

Spectral interferometry for the complete characterisation of near infrared femtosecond and extreme ultraviolet attosecond pulses.

Adam S. Wyatt
Wolfson College, Oxford



Submitted for the degree of Doctor of Philosophy
Michaelmas term 2007

Supervised by
Prof. Ian A. Walmsley

Clarendon Laboratory
University of Oxford
United Kingdom

Abstract

This thesis describes methods for using spectral interferometry for the complete space-time characterisation of few-cycle near-infrared femtosecond pulses and extreme ultraviolet (XUV) attosecond pulses produced via high harmonic generation (HHG).

Few-cycle pulses tend to exhibit one or more of the following: (1) an octave-spanning bandwidth, (2) a highly modulated spectrum and (3) space-time coupling. These characteristics, coupled with the desire to measure them in a single-shot (to characterise shot-to-shot fluctuations) and in real-time (for online optimisation and control) causes problems for conventional characterisation techniques. The first half of this thesis describes a method, based on a spatially encoded arrangement for spectral phase interferometry for direct electric-field reconstruction (SEA-SPIDER). SEA-SPIDER is demonstrated for sub-10 fs pulses with a central wavelength near 800 nm, a bandwidth over 350 nm, and a pulse energy of several nano-Joules. In addition, the pulses exhibit a modulated spectrum and space-time coupling. The spatially-dependent temporal intensity of the pulse is reconstructed and compared to other techniques: interferometric frequency-resolved optical gating (IFROG) and spectral phase interferometry for direct electric field reconstruction (SPIDER). SEA-SPIDER will prove useful in both femtoscience, which requires accurate knowledge of the space-time character of few-cycle pulses, and in HHG, which requires the precise knowledge of the driving pulse for seeding into simulations and controlling the generation process itself.

Pulses arising from HHG are known to exhibit significant space-time coupling. The second half of this thesis describes how spectral interferometry may be performed to obtain the complete space-time nature of these fields via the use of lateral shearing interferometry. Finally, it is shown, via numerical simulations, how to extend the SPIDER technique for temporal characterisation of XUV pulses from HHG by driving the process with two spectrally-sheared driving pulses. Different experimental configurations and their applicability to different laser systems are discussed. This method recovers the space-time nature of the harmonics in a single shot, thus reducing the stability constraint currently required for photoelectron based techniques and may serve as a complimentary method for studying interactions of XUV attosecond pulses with matter.

To Marilyn, Kevin and Magdalena...

Acknowledgements

Firstly I would like to thank my supervisor, Prof. Ian A. Walmsley for his support, wisdom and excellent insight into life (and the underpants of particle physicists. . .). I would also like to thank Ian for allowing me to work in such a diverse group academically, geographically and socially, and for teaching me that physics is not just about little creatures for measuring brief flashes of light, especially those which come in different colours (red, green . . .). I am very grateful to everyone within the ultrafast group for providing me many opportunities to rant, drink and of course discussions on physics, for without which this thesis would never have been completed. Special mention in that respect should go to Antoine Monmayrant, who was a great source of inspiration and excellent at verbal tennis; and to Tobias Witting for all his help in the lab. I would also like to thank my collaborators: Eric Cormier for his simulations; Jon Marangos, John Tisch and Roland Smith for allowing us to use their lab and their scientific input and support; Joseph Robinson and Charles Haworth for building the apparatus and helping me use it only to allow me to complain, Günter Steinmeyer and Gero Stibenz for working on the few-cycle metrology; Jens Biegert, Amelle Zaïr and the rest of the Keller group for allowing us to work on their attosecond system. Without any of these people and their equipment, this thesis would not have been possible and so I am very grateful. I would also like to thank the SPIDER elders Ellen Kosik and Alex Radunsky for teaching me their ways, Piotr for his machine skills and Simon-Pierre for additional SPIDER stuff. Of course a DPhil is not just a purely academic affair, but comes with many social responsibilities, so I would like to thank Pete,

Alex and Daryl especially for making it so enjoyable, and Josh for all those cups of tea, for which no British person could be without. Not to forget the other children: Sharp (Lijian), Felix, Tobias, Offir and Phil and the old fogeys: Christine, Laura, Rob, Antoine, Brian, Matthijs and Jeff for their pub visits, and to the rest of the group. Many thanks to Laurie and Rich for their friendship and support throughout both my postgraduate and undergraduate days, and sticking with me through my many annoyances and Pelle for his crazy Swedish ways. Many thanks to my parents for enabling me to be myself and somehow placing me on the track to success. Finally, special thanks to my wife Magdalena who managed to survive these past four years, for her compassion, understanding and overall support.

Author's publications

Journal publications

1. E. M. Kosik, L. Corner, A. S. Wyatt, E. Cormier, I. A. Walmsley, and L. F. Dimauro. “**Complete characterization of attosecond pulses.**” *Journal Of Modern Optics*, 52(2-3):361–378, January 2005.
2. E. Cormier, I. A. Walmsley, E. M. Kosik, A. S. Wyatt, L. Corner, and L. F. DiMauro. “**Self-referencing, spectrally, or spatially encoded spectral interferometry for the complete characterization of attosecond electromagnetic pulses.**” *Physical Review Letters*, 94(3):033905, January 2005.
3. E. Cormier, L. Corner, E. M. Kosik, I. A. Walmsley, and A. S. Wyatt. “**Spectral phase interferometry for complete reconstruction of attosecond pulses.**” *Laser Physics*, 15(6):909–915, June 2005.
4. A. S. Wyatt, I. A. Walmsley, G. Stibenz, and G. Steinmeyer. “**Sub-10 fs pulse characterization using spatially encoded arrangement for spectral phase interferometry for direct electric field reconstruction.**” *Optics Letters*, 31(12):1914–1916, June 2006.
5. G. Stibenz, C. Ropers, C. Lienau, C. Warmuth, A. S. Wyatt, I. A. Walmsley, and G. Steinmeyer. “**Advanced methods for the characterization of few-cycle light pulses: a comparison.**” *Applied Physics B-Lasers And Optics*, 83(4):511–519, June 2006.
6. A. Zaïr, M. Holler, A. Guandalini, F. Schapper, J. Biegert, L. Gallmann, U. Keller, A. S. Wyatt, A. Monmayrant, I. A. Walmsley, E. Cormier, T. Augustine, J. P. Caumes and P. Salières. “**Quantum Path Interferences in High-order Harmonic Generation**” *Physical Review Letters*, submitted July 2007.

Conferences

1. A. S. Wyatt, I. A. Walmsley, G. Stibenz and G. Steinmeyer. “**Characterization of Sub-10-fs Pulses Using Spatially Encoded Arrangement for SPIDER**” *Conference on Lasers and Electro-Optics*, May 2006
2. A. S. Wyatt, T. Witting, A. Monmayrant, I. A. Walmsley, C. Haworth, J. S. Robinson, J. W. G. Tisch and J. P. Marangos. “**Characterizing Spatio-Temporal Coupling of Extreme Ultraviolet Ultrashort Pulses from High Harmonic Generation**” *Conference on Lasers and Electro-Optics*, May 2007
3. A. Zaïr, M. Holler, A. Guandalini, F. Schapper, J. Biegert, U. Keller, P. Salières, T. Auguste, E. Cormier, A. S. Wyatt, A. Monmayrant and I. A. Walmsley. “**Two-Quantum-Path Interferences in High Order Harmonic Generation**” *Conference on Lasers and Electro-Optics*, May 2007 .
4. T. Witting, A. S. Wyatt, A. Monmayrant, I. A. Walmsley, C. Haworth, J. S. Robinson, J. W. G. Tisch and J. P. Marangos. “**Characterising spatio-temporal coupling of extreme ultraviolet ultrashort pulses from high harmonic generation**” *Conference on Lasers and Electro-Optics / Europe*, June 2007
5. A. Zaïr, M. Holler, A. Guandalini, F. Schapper, J. Biegert, U. Keller, P. Salières, T. Auguste, E. Cormier, A. S. Wyatt, A. Monmayrant and I. A. Walmsley. “**Quantum-path interferences in high order harmonic generation**” *Conference on Lasers and Electro-Optics / Europe*, June 2007 .

Contents

Acknowledgements	vii
Author's publications	ix
List of Figures	xv
List of Tables	xix
List of abbreviations	xxi
1 Introduction	3
1.1 Time-resolved spectroscopy	3
1.2 Attoscience and the UK Attosecond Technology Project	6
1.3 Metrology of ultrashort pulses and high harmonic generation	8
1.4 Thesis outline	11
2 Metrology of ultrashort pulses	15
2.1 Mathematical description of electromagnetic pulses	16
2.1.1 Time-frequency representations of ultrashort pulses	18
2.2 Propagation of electromagnetic pulses	23
2.2.1 Propagation of electromagnetic fields in dielectrics	29
2.2.2 Chromatic Dispersion	31
2.2.3 The role of phase in ultrashort pulses — summary	36
2.3 Space-time coupling of ultrashort pulses	39
2.3.1 Frequency dependent mode size	39
2.3.2 Focusing of a short optical pulse	41
2.3.3 Angular dispersion	42
2.3.4 Spatial chirp	44
2.4 Measuring the wavefront of optical beams	45
2.4.1 Shack-Hartmann wavefront sensor	45
2.4.2 Point diffraction interferometry	46
2.4.3 Lateral shearing interferometry	48
2.4.4 Radial shearing interferometry	49

2.5	Complete temporal characterisation of ultrashort optical pulses	50
2.5.1	Field autocorrelation	51
2.5.2	Intensity and interferometric autocorrelation	52
2.5.3	FROG	54
2.5.4	Chronocyclic tomography	57
2.5.5	SPIDER	60
2.6	Accuracy, precision and consistency	67
2.6.1	RMS field error	69
2.7	Summary	70
3	Complete characterisation of near infrared few-cycle pulses	71
3.1	Metrology of few-cycle pulses	74
3.1.1	Errors in SPIDER pulse reconstructions	78
3.1.2	Carrier-envelope phase measurement in few-cycle pulses	89
3.2	Spatially encoded arrangement for SPIDER	92
3.2.1	Concept	92
3.2.2	Intuitive interferograms	94
3.2.3	SEA-SPIDER reconstruction algorithms	98
3.2.4	Experimental setup for few-cycle pulses	101
3.2.5	SEA-SPIDER alignment	106
3.2.6	SEA-SPIDER results - Few-cycle Oscillator	108
3.2.7	SEA-SPIDER results - CPA	117
3.3	Summary	120
4	Attosecond metrology and high harmonic generation	123
4.1	HHG: Single-atom response	127
4.1.1	Three step semiclassical model	127
4.1.2	Quantum orbit model	133
4.1.3	Time dependent Schrödinger equation	137
4.2	HHG: Macroscopic response	139
4.2.1	Propagation of the intense low frequency field	140
4.2.2	Propagation of the weak harmonic field	141
4.2.3	Generalised phase-matching for HHG	142
4.3	Attosecond metrology via photoelectron spectroscopy	151
4.3.1	Autocorrelation	151
4.3.2	Theory of atomic XUV photoionisation in a laser field	152
4.4	Summary	158
5	Spectral interferometry for characterising space-time coupling of XUV attosecond pulses from high harmonic generation	161
5.1	Experimental setup	166
5.2	Alignment procedure	169
5.3	XUV lateral shearing interferometry results	179
5.4	Summary	185

6	Spectral phase interferometry for direct electric-field reconstruction of XUV attosecond pulses from high harmonic generation	187
6.1	XUV SPIDER via HHG	190
6.1.1	Spectrally encoded XUV-SPIDER	193
6.1.2	Spatially encoded XUV-SPIDER	194
6.2	Shearing the driving fields	196
6.2.1	Nonlinear frequency mixing	196
6.2.2	Pulse shaper	196
6.2.3	Amplitude shaping	198
6.2.4	Space-Time Coupling	200
6.3	Summary	202
7	Conclusion	205
7.1	Outlook	211
A	Maxwell's equations and the electromagnetic wave equation	215
A.1	Maxwell's equations	215
A.2	The wave equation	216
A.3	Analytic representation of the electromagnetic field	217
A.4	Fourier transforms	220
A.4.1	Cylindrically symmetric beams	221
B	Amplitude shaping beamsplitters for XUV-SPIDER	223
	Bibliography	227
	*	

List of Figures

1	Eadweard Muybridge’s 1878 “Horse in motion.”	1
1.1	Simulation of overexposure of a moving image.	4
1.2	The various roles of each institute within the UK attosecond consortium.	7
2.1	Wigner representation and marginals of different optical pulses.	20
2.2	Spectrographic representation and marginals of different optical pulses.	22
2.3	Gaussian beam propagation and spatial phase.	25
2.4	The effect of spectral phase on the temporal electric field.	32
2.5	Space-time coupling of an octave spanning pulse due to a frequency dependent mode size.	38
2.6	Apparent increase in the duration of an ultrashort pulse due to a lens.	42
2.7	Schematic of angular dispersion / pulse front tilt	43
2.8	Spatial chirp of a collimated beam results in pulse front tilt in the focal plane of a lens	45
2.9	Principle of the Shack-Hartmann wavefront sensor	46
2.10	Principle of point diffraction interferometry	47
2.11	Principle of lateral shearing interferometry	48
2.12	The field autocorrelation detected via a Michelson interferometer	52
2.13	Simulated interferometric autocorrelation data	53
2.14	Schematic of a single-shot SHG-FROG setup.	55
2.15	Schematic of the iterative FROG PCGPA algorithm.	56
2.16	Schematic of chronocyclic tomography via rotations of the Wigner distribution of the pulse.	58
2.17	Flow chart of the SPIDER reconstruction algorithm.	62
2.18	Concept of integration or concatenation in SPIDER.	63
2.19	Schematic of a conventional SPIDER concept.	65
3.1	Fractional ionisation of Argon from an ultrashort laser pulse.	72
3.2	Reconstructing the spectral phase over regions of low spectral intensity with SPIDER.	76
3.3	Error in the reconstructed pulse duration for a FTL pulse due to an error in the calibration of the SPIDER delay.	85

