of 640 length are required, reducing the universe by a large factor. (*Not* half, since the total possible members of the universe will reduce from  $4096^{1280}$  to  $4096^{640}$ , i.e. by a factor of  $8.1 \times 10^{2311}$ .) However, this relies on the correctness of the assumptions that no amplitude modulation is required and that applying the same mask to both frames does not

introduce any amplitude modulations.

- **Reduced subset of the LCM:** A grating spreads the pulse spectrally in space and a lens focuses each spectral slice onto a pixel of the LCM placed at the focus of the lens. Ideally the grating and lens arrangement would be such that the spectrum of the pulse fills the entire array, with no pixels unused and no part of the spectrum falling outside the array. In practice it is found that only part of the LCM array is used, resulting in a number of pixels  $n$  being effectively redundant. The settings of these pixels are therefore of no consequence and so the individual can be shorter by length  $n$ , reducing the universe by a factor of  $4096^n$ .
- **Nyquist sampling:** The most important limitation to the individuals is that they fulfil the Nyquist requirement for the pulse-shaper (Section 2.4.1), i.e. the phase change from one pixel to the next must not exceed  $2\pi$ . Individuals that do not meet this requirement are therefore invalid solutions. Determining the resulting decrease in the universe is not simplistic, but clearly it will be significant.

Of these three, the first two are the simplest to implement, although the phase-only case will not be suitable for all applications. These two cases also have the advantage that they limit the universe from the beginning, i.e. invalid individuals are never accessed by the genetic algorithm and so never take up time-cycles during evolution. The third case, related to the Nyquist requirement, is trickier to implement. Instead of being independent, each value within the individual will be dependent on the values either side of it. Also, that dependence is not a simple relation; the phase-change from pixel to pixel represented by a difference in voltage levels changes according to the wavelength controlled by the pixels in question and even changes with the absolute values of the voltages. As a result, a number of possible methods of producing individuals that satisfied the Nyquist requirement were considered.

The most straightforward one would seemingly be to check a new individual to see if it satisfies the Nyquist requirement and discard it if not (a ‘fatal mutation’ approach). Although simple, the disadvantage of this is that unsuitable individuals still take up time-cycles (though less than if their fitness

were evaluated). If an overly-large proportion of the universe consists of these invalid individuals then an excessive amount of time may be spent simply trying to produce valid individuals, slowing the evolution of the algorithm. This may be a continual problem, as it is not clear that a population of valid individuals will not produce invalid ones (see the discussion on cross-over and mutation below).

Another approach considered was to define the values so that they represent purely the degree of phase change from one pixel to the next. The first value in the individual would be the absolute phase change of the first pixel (there is no reason that each value in the individual must represent the same thing — only that the individual as a whole describes a valid state of the system), the second value would be the phase change of the second pixel relative to the first, the third value the phase change of the third pixel relative to the second and so on. A value  $k$  could vary from 0 to  $2m$ , where the phase change represented by the value is given by  $\frac{(k-m)\pi}{m}$ . In this way no pixel can vary more than  $\pi$  from its neighbours, fulfilling the Nyquist requirement. Some routine would then be required to calculate the corresponding voltage levels for each pixel. Additionally the universe can be greatly reduced in size by setting  $m$  to a low number (compared to 4096), at the cost of finesse. (Note that some routine will also be required to ensure that the cumulative phase-shift to be applied to any pixel does not exceed the ability of that pixel.)

Although this approach provides a means of generating an individual guaranteed to meet the Nyquist requirement, there are several aspects that may hamper the genetic algorithm. The influence of the first value of the individual, the starting phase shift, may be too great, since it will have an effect on the voltages of all pixels. In the phase-only system this may not be much of an issue (as the only result will be a time-shift) but in the amplitude-and-phase system with its complicated relationship between both masks this could have drastic consequences that suppress the effect of the remainder of the individual. There is also a crucial problem with the breeding of individuals. When a cross-over point is selected the second part of the new individual will have different absolute values to the parent that the second

part came from, and therefore may not inherit those characteristics that were beneficial, crippling the operation of the algorithm.

It becomes clear that to ensure inheritance occurs correctly between parents and children, it will be necessary to use absolute phase values for the individuals. Ensuring that these individuals fulfil the Nyquist requirement will therefore necessitate subtle changes to the behaviour of the creation, breeding and mutation routines.