3.4	SPIDER pulse reconstruction averaged over many fluctuating pulses	88
3.5	SEA-SPIDER concept	95
3.6	A set of SEA-SPIDER interferograms showing their intuitive nature.	97
3.7	2D FFT SEA-SPIDER inversion routine	99
3.8	1D FFT SEA-SPIDER inversion routine.	101
3.9	SEA-SPIDER setup for few-cycle pulses.	102
3.10	Experimental setup for SEA-SPIDER	105
3.11	Spatially dependent spectral intensity and phase of a few-cycle oscillator, measured via SEA-SPIDER.	109
3.12	Spatially dependent FTL temporal intensity corresponding to a few-cycle oscillator	110
3.13	Spatially dependent temporal intensity of a few-cycle oscillator	110
3.14	Phase of 1 cm of fused silica measured using SEA-SPIDER and calculated via the Sellmeier equation.	113
3.15	Reconstruction of a few-cycle oscillator at the centre of the pulse from two different shears	114
3.16	Wigner distribution of a few-cycle oscillator showing complicated structure in the wings	115
3.17	Reconstructed spectral phase of a few-cycle oscillator using different techniques.	116
3.18	Temporal electric field envelope of a few-cycle oscillator reconstructed from different characterisation techniques	116
3.19	Spectrogram of a few-cycle oscillator showing a post-pulse in the wings of the spectrum	118
3.20	SEA-SPIDER interferograms of a CPA pulse with different amounts of chirp	119
4.1	Three step model of high harmonic generation	127
4.2	Time gating due to the high nonlinearity of the laser induced ionisation rate of Argon.	129
4.3	Classical electron trajectories from strong field ionisation	131
4.4	Schematic of generalised phase-matching for high harmonic generation.	146
4.5	phase-matching maps for HHG.	148
4.6	Interference fringes showing the intensity dependent dipole phase.	150
4.7	Three different regimes of performing characterisation of XUV attosecond pulses via PES.	154
5.1	Mach–Zehnder interferometer used to generate two identical few-cycle pulses	166
5.2	Experimental setup for performing XUV LSI via HHG.	168
5.3	Beam pointing alignment device	172
5.4	Phase difference between the two arms of the interferometer measured via SI	173
5.5	Measuring the wavefront aberrations of an interferometer	174
5.6	A simple adaptive mirror to compensate a spherical wavefront	175
5.7	The dependence of the spatial-spectral intensity of high harmonics as a function of the focal position relative to the gas jet along the propagation direction	176
5.8	Spatial-spectral distribution of harmonics from each arm of the interferometer	177

5.9	Propagation of XUV fields for performing lateral shearing interferometry . . .	178
5.10	Walk-off procedure for aligning XUV beams	178
5.11	Spatially and spectrally resolved interference pattern of XUV radiation generated from two sources of HHG	181
5.12	Extracted amplitude and phase from an XUV LSI measurement on harmonic 11 from argon in a gas target.	182
5.13	Extracted space-frequency phase from the central region of the beam.	184
6.1	Simulation of the single atom response for different driving frequencies	190
6.2	Overlap of harmonics driven with different central frequencies	192
6.3	Experimental setup for performing spectrally encoded XUV-SPIDER	193
6.4	Simulation of spectrally encoded XUV-SPIDER for different driving intensities.	194
6.5	Experimental setup for performing spatially encoded XUV-SPIDER	195
6.6	Simulation of spatially encoded XUV-SPIDER and the Fourier filtering process	195
6.7	Implementation of spectrally encoded SPIDER using a pulse shaper between the oscillator and CPA lasers	197
A.1	Analytic representation of a real signal.	219
B.1	Reflectivity curve for a broadband 50:50 beamsplitter	224
B.2	Reflectivity curve for a broadband beamsplitter which is slightly linearly dependent on frequency	224
B.3	Reflectivity curve for a broadband beamsplitter which is strongly linearly dependent on frequency	225

List of Tables

3.1	Pulse intensity error between different reconstruction methods.	117
3.2	FWHM pulse durations from different methods.	117
4.1	Ionisation potential of some Noble gases used for HHG.	128

List of abbreviations

Abbreviation	Description
AOM	acousto-optic modulator
AOPDF	acousto-optic programmable dispersive filter
BBO	β -barium borate
CCD	charged coupled device
CEO	carrier envelope offset
CEP	carrier envelope phase
CPA	chirped pulse amplification
CRAB	complete reconstruction of attosecond bursts
CW	continuous wave
EOM	electro-optic modulator
FEL	free electron laser
FFT	fast Fourier transform
FROG	frequency-resolved optical gating
FTL	Fourier transform limited
FWHM	full width half maximum
FOD	fourth order dispersion
GD	group dispersion
GDD	group delay dispersion
GV	group velocity
GVD	group velocity dispersion
HCFPC	hollow-core fibre pulse compressor
HHG	high harmonic generation
IAC	interferometric autocorrelation
IFFT	inverse fast Fourier transform
LSI	lateral shearing interferometry
ND	neutral density
NIR	near infrared
NOPA	non-collinear optical parametric amplifier

Abbreviation	Description
OPA	optical parametric amplifier
OPCPA	optical parametric chirped pulse amplifier
PCGPA	principle component generalised projection algorithm
PDI	point diffraction interferometry
PE	photoelectron
PES	photoelectron spectroscopy
PFT	pulse front tilt
QOM	quantum orbit model
RABBITT	reconstruction of attosecond beating by interference of two-photon transitions
RMS	root mean square
RSI	radial shearing interferometry
SCT	simplified chronocyclic tomography
SEA-	spatially encoded arrangement
SFA	strong field approximation
SFG	sum frequency generation
SHG	second harmonic generation
SHWS	Shack–Hartmann wavefront sensor
SNR	signal to noise ratio
SPIDER	spectral phase interferometry for direct electric-field reconstruction
STC	space-time coupling
TDSE	time-dependent Schrödinger equation
Ti:Sapph	titanium sapphire
TOD	third order dispersion
TRS	times-resolved spectroscopy
UV	ultraviolet
Vis.	visible
XUV	extreme ultraviolet

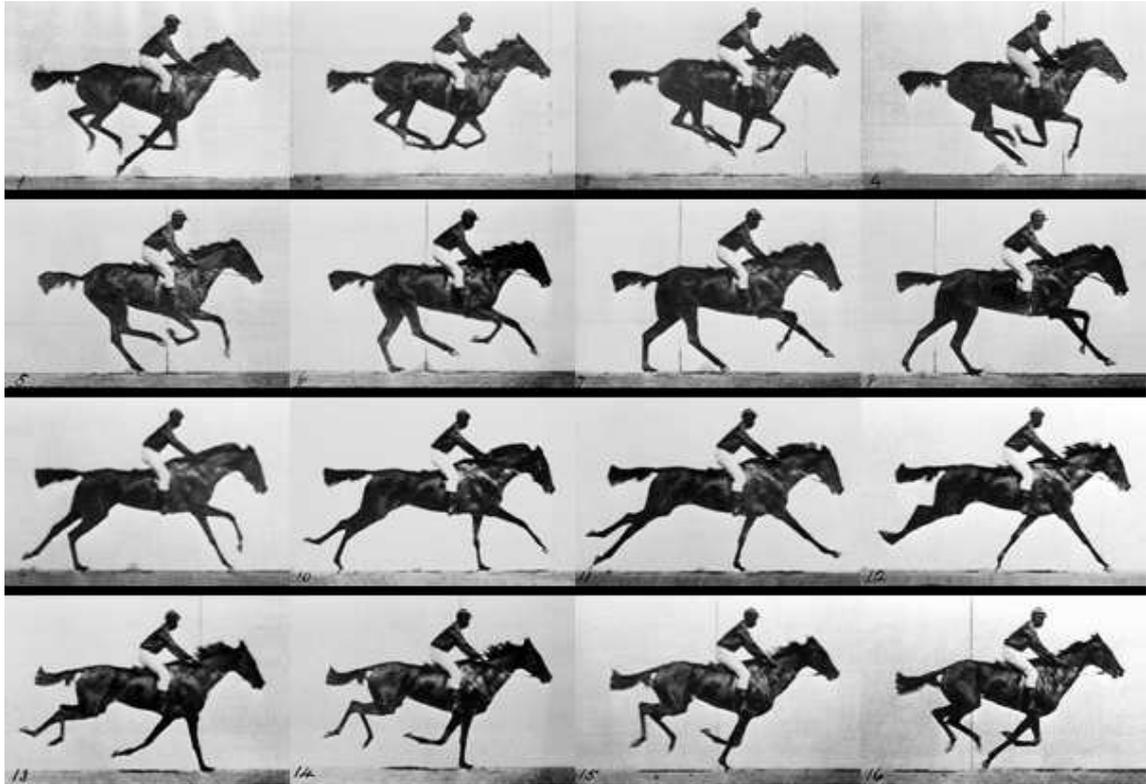


Figure 1: Eadweard Muybridge's 1878 "Horse in motion."

Introduction

1.1 Time-resolved spectroscopy

In 1878, Eadweard Muybridge developed the method of freeze-frame time-lapse photography by taking snapshots of a horse running (figure 1), which was the result of developments in three key areas: (1) synchronisation between the dynamics and the acquisition, (2) gating the acquisition to limit the exposure time and (3) improved detection. Synchronisation is essential to ensure the image is captured at the correct time, in this case via the use of trip-wires to trigger the camera. Gating is necessary to ensure that the system appears to be frozen during the exposure time of the detection, which was initially achieved via mechanical shutters to physically limit the exposure time. For example, if the timescale of the system dynamics being observed has a duration on the order of, or even longer than, the exposure time, then the image will become blurred (figure 1.1). Finally, the detection system needs to be sensitive enough to capture the image within the limited exposure time at the wavelength being recorded, which at that time required developments in photographic film.

Such a setup as devised by Muybridge could be seen as the first steps towards time-resolved spectroscopy (TRS), which is the study of dynamical processes via the interaction of light with matter.^[1] A laboratory TRS experiment is typically based around a pump-

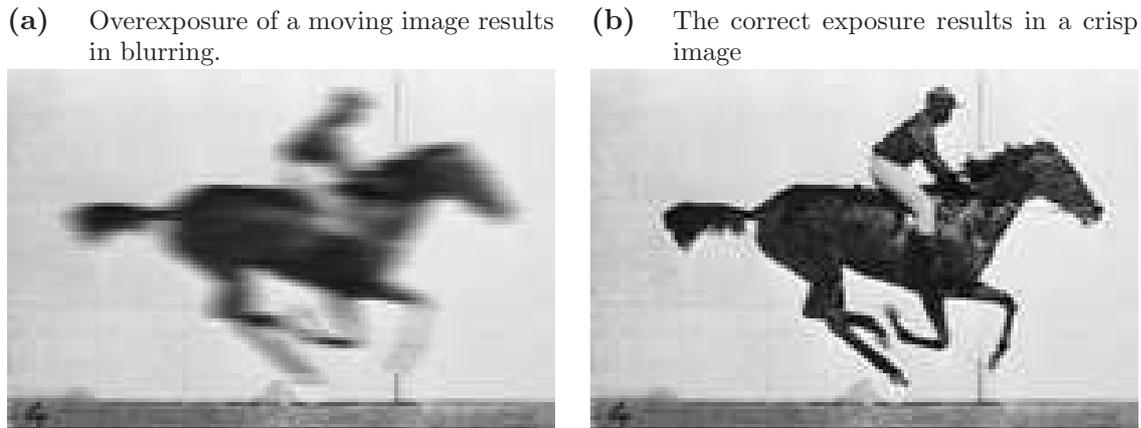


Figure 1.1: Simulation of overexposure of a moving image.

probe geometry and can be considered as an advanced form of flash photography. Although slightly different, TRS still requires the three features stated above. In a time-resolved pump-probe experiment, a dynamical system is driven by a pump (usually a laser pulse) and then probed by a short pulse at some time interval later. By varying the time interval between the pump and probe, it is possible to build up a movie of the dynamics, similar to a flick book, which can be then used to measure how the system relaxes back into equilibrium. Transient absorption spectroscopy is a pump-probe based experiment in which the absorption of specific wavelengths of light from the probe field is monitored as a function of time after a system is driven.^[2-4] In this case, the time resolution is set by the interaction of the probe with the system and the timescale for this to change. Time resolved fluorescence spectroscopy is another form of a pump-probe experiment.^[5-7] However, this time the excited system will emit light that varies as a function of time as the system relaxes back into equilibrium. The time resolution in this case is set by the detection method. If detected via a fast detector such as a fast photodiode and fast electronics, this is limited to nanoseconds or slower. Streak cameras^[8] offer a temporal resolution of picoseconds or slower. Gating the fluorescence with a short pulse provides the best resolution, which is set by the gate pulse duration; this currently can be as little as several femtoseconds ($1 \text{ fs} = 10^{-15} \text{ s}$) and most recently as several hundred attoseconds ($1 \text{ as} = 10^{-18} \text{ s}$). In either pump-probe form, the pulse duration must be much less than

the timescale of the motion to be captured. It is normally also necessary for the pump duration to be on the same timescale or less than the dynamics being driven to prevent ‘temporal smearing’ and to act as a well-defined clock to the probe pulse.

Initially the minimum exposure time for spectroscopy was limited by the use of mechanical shutters, but the desire to capture on a shorter timescale necessitated even shorter exposure times. One method is to limit the pulse duration used to probe the system; this is the idea behind flash photography: the image is acquired only for the pulse duration, thus ‘freezing’ the motion if the pulse duration is short enough. The development of flash lamps allowed motion on the microsecond timescale to be recorded. Nanosecond timescales were reached using pulsed ruby lasers, followed by picosecond and subsequently femtosecond timescales using mode-locked lasers. Nanosecond and picosecond timescales can now also be obtained using modern opto-electronics (e.g. Pockels cells or CCD chips).^[9,10] With femtosecond pulses, it is possible to track molecular motion, and thus build up real-time movies of chemical reactions. The success and importance of this field was illustrated when, in 1999, Ahmed Zewail was awarded the Nobel Prize in chemistry “for his studies of the transition states of chemical reactions using femtosecond spectroscopy”.^[3,4,11] In his work, it was possible to uniquely determine the exact makeup of intermediates in photosynthesised organic chemical reactions by measuring the bond lengths through time-resolved absorption spectroscopy. This was possible through the use of 100 fs duration optical pulses because molecular motion occurs on a timescale of fractions of a picosecond and dominates the timescale of the chemical reaction. Further examples of time-resolved femtosecond spectroscopy can be found in chemistry (e.g. coherent control)^[12], materials science (e.g. micro machining)^[13] and physics (e.g. non-equilibrium dynamics)^[14], illustrating its importance in a wide range of disciplines.

Electrons, being several orders of magnitude smaller than nuclei, tend to exhibit dynamics on the order of several hundred attoseconds to several femtoseconds. Thus in order to track electronic motion, it is necessary to be able to produce and utilise attosecond

pulses.^[15] Such a possibility promises to provide the same revolution in the understanding of non-equilibrium electron and nuclear motional dynamics in atoms, molecules and solids as femtoscience has provided for molecules.

1.2 Attoscience and the UK Attosecond Technology Project

Attoscience is the study of dynamics on the timescale of 10^{-15} – 10^{-18} seconds. It has been enabled by the development of light sources that facilitate controllable generation of bursts of electromagnetic radiation with durations in the range 10^{-15} – 10^{-17} seconds. Such sources have emerged from more than a decade's study of the interaction of high intensity laser radiation with atoms, leading to a table-top source of attosecond pulses in the extreme ultraviolet (XUV) region of the spectrum. The key applications of this research are in the study of dynamics of highly nonequilibrium electronic motion in atoms, molecules and solids, as well as excited-state vibrational motion in molecules. A deeper understanding of these short-timescale events, and their interplay in determining macroscopic properties of materials will provide the knowledge necessary for new technologies.

Although progress is well advanced in some areas of the discipline, there is still much work to be done before attosecond spectroscopy becomes a standard laboratory tool. For example, the best approaches to studying various dynamics have yet to be established, and new ideas for time-resolved spectroscopic configurations, yielding different information than standard photoelectron spectroscopy, are beginning to be explored.

The UK has set up an attosecond consortium, funded by the *Research Councils UK, Basic Technology* grant, consisting of several UK institutes with the aim of developing young researchers in this emerging field and to advance the field towards the study of attosecond dynamics in multiple traditional fields such as chemistry, biology, materials science, surface science and physics. The consortium consists of seven universities and research institutes (Imperial and University Colleges, the universities of Oxford, Reading,

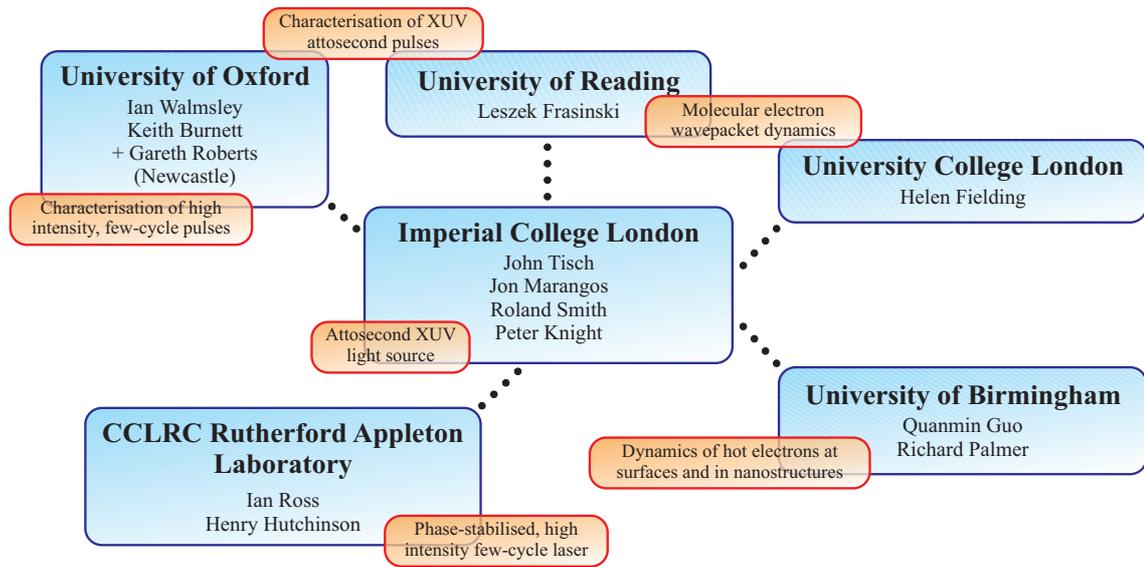


Figure 1.2: The various roles of each institute within the UK attosecond consortium.

Birmingham and Newcastle and CCLRC Rutherford Appleton Laboratory) to develop the source, metrology and applications to attosecond spectroscopy of atoms and molecules (figure 1.2).

The role of Oxford within the consortium is to develop novel methods of metrology for attosecond sources. Although there have been great advances in both the metrology and development of attosecond sources, they have only just started to be used in applications of attosecond pulses. It has been the goal of Oxford and the consortium to simplify the measurement procedure and thus open up attoscience to the general scientific community, rather than restricting it to the few experienced researchers in this field.

Metrology is crucial for attoscience in the development of both sources and applications, it being necessary to optimise the generation as well as providing new experimental capabilities. Accurate metrology is important for applications because the amount of information on the dynamics of the system that can be extracted from measurements is limited by the knowledge of the pulses used to probe the system.

A secondary aspect of Oxford's role is to develop metrology for the few-cycle drive laser used to generate attosecond pulses. It is critical to have an extremely accurate and

robust method that allows online optimisation of the system, and to have data to be used in simulations of the attosecond generation process. This in itself is still an interesting and active area of research, for example relating the microscopic and macroscopic effects during the generation and to test and develop theories beyond generation in simple atoms.

1.3 Metrology of ultrashort pulses and high harmonic generation

Femtosecond pulses are becoming more common in many laboratories worldwide due to the improvements in both sources and metrology. For example, it is possible to generate femtosecond pulses ranging from the near infrared (NIR), through the visible (Vis.) and into the ultraviolet (UV) spectral region with durations as short as a few optical periods. It is also possible to arbitrarily shape the temporal electric field of these pulses. For these wavelengths, it is not possible to measure the time evolution of the electric field due to the high carrier frequency. Not only that, but it is not even possible to measure the temporal intensity directly because the pulse duration is far too short (e.g. electronic devices typically have response times of nanoseconds or more, compared to pulses with sub-picosecond durations). However, it is possible to completely characterise the pulse by measuring both the spectral intensity and spectral phase, which can be done via the use of slow detectors (i.e. ones with a response time many orders of magnitude longer than the pulse duration being measured). Many different techniques exist, and these fall into two broad categories: (1) referenced based and (2) self-referencing.^[16]

If there exists a well-characterised pulse, then it is possible to use this as a reference for characterising an unknown pulse. In the case when the spectrum of the reference pulse fully encompasses that of the unknown pulse, then it is possible to utilise linear techniques, offering the advantages of the highest available sensitivity and the ability to measure complex pulse shapes. In the case when the spectrum of the known pulse does not fully cover that of the unknown pulse, or even not at all, then it is necessary to use nonlinear methods, reducing the sensitivity of these techniques.

Often a known reference pulse is not available, which is obviously the case for the shortest pulses. In this case, it is necessary to perform self-referencing measurements, i.e. the pulse to be measured is referenced to itself. The details of these methods have been formalised by Walmsley and Wong for use with square-law integrating (i.e. slow) detectors.^[17,18] Self-referenced techniques require the use of a time-nonstationary filter, i.e. a device whose response depends on the arrival time of the incident pulse. Currently these filters are implemented through nonlinear optics. Thus they have lower sensitivity and require a medium that is nonlinear at the pulse wavelength and intensity. In addition, these methods tend not to perform adequately for complicated pulses (e.g. pulses with well-separated spectral regions, or trains of pulses).

At present, it is not possible to produce *and* measure attosecond pulses routinely and easily. This is mainly because the bandwidth required is extraordinarily large and the mean wavelength of the pulse is in the XUV spectral region (10–100 nm). The first restriction arises because the spectral bandwidth required to support a pulse is inversely proportional to its duration. This presents problems because no medium is currently known to have gain covering this bandwidth. In addition, current optics typically are not transparent and have significant dispersion over such a spectral range. The second condition arises because the shortest pulse duration possible for propagating waves is equal to the period of the carrier wave. It is extremely difficult to generate such wavelengths out of a laser directly and there are only limited optics that can be used to manipulate these pulses. These two requirements (central wavelength and bandwidth) pose many practical and physical limitations on the generation, manipulation and detection of attosecond pulses. Thus it is necessary to have as wide array of techniques and geometries as possible that can either complement each other, or provide solutions in different parameter spaces.

The primary method of creating attosecond pulses is to generate XUV light via the interaction of an intense, phase-stabilised few-cycle laser pulse with a gas — a process called high harmonic generation. Although characterisation techniques exist for measur-

ing these ultrafast driving pulses, for full-width-half-maximum pulse durations less than 10 fs, corresponding to a few optical periods, their accuracy is reduced compared to longer duration pulses due to certain features associated with few-cycle pulses. For example, sub-10 fs pulses tend to exhibit one or more of the following: (1) the spectrum is ultra-broadband (e.g. spans over an octave), (2) the spectrum is highly modulated, (3) the pulse exhibits space-time coupling and (4) it is desirable to have single-shot acquisition (in order to characterise pulse-to-pulse fluctuations). The first half of this thesis formulates methods of characterising ultrashort laser pulses and describes a characterisation technique based on spectral shearing interferometry that I have developed. I then show how this method can be used to measure sub-10 fs near-infrared pulses, overcoming the problems with conventional techniques due to the properties of these pulses as stated above. I then compare the reconstructed electric field to reconstructions from other ‘conventional’ techniques, highlighting the features that pose difficulties for these conventional methods. Accurate metrology of these pulses is critical for developing an understanding of attosecond sources based on high harmonic generation due to the high nonlinearity of the process with respect to the driving field.

There are currently two different schemes for the spectral characterisation of XUV pulses: (1) the detection of the velocity distribution of electrons emitted through photoionisation from a gaseous or solid target and (2) the direct detection of the photon energy via an XUV spectrometer. Previous metrology techniques for XUV attosecond pulses have been based around the former detection scheme. Photoelectron spectrometry (PES) tends to have very low efficiency, resulting in long acquisition times. Also, the current PES schemes require excessive amounts of data acquisition over several hours. Thus extremely stable laser systems are required to reduce errors in the reconstruction. Photoelectron spectrometers are also expensive and quite complex, requiring expertise to operate them. Direct detection of the photon energy through an XUV spectrometer generally has much higher efficiency, and can measure a two-dimensional parameter space (e.g. frequency and

one-dimensional spatial co-ordinate), thus possibly enabling single-shot measurements of the pulse's temporal structure and space-time coupling.

The second half of this thesis is concerned with metrology of extreme ultraviolet, attosecond pulses from high harmonic generation via direct XUV interferometry. The developments that I have made to enable self-referencing spectral interferometry of XUV attosecond pulses by controlling the high harmonic generation process will be presented. Demonstration of lateral shearing interferometry, a linear method based on spectral interferometry, for measuring the complete frequency-dependent spatial phase and a method of performing spectral shearing interferometry for measuring the temporal electric field will be discussed.

1.4 Thesis outline

Chapter 2 will outline the mathematical definition of electromagnetic pulses, including space-time coupling. A few examples will be used to put this into a physical context. This chapter will also outline the conventional methods used for femtosecond metrology.

Chapter 3 discusses the problem of characterising few-cycle near infrared pulses and highlights the shortcomings of conventional techniques. Then I describe the development of a novel metrology scheme that is a variant of one particular technique based on spectral shearing interferometry and demonstrate it for few-cycle pulses.

Chapter 4 describes the process of high harmonic generation theoretically. I also provide a brief overview of current methods for characterisation of the generated field.

Chapter 5 outlines the experimental setup and measurement results of spectrally-resolved wavefront measurements of high harmonics and more specifically a method of measuring their space-time coupling via spectrally-resolved lateral shearing interferometry. This method can be considered as a simple interferometric measurement that reconstructs

the complete phase over the whole frequency–1D spatial co-ordinate, thus identifying the problems associated with self-referencing spectral interferometry for XUV pulses from HHG.

Chapter 6 describes different methods and geometries for performing spectral shearing interferometry of XUV pulses from high harmonic generation. Simulations of spectrally-sheared harmonics will be presented to provide the theoretical framework for the feasibility of performing spectral shearing interferometry in the XUV region. Different methods relating to generating the spectral shear will be presented along with the applicability of each technique to different operating regimes.

Chapter 7 will conclude the thesis with a summary of what has been achieved and what can be expected in the future.

Author’s contribution

The work and ideas presented in this thesis, unless specifically stated, were carried out by the author at the university of Oxford under the supervision of Prof. I.A. Walmsley. However, as Oxford does not have any facilities for the experimental work presented in this thesis, this was carried out in collaboration with a number of different universities. Oxford is a node in the European funded XTRA network for the training of young researchers. Many of the experiments on high harmonic generation were performed in collaboration with other nodes within this network at their laboratories, specifically the CELIA laser facility at the university of Bordeaux, France and the U. Keller group at ETH, Zürich, although no results from these groups have been presented here.

The results presented in chapter 3 were performed at the Max-Born-Institute, Berlin, Germany in collaboration with Dr. G. Steinmeyer and PhD student G. Stibenz. The author designed and built the SEA-SPIDER presented in this chapter, analysed the data and wrote the paper on “Sub-10 fs pulse characterisation using spectrally encoded arrangement

for spectral phase interferometry for direct electric-field reconstruction.”^[19] The experiments on IFROG and SPIDER and their analysis presented in this chapter were performed by G. Stibenz.^[20] The discussions on few-cycle metrology presented within this chapter are from the author.

The experimental results presented in chapters 5 and 6 were performed using the attosecond laser facility at Imperial College, London under the supervision of Prof. J. Marangos and Dr. J. W. G. Tisch in conjunction with the post-doctoral researcher J. Robinson and PhD student C. Haworth from Imperial college and PhD student T. Witting from the university of Oxford. The author was responsible for the planning, implementation and analysis of these experiments, although the acquisition and running of the experiments were performed in conjunction with the previously name researchers. The simulations presented in chapter 6 were performed by E. Cormier from the university of Bordeaux on the M3PEC pole for parallel computing at Université Bordeaux 1, and I was involved in the analysis of the data and the writing of the papers^[21–23].

The author would also like to acknowledge stimulating discussions with Dr. A. Monmayrant, Dr. L. Corner from the university of Oxford and Dr. R. Smith from Imperial College, London in relation to the work presented in this thesis.

Metrology of ultrashort pulses

In time-resolved spectroscopy experiments, the electromagnetic field is a fundamental entity in Maxwell's theory and the amount of knowledge of this limits the amount of information one can extract from such an experiment. Typically the intensity of the radiation is small enough that magnetic field effects can be neglected (as will be done throughout this thesis) and thus it is the time-variation of the electric field at some point or set of points in space that one wants to characterise.

Ideally one would like to measure the electric field at the location of the experiment (i.e. *in situ*). However, this is typically not possible. Thus the electric field is measured elsewhere and the field is calculated numerically at the desired location. The use of ultrashort pulses means that it is not possible to measure the time-variation of the intensity envelope directly because the pulse duration is shorter than currently available detector response times. For example, electronic detectors have response times on the order of picoseconds to nanoseconds, whereas I deal with pulses with durations less than 100 fs and carrier frequencies with a period of several femtoseconds (i.e. in the optical regime).

In order to gain insight into how it is possible to fully characterise the electric field at the site of the experiment, I outline the mathematical representation of electromagnetic pulses. This will show how one can determine the electric field in both space and time at an arbitrary location. I then outline a few of the most common techniques used for

measuring the spatial and temporal properties of electromagnetic pulses, and identify some advantages and disadvantages of each.

2.1 Mathematical description of electromagnetic pulses

The properties of all electromagnetic fields are governed by Maxwell's equations (appendix A), and thus ultimately it is the solutions to these equations that one seeks. Electromagnetic radiation allows energy to propagate in space via the electromagnetic field and is described by the electromagnetic wave equation derived directly from Maxwell's equations. The wave equation in vacuum is

$$\left(\nabla^2 - \frac{1}{c^2} \frac{\partial^2}{\partial t^2}\right) \boldsymbol{\xi} = 0, \quad (2.1)$$

where c is the speed of light in vacuum and $\boldsymbol{\xi}$ is the electric field vector. Equation (2.1) represents three decoupled equations that describe the evolution of the individual components of the electric field vector in Cartesian co-ordinates. Gauss' law in vacuum, $\nabla \cdot \boldsymbol{\xi} = 0$, results in a spatially bounded field containing a longitudinal electric-field component (i.e. in the propagation direction).^[24,25] However, this only becomes significant for large numerical apertures ($\gtrsim 0.6$), or equivalently when the spatial confinement is small (on the order of a wavelength) and thus can usually be neglected in many laboratory situations.^[25,26] Hence, for *linearly* polarised light, one can consider the scalar value, ξ , which represents the electric field strength along the polarisation co-ordinate. Individual polarisation components can be considered independently in the case of elliptically polarised light. Although multiple solutions to equation (2.1) exist, I only consider, for reasons that will become apparent later, plane wave solutions of the form

$$\xi(t, \mathbf{r}) = \xi_0 \cos(\mathbf{k} \cdot \mathbf{r} - \omega t + \tilde{\phi}(\omega, \mathbf{k})), \quad (2.2)$$

where ω is the angular frequency (in radians per second); $\mathbf{k} = k_x \hat{\mathbf{x}} + k_y \hat{\mathbf{y}} + k_z \hat{\mathbf{z}}$ is the wave vector (in radians per metre); $\mathbf{r} = (x, y, z)$ is the position co-ordinate (in metres); and $\tilde{\phi}(\omega, \mathbf{k})$ is a phase angle (units of radians), the importance of that will be discussed later. The angular frequency and wave vector are constrained by the dispersion relation in vacuum

$$\mathbf{k}^2 = \frac{\omega^2}{c^2}, \quad (2.3)$$

where c is the speed of light in vacuum.

Due to the linearity of the wave equation, any linear superposition of plane waves with different wave vectors or (positive semi-infinite) frequencies is also a solution. In fact, equation (2.2) cannot represent a physical electromagnetic wave because the electric field has infinite extent in both space and time. Thus, all physical radiation has to be a superposition of an infinite set of plane waves

$$E(t, \mathbf{r}) = \int_{-\infty}^{\infty} d\omega \int_{-\infty}^{\infty} d\mathbf{k} |\tilde{E}(\omega, \mathbf{k})| e^{i\tilde{\phi}(\omega, \mathbf{k})} e^{-i\omega t} e^{i\mathbf{k} \cdot \mathbf{r}}, \quad (2.4)$$

where for mathematical simplicity I have introduced the analytic electric field, $E(t, \mathbf{r}) = |E(t, \mathbf{r})| e^{i\phi(t, \mathbf{r})}$ (see appendix A.3 for further details). The physical electric field is twice the real part of the analytic field

$$\boldsymbol{\xi}(t, \mathbf{r}) = 2\Re[\mathbf{E}(t, \mathbf{r})]. \quad (2.5)$$

From equation (2.4), it is clear that any pulse can be represented by its complex valued analytical frequency–angular spectrum, $\tilde{E}(\omega, \mathbf{k}) = |\tilde{E}(\omega, \mathbf{k})| e^{i\tilde{\phi}(\omega, \mathbf{k})}$. The temporal evolution of the electric field at a given point \mathbf{r}_0 is found by the Fourier transform of the

frequency spectrum at that point

$$\begin{aligned} E(t, \mathbf{r}_0) &= \int_{-\infty}^{\infty} d\omega \tilde{E}(\omega, \mathbf{r}_0) e^{-i\omega t} \\ &= \mathfrak{F} \left\{ \tilde{E}(\omega, \mathbf{r}_0); \omega \rightarrow t \right\}. \end{aligned} \quad (2.6)$$

The spatial distribution in a plane at any given position z_0 and time t_0 is given by the (scaled) two dimensional inverse Fourier transform of the angular distribution at that position and time

$$\begin{aligned} E(t_0, x, y, z_0) &= \int_{-\infty}^{\infty} dk_x \int_{-\infty}^{\infty} dk_y \tilde{E}(t_0, k_x, k_y, z_0) e^{ik_x x} e^{ik_y y} \\ &= (2\pi)^2 \mathfrak{F}_2^{-1} \left\{ \tilde{E}(t_0, k_x, k_y, z_0); k_x \rightarrow x, k_y \rightarrow y \right\}. \end{aligned} \quad (2.7)$$

Note that I have conformed to the mathematical convention for the definitions of the Fourier transforms, which have the opposite signs to those given in conventional optics (details are given in appendix A.4). I have also assumed, without loss of generality, that the detector's active area is placed parallel to the x-y plane.

As the electric field is represented by a complex quantity, the electric field distribution is not only dependent on the spectral intensity, but also the spectral phase. Hence it is necessary to be able to measure both the spectral intensity and spectral phase in order to fully characterise the temporal evolution of the field. The spectral phase has physical significance, which will be discussed in the following sections.

2.1.1 Time-frequency representations of ultrashort pulses

Time-frequency representations of ultrashort electromagnetic pulses are useful for visualising the pulse structure, as well as for describing the action of optical devices on the pulse. Two distributions that are very often encountered in ultrashort optics are the Wigner distribution and the spectrograph, which are summarised below.

Wigner representation of ultrashort pulses

The Wigner distribution for electromagnetic pulses is a real function in the time-frequency phase space. It is useful in describing the action of devices and nonlinear propagation of ultrashort pulses. Due to the distribution containing both temporal and spectral information, it is commonly referred to as the *chronocyclic representation*. For the complete formalism for this representation, refer to Paye.^[27] The Wigner function, $W(E; t, \omega)$ of the complex electric field $E(t)$ is given by the scaled inverse Fourier transform of the mutual intensity

$$W(E; t, \omega) = 2\pi\mathfrak{F}^{-1} \left\{ E \left(t + \frac{t'}{2} \right) E^* \left(t - \frac{t'}{2} \right); t' \rightarrow \omega \right\}. \quad (2.8)$$

The Wigner distribution can also be calculated in the spectral domain from the complex spectrum $\tilde{E}(\omega)$

$$W(E; t, \omega) = 2\pi\mathfrak{F} \left\{ \tilde{E} \left(\omega + \frac{\omega'}{2} \right) \tilde{E}^* \left(\omega - \frac{\omega'}{2} \right); \omega' \rightarrow t \right\}. \quad (2.9)$$

The factors of 2π arise from the definitions used by Paye ($\tilde{\varepsilon}(\omega) = 2\pi\tilde{E}(\omega)$ where $\tilde{\varepsilon}(\omega)$ is the quantity used by Paye).^[27]

The Wigner distribution is always real, but not necessarily positive. The marginals of the Wigner distribution give the spectral and temporal intensities

$$\begin{aligned} \tilde{I}(\omega) &= \int_{-\infty}^{\infty} dt W(E; t, \omega) \\ &= 4\pi^2 |\tilde{E}(\omega)|^2. \end{aligned} \quad (2.10)$$

$$\begin{aligned} I(t) &= \frac{1}{2\pi} \int_{-\infty}^{\infty} d\omega W(E; t, \omega) \\ &= |E(t)|^2. \end{aligned} \quad (2.11)$$

The integral over the whole distribution phase space yields the pulse energy

$$E_n = \frac{1}{2\pi} \int_{-\infty}^{\infty} dt \int_{-\infty}^{\infty} d\omega W(E; t, \omega). \quad (2.12)$$

It is important to note that both marginals and the energy will always be positive values even though the Wigner distribution may contain negative values. Although the Wigner distribution contains all information about the pulse, care is needed in its interpretation,

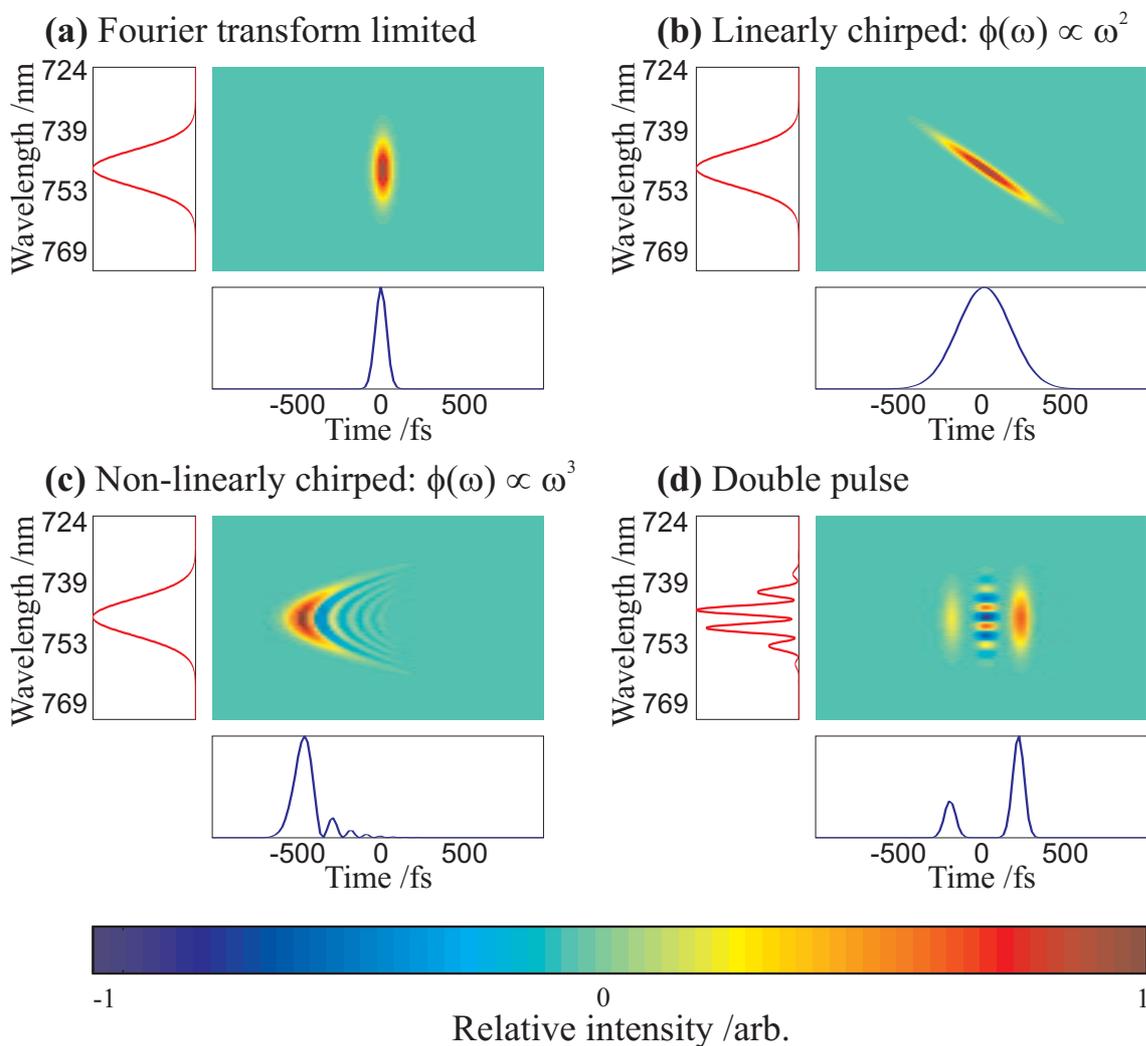


Figure 2.1: Wigner representation and marginals of different optical pulses. Pulses (a)-(c) have the same spectral intensity, but with flat, quadratic and cubic spectral phase respectively. (d) Represents a Fourier limited double pulse train. Each plot consists of the Wigner distribution (image), spectrum (left marginal, red curve) and temporal intensity (bottom marginal, blue curve).

in the sense that it is not a physical representation of the pulse. For example, a slice through the Wigner distribution at any given time is not equal to the spectrum at that time. In addition, the distribution may contain structure that is non-physical (e.g. fast oscillations), which are not seen in the marginals, e.g. figure 2.1 (c).

The use of the Wigner distribution allows one to use many mathematical properties, typically derived in quantum mechanics. The distribution itself is very intuitive, showing the carrier frequency dependence of the pulse with time. Figure 2.1 illustrates four different Wigner distributions of the same spectrum but with different phase. One very important use of the Wigner distribution for optical pulses is in tomography, which can be used to reconstruct the complex electric field via manipulations of the Wigner distribution (section 2.5.4), or for nonlinear propagation.

Spectrographic representation of ultrashort pulses

Although the Wigner distribution is a very useful mathematical tool in ultrashort optics, many pulse shapes lead to complicated distributions that are not so intuitive. For example, the distribution may contain negative regions or interference-like artifacts as shown in figures 2.1 (c) and (d). The spectrographic representation of ultrashort pulses provides a more intuitive picture of the pulse in these cases. For simple pulse shapes, the spectrographic and chronocyclic representations may be identical. The spectrographic representation, $S(E, G; t, \omega)$ is formed from the time-gated spectrum of the pulse as a function of the delay, τ between the pulse and the gate

$$S(E, G; \tau, \omega) = |\mathfrak{F}^{-1}\{E(t)G(t-\tau); t \rightarrow \omega\}|^2, \quad (2.13)$$

where $E(t)$ is the complex pulse and $G(t)$ is the gate function. In general, the gate may also be a complex function, but is typically a (Gaussian) amplitude filter only. From equation (2.13), it is clear that the spectrographic representation is formed from the convolution of the pulse and gate functions, and thus requires deconvolution of the gate function

same pulses used in figure 2.1. In many situations, it *is* possible to recover this phase information provided the gate function is known. However, certain pulse shapes lead to identical spectrograms. For example, any phase offset between the two pulses in the pulse train shown in figure 2.2 will give the same spectrograph.

2.2 Propagation of electromagnetic pulses

In the previous section I outlined the model that I use to describe ultrashort pulses. Using equation (2.7), I will show how one can describe the propagation of a beam in space and how common linear optics are related to this description of the beam. Through this description, I highlight the importance the spatial phase plays in beam propagation. I then turn to a discussion of how a pulse propagates in a medium, using the same formalism presented in the previous section. This will then allow a discussion of where a spectral phase can arise and how this plays a role in the temporal evolution of optical pulses via equation (2.6).

The simplest description of the propagation of light occurs in *geometric optics*, which considers the light to propagate along trajectories with a defined direction and starting point. These trajectories are known as optical rays. The propagation is then defined by the reflection and refraction of these rays at surfaces. Such a treatment is sufficient to describe optical imaging and chromatic dispersion but fails to incorporate other important phenomena such as diffraction and interference. Another simplification, often made in simple geometric optics, is the *paraxial approximation*, which assumes that all rays travel at small angles, θ , to the chief ray of a bundle, so that the approximation $\tan \theta \simeq \sin \theta \simeq \theta$ can be applied. This is referred to as *paraxial optics*.

Gaussian optics is an expansion of paraxial optics that partially takes into account diffraction of optical beams. In the first order treatment, the beam is assumed to have a

Gaussian spatial profile

$$E(r, z) = E_0 \frac{w_0}{w(z)} \exp \left[- \left(\frac{r}{w(z)} \right)^2 \right] \exp \left[-i \left(kz + k \frac{r^2}{2R(z)} - \zeta(z) \right) \right], \quad (2.14)$$

where

$$w(z) = w_0 \sqrt{1 + \left(\frac{z}{z_R} \right)^2} \quad (2.15)$$

is the beam half width at $E(r, z) = E_0 e^{-1}$ (which is equivalent to the half-width at $I(r, z) = I_0 e^{-2}$ for the intensity); $w_0 = w(z=0)$ is the minimum value for the beam half-width and is known as the *beam waist*.

$$R(z) = z \left[1 + \left(\frac{z_R}{z} \right)^2 \right] \quad (2.16)$$

is the radius of curvature due to propagation and accounts for diffraction of the beam.

$$\zeta(z) = \tan^{-1} \left(\frac{z}{z_R} \right) \quad (2.17)$$

is the Gouy phase^[28–30].

$$z_R = \frac{\pi w_0^2}{\lambda} = \frac{k w_0^2}{2} \quad (2.18)$$

is the Rayleigh range, which is the distance from the position of the beam waist to the point at which the beam width increase by a factor of $\sqrt{2}$, i.e. $w(\pm z_R) = w_0 \sqrt{2}$. Equation (2.14) is a solution to equation (2.1) in the paraxial limit which assumes

$$2ik \frac{\partial}{\partial z} \tilde{E}(\omega, \mathbf{r}) \gg \frac{\partial^2}{\partial z^2} \tilde{E}(\omega, \mathbf{r}). \quad (2.19)$$

Figure 2.3 shows the propagation of a Gaussian beam through a focus, highlighting the beam radius, spatial phase, beam wavefronts and wave vectors.

Gaussian beams arise naturally in the treatment of cavities and are thus extremely useful in laser physics. Equation (2.14) refers to the lowest order transverse electric-

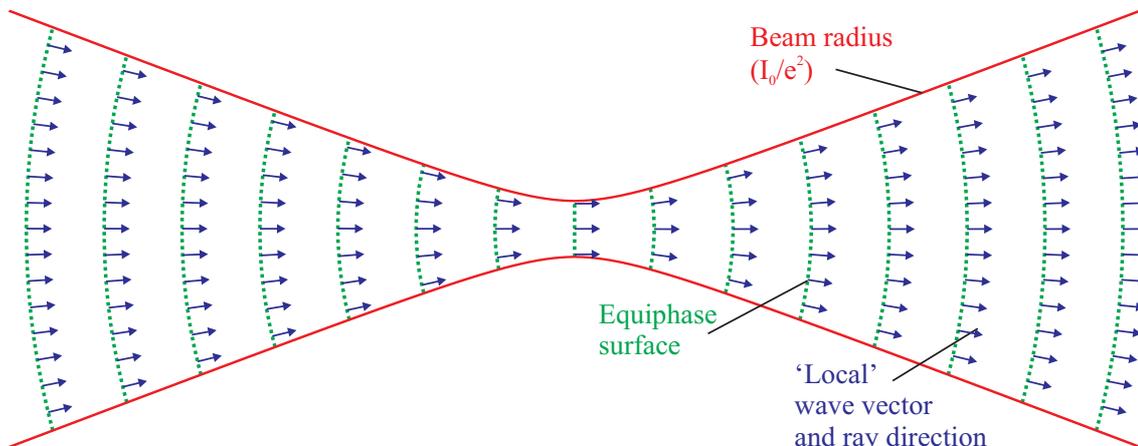


Figure 2.3: Gaussian beam propagation and spatial phase. A Gaussian beam is propagated according to either equation (2.14) or equation (2.23). The beam radius at $1/e^2$ of the maximum intensity along the propagation direction is drawn in red, the lines of equal phase are drawn as dotted green and the local wave vectors (equal to the gradient of the phase and direction of the optical rays) are drawn as blue arrows.

magnetic (TEM_{00}) field mode. However, higher-order modes can be obtained via the use of Hermite Gaussian beams, allowing for the propagation of more complex beam profiles.^[31] The use of Gaussian beams allows one to use matrix optics, which determines the propagation of a beam through linear optical elements via a set of simple matrix operations. The form of the matrices used in Gaussian optics are exactly the same as for geometrical optics.^[32]

Due to the nature of generating extremely short pulses, down to a few optical periods, the spatial profile of such pulses do not typically have a Gaussian envelope and can exhibit spatial phase not in the form given in equation (2.14). Thus it is necessary to consider *physical optics*, which models the propagation exactly. Such a treatment allows one to obtain complete information on effects such as interference and polarisation as well as aberrations and more complex phenomena, e.g. space-time coupling (STC). The full electric field must be modelled, including a phase angle, $\phi(\mathbf{r})$. It is the propagation of complex wavefronts through optical systems that determine the pulse shape, usually in the form of the electric field derived in section 2.1 and more specifically equation (2.7) under the scalar approximation.

For any given frequency in a pulse, it is possible to calculate its propagation according to equation (2.4). However, the wave vector is constrained by the free-space dispersion relation

$$k = \frac{\mathbf{k}}{\hat{\mathbf{k}}} = \frac{\omega}{c} \quad (2.20)$$

and the wave vector parallel to the propagation direction is constrained by

$$k_z = \sqrt{k^2 - k_x^2 - k_y^2} \quad (2.21)$$

$$\simeq k - \frac{k_x^2 + k_y^2}{2k} \quad \text{for } k^2 \gg (k_x^2 + k_y^2). \quad (2.22)$$

Expanding equation (2.21) to only two orders, (2.22), is equivalent to the paraxial approximation. Thus, for a given frequency, the electric field strength at any point in space is given by

$$E(\omega, x, y, z) = 4\pi^2 \mathfrak{F}_2^{-1} \left\{ |\tilde{E}(\omega, k_x, k_y, z_0)| e^{i\tilde{\phi}(\omega, k_x, k_y, z_0)} e^{i\sqrt{k^2 - k_x^2 - k_y^2}z}; k_x \rightarrow x, k_y \rightarrow y \right\}. \quad (2.23)$$

Hence, diffraction is described by an elliptical phase in the angular distribution, reducing to a quadratic dependence in both the angular and spatial distributions in the paraxial or far-field approximation:

$$\tilde{\phi}_{\text{diff}}(\omega, \mathbf{k}) \simeq -\frac{L}{2k} (k_x^2 + k_y^2) \quad (2.24)$$

$$\phi_{\text{diff}}(\omega, \mathbf{r}) \simeq \frac{k}{2L} (x^2 + y^2), \quad (2.25)$$

where L is the propagation distance; $\tilde{\phi}_{\text{diff}}(\omega, \mathbf{k})$ is the angular phase due to diffraction in the paraxial limit and in the pulse frame; and $\phi_{\text{diff}}(\omega, \mathbf{r})$ is the spatial phase due to diffraction in the pulse frame and in the far-field limit such that $L \gg \Delta x^2/\lambda$, where Δx is the spatial extent of the source.

In order to have complete knowledge of the beam wavefront, it is necessary to measure both the spatial amplitude profile and the spatial phase distribution. As an example,

consider the effects the distorted surface of a swimming pool has on light propagating through it. Even though the amount of light transmitted through the surface remains approximately constant, shadows are seen at the bottom of the swimming pool. This is due to the phase accumulated as light passes through the uneven surface and then propagated to the ‘far-field’ (i.e. bottom of the swimming pool). Thus it is clear that phase distortions can give rise to significantly different intensity distributions, especially after propagation. These effects, and the physical interpretation of the phase, will now be considered in more detail.

Physically, the spatial gradient of the spatial phase at a point is normal to the local wavefront at that point — the blue arrows in figure 2.3 shows the local wave vectors (i.e. the local wavefront normal). Tilt (i.e. propagation at an angle to the optical axis) is defined by a linear spatial phase

$$E(x, y) = |E(x, y)|e^{ikx \sin \theta}$$

$$\tilde{E}(k_x, k_y) \rightarrow \tilde{E}(k_x - k \sin \theta, k_y), \quad (2.26)$$

where θ is the angle that the plane-wave propagation vector makes with the z-axis in the x-z plane. Note that in the case of a large obliquity angle, the spatial amplitude will be different from an identical beam at normal incidence due to a geometrical stretch from non-normal incidence, and possibly diffraction in the case of a small beam. However, the beam can still be described in the form of (2.26). As it is normally necessary to only consider near-normal incidence, the spatial amplitude is approximately constant as a function of the incidence angle with respect to the detector.

The effect of a focusing optic is modelled by a nonlinear spatial phase. For example, a parabolic mirror would induce a quadratic spatial phase:

$$\phi_{\text{pm}}(\omega, x, y) = -k \left[\frac{x^2 + y^2}{2f} + x \sin 2\theta \right], \quad (2.27)$$

where f is the focal length of the mirror and θ is the angle of incidence in the x - z plane (assuming no tilt in the y - z plane). A spherical mirror would induce a spherical phase:

$$\phi_{\text{sm}}(\omega, x, y) = -k \left(f - \sqrt{f^2 - (x^2 + y^2)} + x \sin 2\theta \right). \quad (2.28)$$

It is also necessary to take into account the change in propagation direction:

$$\tilde{\xi}(\omega, \mathbf{k}) \cos \left[\mathbf{k} \cdot \mathbf{r} - \omega t + \tilde{\phi}(\omega, \mathbf{k}) \right] \leftrightarrow \tilde{\xi}(\omega, \mathbf{k}) \cos \left[\mathbf{k} \cdot \mathbf{r} + \omega t + \tilde{\phi}(\omega, \mathbf{k}) \right] \quad (2.29)$$

and the change in handedness (right handed circular \leftrightarrow left handed circular) upon reflection. In the paraxial approximation, equation (2.28) reduces to a parabolic dependence.

A *thin* lens can also be modelled by a phase function (neglecting losses due to surface reflections) that is given by

$$\phi_1(\omega, x, y) = k [\Delta_0 + (n - 1) \Delta(x, y)], \quad (2.30)$$

where Δ_0 is the thickness of the lens at the centre, k is the magnitude of the wave vector in vacuum, n is the refractive index of the lens (see section 2.2.1) and $\Delta(x, y)$ is the thickness function of the lens. For a spherical lens with radii of curvatures r_1 for the front surface and r_2 for the back surface, where r_i is positive if the surface is convex relative to the propagation direction. The thickness function is given by

$$\Delta(x, y) = \Delta_0 - r_1 \left(1 - \sqrt{1 - \frac{x^2 + y^2}{r_1^2}} \right) + r_2 \left(1 - \sqrt{1 - \frac{x^2 + y^2}{r_2^2}} \right). \quad (2.31)$$

The focal length of the lens is given by

$$\frac{1}{f} = (n - 1) \left(\frac{1}{r_1} - \frac{1}{r_2} \right). \quad (2.32)$$

It is clear that the spatial phase due to diffraction, (2.25), in the far-field limit can be

fully compensated by the quadratic phase introduced by a parabolic mirror, or spherical focusing optic in the paraxial limit. Wavefront aberrations can be modelled by a polynomial spatial phase of higher than second order, and can be compensated to some degree via a deformable mirror.

2.2.1 Propagation of electromagnetic fields in dielectrics

When an electromagnetic field propagates through a medium, the medium reacts to it. Considering the electric field only (i.e. assuming a negligible effect from the magnetic field), the response of the medium is represented by its polarisation, \mathbf{P} . The polarisation is the sum of the induced electric dipole moments of the atoms inside the material due to the applied oscillating electric field. The strength of the dipole can be modelled by an electron oscillating in an atomic or lattice potential. The potential can be described by a Taylor expansion; the order of expansion being dependent on the oscillation amplitude, and therefore on the electric field strength. For small electric field strengths, the oscillations are small and so the potential can be modelled by an harmonic potential, resulting in a linear dependence of the polarisation with electric field strength. As the electric field strength increases, the potential becomes anharmonic, resulting in a nonlinear dependence of the polarisation with the electric field. The polarisation can be described by a Taylor expansion

$$\mathbf{P} = \varepsilon_0 \left(\chi^{(1)} \boldsymbol{\xi} + \chi^{(2)} \boldsymbol{\xi}^2 + \chi^{(3)} \boldsymbol{\xi}^3 + \dots \right), \quad (2.33)$$

where $\chi^{(n)}$ is the n^{th} order susceptibility of the medium and ε_0 is the permittivity of free space. In an anisotropic medium, the polarisation and electric field are not necessarily aligned. In general therefore, the electric susceptibility is a tensor, resulting in birefringence for example; further details can be found in Born and Wolf (§15).^[33] Comparison of the polarisation (2.33) with the electric displacement, $\mathbf{D} = \varepsilon_0 \mathbf{E} + \mathbf{P}$, and material equation, $\mathbf{D} = \varepsilon \mathbf{E}$ (appendix A) results in a material permittivity, $\varepsilon = \varepsilon_r \varepsilon_0 = \varepsilon_0 (1 + \chi)$, where ε_r is the relative permittivity.

Incorporating the polarisation of the material into Maxwell's equations results in a modified wave equation for a homogeneous isotropic medium given by

$$\left(\nabla^2 - \frac{1}{c^2} \frac{\partial^2}{\partial t^2}\right) \boldsymbol{\xi} = \frac{1}{\varepsilon_0 c^2} \frac{\partial^2 \mathbf{P}}{\partial t^2}. \quad (2.34)$$

At low intensities, such that the polarisation can be reduced to the linear case, the wave equation is also linear. This results in plane-wave solutions, equation (2.2), as before but with a modified dispersion relation given by

$$\mathbf{k}^2 = \left(\frac{n\omega}{c}\right)^2 \quad (2.35)$$

where $n = \sqrt{\varepsilon_r} = \sqrt{1 + \chi} = c/v$ is the refractive index of the medium (in the case of negligible magnetisation), and describes the phase velocity, v , relative to the speed of light in a vacuum, c . In general, the susceptibility and hence refractive index is frequency dependent, leading to *chromatic dispersion*.

At higher electric field intensities, such that the polarisation, \mathbf{P} , varies nonlinearly with the electric field, the wave equation is nonlinear. In the perturbative regime — the electric field strength is relatively small compared to the binding field of the atom (e.g. an electric field strength which results in an intensity $I \ll 10^{15} \text{ Wcm}^{-2}$ for linearly polarised light) — the nonlinear wave equation leads to low-order harmonics of the fundamental frequency, self focusing and Raman scattering. Nonlinear optics can both be detrimental to, and utilised in, optical systems. A detailed discussion of nonlinear optics can be found in Boyd.^[34]

If the electric field strength is large enough that it can exert a force comparable to the atomic binding force (so that $I \sim 10^{15} \text{ Wcm}^{-2}$ for linearly polarised light), then it is possible for the uppermost bound electrons to undergo tunnelling or over-the-barrier ionisation. The dynamics of such a system is the subject of strong field physics. High harmonic generation (HHG) is a phenomenon that results from intense laser-matter interactions and can

be used as a method of generating short-wavelength ultrashort bursts of electromagnetic radiation. The details of HHG will be discussed in more detail in chapter 4.

At extremely high field intensities ($I \gtrsim 10^{18} \text{ Wcm}^{-2}$), the ionised electrons can be accelerated to relativistic speeds, leading to relativistic nonlinear optics. Many of the assumptions used in the strong field regime breakdown. A review of applications for ultra high-intensity phenomena is given by Ledingham *et al.*^[35]

2.2.2 Chromatic Dispersion

Chromatic dispersion arises due to a frequency dependent phase velocity in the case of material or waveguide dispersion, from a frequency dependent optical path length in the case of pulse shapers (e.g. stretchers and compressors), via interference (e.g. Gires–Tournois interferometer^[36]) or even from nonlinear propagation.^[37]

In the case of material dispersion, the refractive index of the material will be a function of frequency, and is typically described by a formula known as the Sellmeier equation:^[38]

$$n(\lambda) = \sqrt{1 + \sum_j \frac{B_j \lambda^2}{\lambda^2 - C_j}}. \quad (2.36)$$

The formula is derived from the summation of different absorption resonances with a coupling strength B_j at a wavelength $\sqrt{C_j}$. The coefficients are found empirically by fitting the formula to measured data over a given range. The smooth function allows differentiation to obtain dispersion coefficients. However, the range of wavelengths for which it is valid is limited, typically covering the visible spectrum, and large errors occur near the resonant frequencies of the medium. Note that many companies use slightly different, yet equivalent, forms of the formula, resulting in different values for the coefficients.

The accumulated spectral phase of a pulse due to propagation through a medium is calculated by the product of the wave vector, k , with the length of the medium, L . If the magnitude of the wave vector is frequency dependent, then a chirp is induced on the pulse (i.e. a change in the instantaneous frequency as a function of time) as well as a temporal

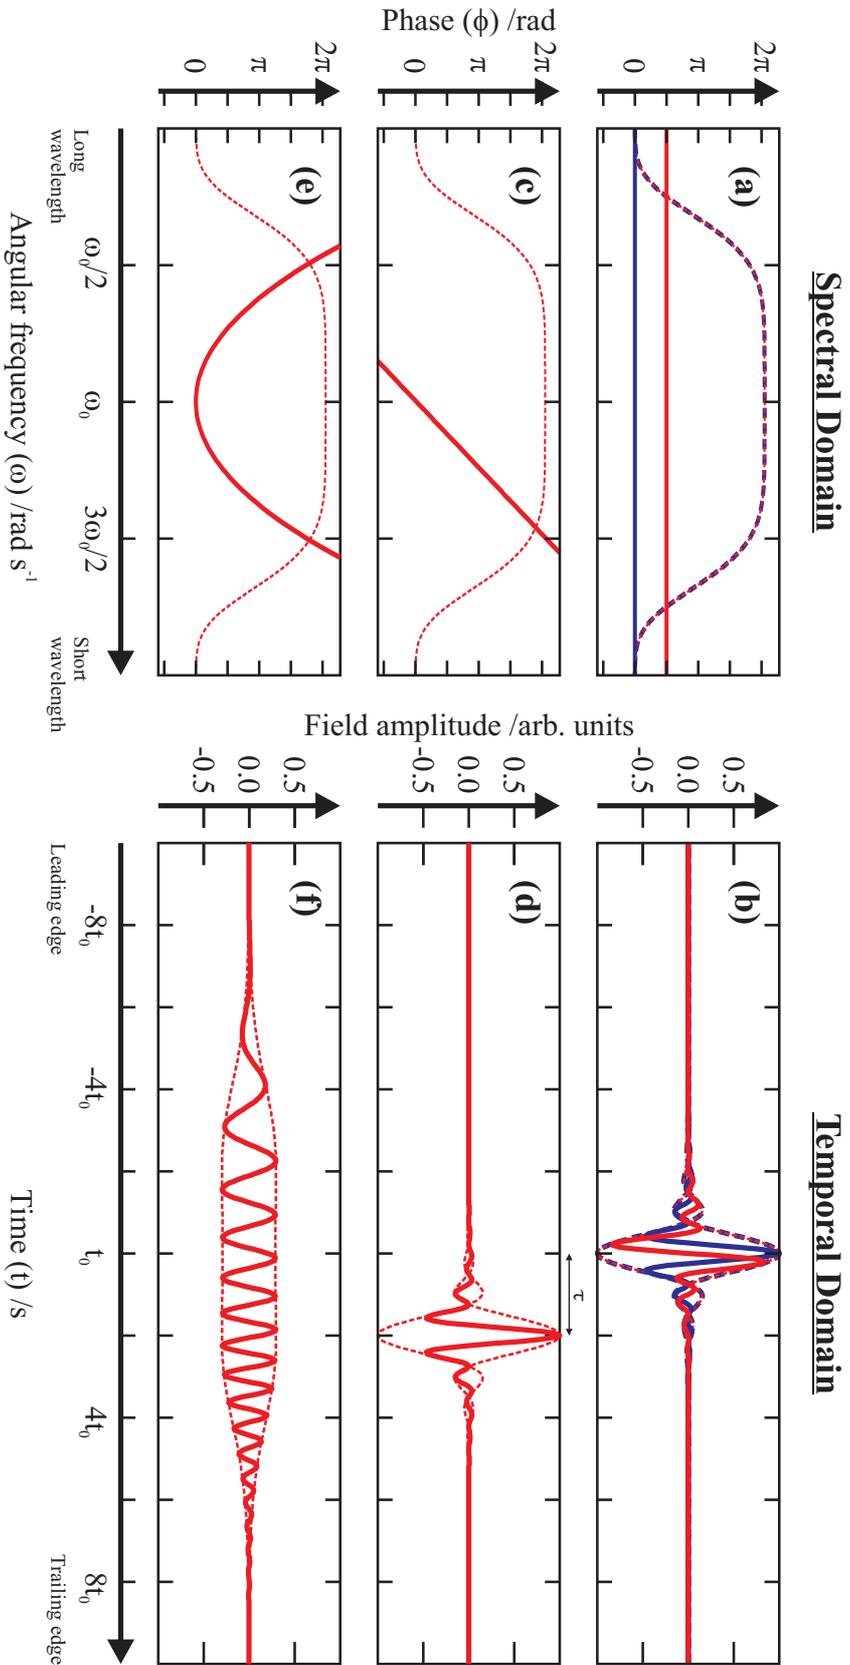


Figure 2.4: The effect of spectral phase on the temporal electric field. Left column: spectral amplitude (dashed lines) and spectral phase (solid lines). Right column: temporal envelope (dashed lines) and temporal electric field (solid lines). (a) A change in the absolute phase, $\tilde{\phi}(\omega) = 0$ (red) and $\tilde{\phi}(\omega) = \pi/2$ (blue) causes a shift in the peak of the electric field relative to the peak in the envelope as shown in (b). (c) A linear spectral phase $\tilde{\phi}(\omega) = \omega\tau_g$ results in a group delay, τ_g , of the pulse, shown in (d). (e) A quadratic spectral phase, $\tilde{\phi}(\omega) = \tilde{\phi}''/2(\omega - \omega_0)^2$ causes a temporal broadening of the pulse, shown in (f). A quadratic spectral phase also causes a linear sweep in frequency (known as linear chirp) as a function time; this is clearly shown by the temporal field (solid red) in (f).

broadening of the envelope function. The accumulated phase is typically expressed as a Taylor expansion about a reference frequency, ω_0 :

$$\begin{aligned}\tilde{\phi}(\omega) &= k(\omega)L \\ &= \frac{n(\omega)\omega}{c}L \quad (\text{for material dispersion}) \\ &= \sum_m \frac{1}{m!} \tilde{\phi}^{(m)}(\omega_0) (\omega - \omega_0)^m.\end{aligned}\quad (2.37)$$

where $\tilde{\phi}^{(n)}(\omega_0)$ is the n^{th} derivative of the phase with respect to frequency, evaluated at the reference frequency. The coefficients have units of $[\text{s}]^n/[\text{rad}]^{n-1}$, or more commonly $[\text{fs}]^n/[\text{rad}]^{n-1}$. However, the $[\text{rad}]$ term is often not quoted, and for simplicity I quote values in terms of $[\text{fs}]^n$, and it is assumed that the angular frequency (compared to natural frequency) is being used. It is also possible to express the temporal phase as a Taylor expansion:

$$\phi(t) = \sum_m \frac{1}{m!} \phi^{(m)}(t_0) (t - t_0)^m, \quad (2.38)$$

where t_0 is a reference time. The lower order terms have simple physical interpretations, which will be discussed below.

The absolute spectral phase (zeroth order term), $\tilde{\phi}^{(0)}(\omega_0) = \tilde{\phi}_0$ is the phase accumulated by the reference frequency ω_0 . For simple pulse shapes, this relates the location of the reference frequency to the peak of the pulse intensity. The absolute temporal phase, known as the *carrier envelope offset* (CEO) or *carrier envelope phase* (CEP), $\phi^{(0)}(t_0) = \phi_0$, relates the location of the peak of the electric field to the peak of the pulse intensity envelope. For a given pulse shape (i.e. for a given spectral envelope and spectral phase), a change of the absolute spectral phase does not affect the temporal intensity, but will cause a shift of the CEP equal to the change of the absolute spectral phase. For example, figure 2.4 (b) shows two pulses with equal spectra, except for a difference in the absolute spectral phase as shown in figure 2.4 (a). For this pulse shape, zero absolute spectral phase results in a ‘cosine’ pulse shape (shown in blue in the figure), such that the peak of the electric field coincides with the peak of the pulse intensity. A shift of $\pi/2$ rad in the spectral phase

causes an equal shift in the CEP, resulting in a ‘sine’ pulse (shown as red in the figure): the electric field is zero at the peak of the pulse intensity and two anti-aligned peaks in the electric field strength with equal magnitude occur either side of the intensity envelope maximum. For a pulse containing many cycles, this term is often neglected because it does not play a role in the interaction of pulses with matter. When using few-cycle pulses (i.e. pulses with one or two cycles within the intensity envelope) with phenomena sensitive to the oscillating electric field, this term is very important because the value of the peak electric field varies significantly with the CEP,^[39]. This is evident in figure 2.4 (b) where the peak magnitude of the electric field strength is less for the sine pulse compared to the cosine pulse for the same peak intensity. Since the absolute phase is generally only a concern for few-cycle pulses, a discussion of measurement and stabilisation of the absolute phase will be presented in section 3.1.2.

The linear term in the spectral phase, $\tilde{\phi}^{(1)}(\omega_0) = \tau_g$, is the time taken for a wave packet centred at the reference frequency to propagate between reference frames. For example, a pulse with a linear spectral phase, $\tilde{\phi}(\omega) = \omega\tau_g$ (figure 2.4 (c)), results in the pulse arriving at time $t_0 = \tau_g$ (figure 2.4 (d)). Hence, this term is defined as the *group delay* (GD), although *group velocity* (GV) is often used. Strictly speaking, the group velocity, v_g , given by

$$\begin{aligned} v_g &= \left(\frac{\partial k}{\partial \omega} \right)^{-1} \\ &= L \left(\frac{\partial \tilde{\phi}}{\partial \omega} \right)^{-1}, \end{aligned} \quad (2.39)$$

is the velocity at which the pulse envelope, hence the energy in the absence of any resonances, travels through the medium. This term is not accessible by most characterisation techniques. However, it is also not considered of importance because a time reference is usually not critical. In pump-probe experiments however, the timing between the pump and probe usually *is* of importance. This information can be obtained via a cross-correlation of the pump and probe pulse in a nonlinear medium.

If the phase velocity is frequency dependent, as is the case for materials, then each frequency obtains a different delay, resulting in *group delay dispersion* (GDD). This is quantified by the second order term, $\tilde{\phi}^{(2)}(\omega_0) = \frac{\partial^2 \tilde{\phi}(\omega)}{\partial \omega^2} \Big|_{\omega_0} = \frac{\partial \tau_g}{\partial \omega}$, in equation (2.37). The net effect of GDD is to temporally broaden the pulse envelope, by inducing a sweep of the carrier frequency in time, as shown in figure 2.4 (e). Note that the lower frequencies (i.e. ‘red’) arrive before the higher frequencies (i.e. ‘blue’).

The instantaneous frequency can be defined as the gradient of the temporal phase:

$$\omega_{\text{inst}}(t_0) = \phi^{(1)}(t_0) = \frac{\partial \phi(t)}{\partial t} \Big|_{t=t_0}, \quad (2.40)$$

and quantifies the average number of oscillations in the electric field in a given time interval, in the limit of the time interval tending to zero. The instantaneous frequency is useful for describing the change in the carrier frequency of the pulse as a function of time. For a GDD that dominates the spectral phase, and is large compared to the inverse of the bandwidth ($\tilde{\phi}^{(2)} \gg \Delta\omega^{-2}$, where $\Delta\omega$ is the pulse bandwidth), the instantaneous frequency varies linearly as a function of time. Thus the temporal envelope maps out the spectral envelope. The variation in the group velocity as a function of frequency is given by the *group velocity dispersion* (GVD), defined as:

$$\begin{aligned} \text{GVD} &= \left(\frac{\partial^2 k}{\partial \omega^2} \right)^{-1} \\ &= L \left(\frac{\partial^2 \tilde{\phi}}{\partial \omega^2} \right)^{-1}. \end{aligned} \quad (2.41)$$

Most materials exhibit normal dispersion ($\frac{\partial \tau_g}{\partial \omega} > 0$) at optical wavelengths, resulting from a decrease in group velocity with increasing frequency and hence an increase in the instantaneous frequency with time. Anomalous dispersion occurs when $\frac{\partial \tau_g}{\partial \omega} < 0$, resulting from an increase in group velocity with frequency and hence a pulse with a decreasing instantaneous frequency with time. Anomalous dispersion normally occurs near the short wavelength side of an absorption resonance, and can occur in nonlinear interactions or from certain optical components (e.g. prism/grating compressors).

It is often useful to consider a Taylor expansion of the spectral phase as it enables one to consider the effects of materials and optical components. For relatively low bandwidths and relatively simple spectral phase, it is often sufficient to only consider up to second order. However, as the bandwidth becomes large, it is also necessary to consider higher order terms. These terms no longer have specific names and are referred to by their order number, e.g. third-order dispersion (TOD) and fourth-order dispersion (FOD). A pulse containing higher order dispersion is said to be nonlinearly chirped, which means the instantaneous frequency varies nonlinearly with time. It is often not necessary or practical to consider beyond sixth-order dispersion (although most of the time, only up to FOD is considered). Analysis of such complicated pulses tends to be done numerically with the complete phase, as measured by a characterisation device.

It is worth noting that there are many occasions when a Taylor expansion is not a sufficient description of the spectral phase. Instances when this is the case include near a resonance, during some nonlinear processes such as white light generation^[40,41], or when a complex spectral phase is applied manually via a pulse shaper. Reasons why the Taylor expansion may be insufficient include occasions when there is an abrupt change in phase (e.g. near resonance) or when the phase exhibits multiple oscillations (e.g. white light generation or from chirped mirrors). However, the dominant dispersion is typically due to GDD, and this is normally of primary concern. In pulse shaping, it is often possible to apply an almost arbitrary spectral phase — limited by the spectral resolution of the device — which may not necessarily be expressed as a low-order Taylor expansion.

2.2.3 The role of phase in ultrashort pulses — summary

Equation (2.4) shows that it is possible to represent the spatio-temporal evolution of an optical pulse by its Fourier (i.e. frequency and angular) spectral amplitudes, which are in general complex values consisting of an envelope and phase. The envelope represents the (square root of the) intensity of the Fourier component and the gradient of the phase relates each Fourier component in physical time or space.

For example, consider the spatial–angular distributions in terms of ray optics. The envelope function for the angular distribution represents the strength for each directional ray component. The gradient of the angular phase represents the start of each ray in space. Conversely, the envelope function of the complex spatial distribution represents the electric field spatial flux, whereas the gradient of the spatial phase represents the local wavefront normal at that position (i.e. the direction of the ray at that point).

This idea is also true in the temporal–frequency representation. The spectral envelope represents the strength of the frequency components, whereas the gradient of the spectral phase determines the ‘arrival time’ of a small wavepacket centred around that frequency component. Conversely, the temporal envelope represents the electric field flux as a function of time. The gradient of the temporal phase gives the ‘instantaneous’ frequency. However, it is important to note that these concepts are meaningless in the case of an envelope function that varies on a scale of the carrier frequency or wavelength.

A measurement of the intensity is insufficient to numerically propagate a pulse, as it is also necessary to relate the spatial phase of each component with respect to all others. This requires a measurement of the spatial phase. In ultrafast optics, both spectral intensity and phase information is required in order to reconstruct the temporal electric field. It is the measurement and manipulation of the spectral phase (i.e. dispersion compensation), which has allowed progress in ultrafast optics to date. Continual progress and development will inevitably be based on extending these ideas to even shorter durations or more complex pulse shapes for use in time-resolved spectroscopy or coherent control experiments, for example.

The previous section outlined some common examples of phase accumulated during propagation and illustrated their effects on the pulse. However, the phase accumulated can be much more complicated, especially in the case of ultrashort pulses, and may not be modelled by means of a simple Taylor expansion. Further information on the role of phase in ultrashort pulses can be found in Walmsley *et al.*^[37]

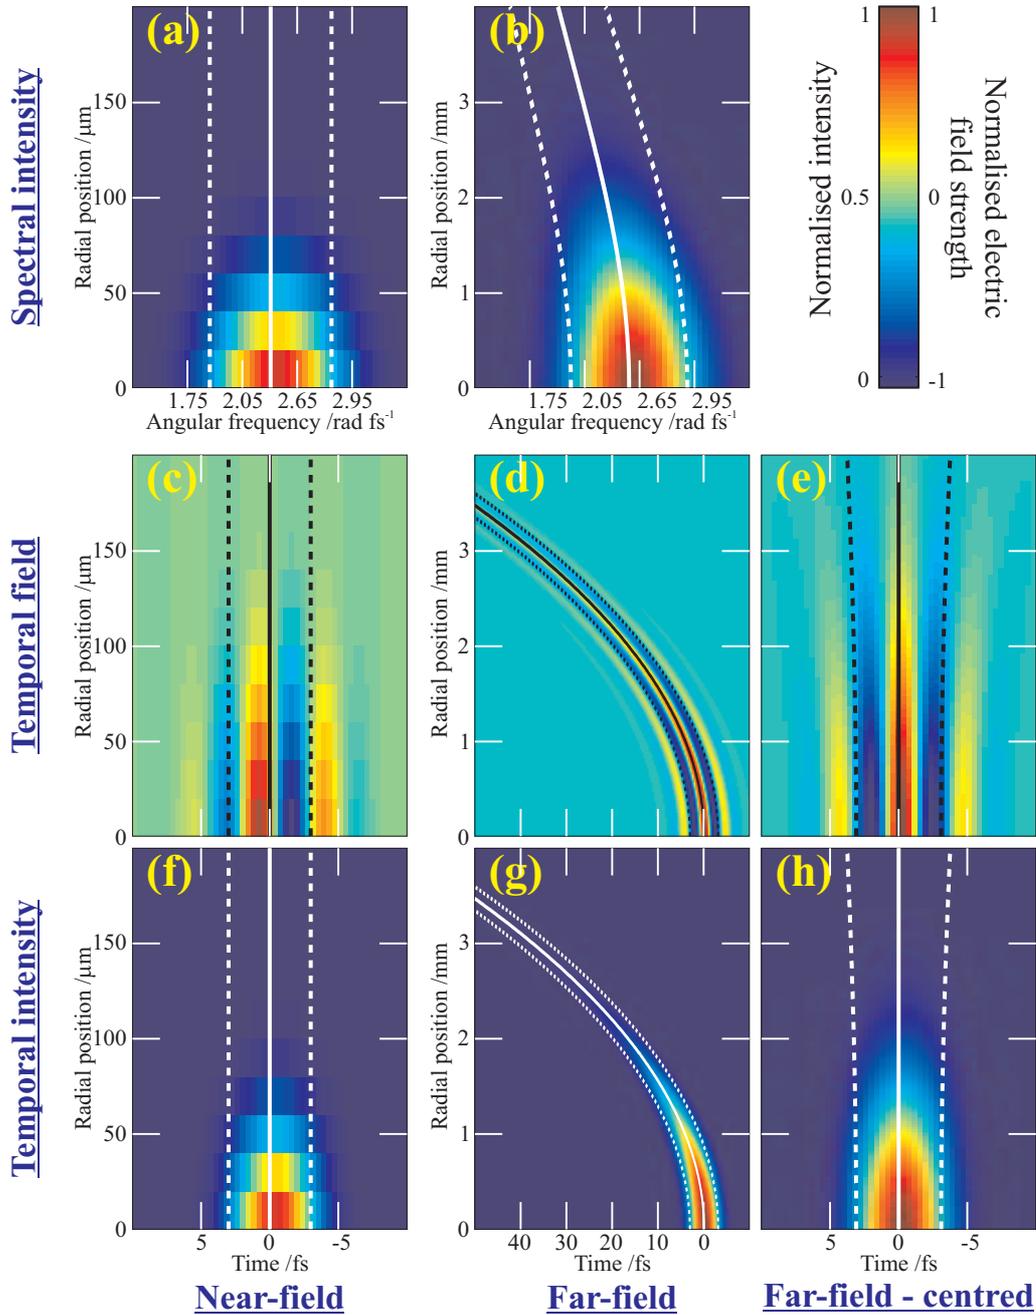


Figure 2.5: Space-time coupling (STC) of an octave spanning pulse due to a frequency dependent mode size (FDMS). The left column shows a factorable Gaussian beam (i.e. one written in the form of equation (2.42)) at the focal plane / near-field. The far-field (at a distance 0.4 m from the focal plane) is shown in the two right-hand columns. The middle column shows the delay as a function of radial position. The far right column has no delay. The top-most row shows the spectral intensity, the middle row shows the electric field in time and the bottom row shows the temporal intensity. The solid lines show the position of the peak intensity. The dotted lines show the position when the intensity is e^{-2} of the peak intensity. Note the different scale for the radial axis for the far field compared to the near field results in a smoother image.

2.3 Space-time coupling of ultrashort pulses

The description of optical pulses so far has only considered the spatial and temporal field distributions separately. This is sufficient provided the temporal profile is independent of position, i.e.:

$$E(t, x, y) = E_{\text{time}}(t) E_{\text{space}}(x, y). \quad (2.42)$$

However, equation (2.42) is not always satisfied, as the temporal evolution of the field may depend on position, i.e. the pulse is said to exhibit space-time coupling (STC). Equation (2.42) also implies, when considering the complex amplitudes, that if a pulse does not exhibit STC then it must be possible to write the pulse in a factorable form in any combination of temporal–frequency and spatial–angular representations. Thus, STC occurs when there is a spatial dependence on either the pulse spectrum or the spectral phase, or both. STC is often used to perform temporal shaping of optical pulses, e.g. in a 4-f pulse shaper.^[37,42] However, STC can also arise due to misalignment of optical setups, from propagation of large bandwidths or from nonlinear effects. Although STC can arise in a complicated form, there are several forms that are often encountered or utilised, which I outline below.

2.3.1 Frequency dependent mode size

It is clear that a pulse will exhibit STC after diffraction due to the fact that the beam size is frequency dependent. This is termed frequency-dependent mode size (FDMS). Usually, this effect is only noticeable for octave spanning bandwidths or more (i.e. few cycle pulses). Figure 2.5 illustrates this phenomenon for a few-cycle pulse upon propagation in free-space. The initial pulse can be written in the form given in equation (2.42) where both E_{time} and E_{space} are Gaussian functions with no phase, and the pulse exhibits cylindrical symmetry in space. The pulse bandwidth is equal to 1.33 rad/fs at a central frequency of 2.35 rad/fs (800 nm) resulting in a Fourier transform limited pulse duration of 6 fs (~ 2 cycles), — all widths quoted correspond to the full width at e^{-2} of the peak intensity. The beam waist,

as defined in section 2.2, is equal to $50 \mu\text{m}$. The initial spatial-spectral intensity is shown in figure 2.5 (a). The initial temporal electric field and temporal intensity as a function of position are shown in (c) and (f) respectively.

The few-cycle pulse shown in (a) is propagated a distance of 0.4 m (corresponding to the far-field) in free space according to equation (2.23). It is clear that all pulses will exhibit a *group front* delay along the radial co-ordinate compared to the centre of the beam — the centre portion of the pulse will arrive before the off-axis portions, as shown in figures 2.5 (d) and (g). This is due to the off-axis beams traversing a longer path. In addition however, the carrier frequency will decrease as a function of the radial co-ordinate, a result of the frequency dependent diffraction. This is evident from equations 2.20 and 2.23, or equations 2.14 to 2.18 and is clearly shown in figure 2.5 (b). A further effect which is not so obvious is that the pulse duration, and even pulse shape in the case of a complicated pulse spectrum and spatial profile, is also radially dependent. Figures 2.5 (e) and (h) show the temporal electric field and intensity respectively with zero group delay, i.e. $E(t, r) \rightarrow E(t - t_0(r), r)$, where $t_0(r)$ is the radially dependent group delay. It is clear that the pulse duration increases as a function of radial position. This effect can result in a reduction of the temporal resolution in certain time-resolved experiments.

The exact effect propagation has on the pulse depends on the pulse shape, spatial profile and the initial STC. Thus it is necessary to consider the pulse generating conditions and the beam size at the site of the measurement relative to the generation. For example, a confocal cavity results in a mode size given by

$$w(z) = \sqrt{\frac{2z_0}{k} \left[1 + \left(\frac{z}{z_0} \right)^2 \right]}, \quad (2.43)$$

where z_0 is a constant given by the cavity length, resulting in a FDMS. However, other optical elements inside the cavity, such as dispersion compensation, can result in more complicated mode structure.^[43] Typical pulse compression techniques (e.g. hollow-core fibre pulse compression and filamentation) can result in even more complicated structures.

Although the STC effects are mostly noticeable far off-axis (e.g. $r > 2w(z)$) where the intensity is relatively low, they may be observed under certain conditions such as after generating an annular beam or in extremely intense pulses.

2.3.2 Focusing of a short optical pulse

When a short pulse is focused by a lens, portions of the wavefront at the edge of the lens will experience a different thickness of material than the portions at the centre of the lens, and thus will experience different dispersion and group delays. Portions of the wavefront that traverse the lens near the edge will arrive at the focus before portions that traverse the lens through the centre. The group front of the beam will no longer be a spherical wave centred at the focus. The net result is that if one was to take into account all the converging rays, the apparent pulse duration at the focus can be dramatically larger than the effective duration of each ray.

This phenomenon was first pointed out by Bor using simple ray tracing calculations.^[44] The radially dependent change in propagation time, $\Delta T(r)$, of a ray from a plane before the lens to the focal point can be calculated as

$$\Delta T(r) = -\frac{r^2}{2cf^2} \bar{\lambda} \frac{df}{d\lambda}, \quad (2.44)$$

where r is the radial co-ordinate, f is the focal length of the lens at the mean wavelength $\bar{\lambda}$, and $\frac{df}{d\lambda}$ is the variation in the focal length as a function of wavelength. A singlet lens of refractive index $n(\lambda)$ has $\frac{df}{d\lambda} = [f/(n-1)] \frac{dn}{d\lambda}$.

As an example, a singlet lens made of BK7 with a focal length of 10 cm will introduce a 50 fs delay between the paraxial ray and a ray 1 cm off-axis. This is obviously significant for few-cycle pulses. This phenomenon can also be described in the frequency domain, which results in *chromatic aberration* (i.e. a frequency dependent focal length and mode size). Using an achromat optimised for the correct wavelength can reduce this group delay. However, the pulse will also experience GDD due to the material itself. This will result

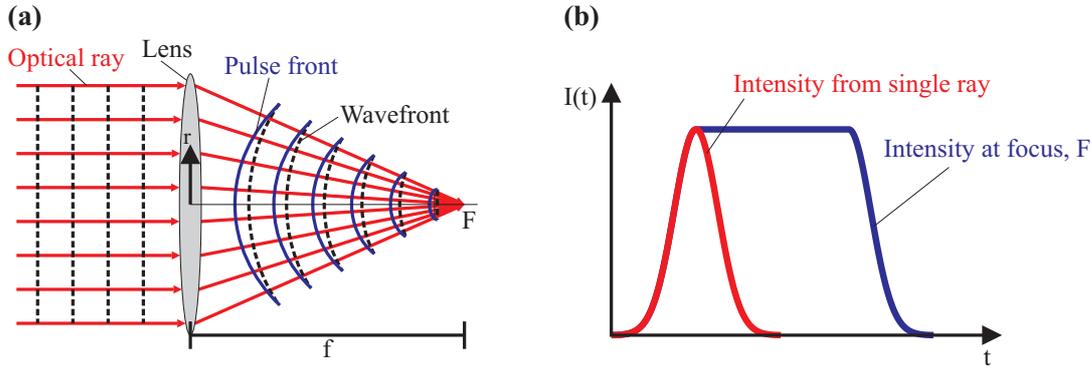


Figure 2.6: Apparent increase in the duration of an ultrashort pulse due to a lens. (a) A singlet lens focuses a collimated beam (the optical rays are shown as red arrows) at F . The group front (blue) does not coincide with the spherical wavefront (dashed black) and thus the temporal intensity will be considerably longer at the focus than before the lens (b).

in the pulse stretching in time due to material dispersion alone, which is significant for few-cycle pulses. For example, a 10 fs pulse is stretched by a factor of four by several millimetres of BK7 (i.e. at the centre of the lens). The dispersion at the edge of the lens is less, which results in some STC of the pulse. However, this effect is considerably less than the ‘stretch’ due to the difference in the group delay and phase delay (figure 2.6).

2.3.3 Angular dispersion

Angular dispersion (AD) occurs when there exists a frequency dependent propagation direction, typically generated from prisms or gratings as shown in figure 2.7 (a). Angular dispersion is also equivalent to *pulse front tilt* (PFT): the arrival time of the pulse varies as a function of position and is shown in figure 2.7 (b). Consider an optical pulse, $\tilde{E}(\omega, k_x)$, that is linearly angularly dispersed in the x -direction:

$$\tilde{E}_{AD} = \tilde{E}(\omega, k_x - \alpha\omega), \quad (2.45)$$

where α quantifies the angular dispersion (which is assumed to be a constant in this case, but in general can be a function of frequency) and \tilde{E}_{AD} is the angularly dispersed pulse.

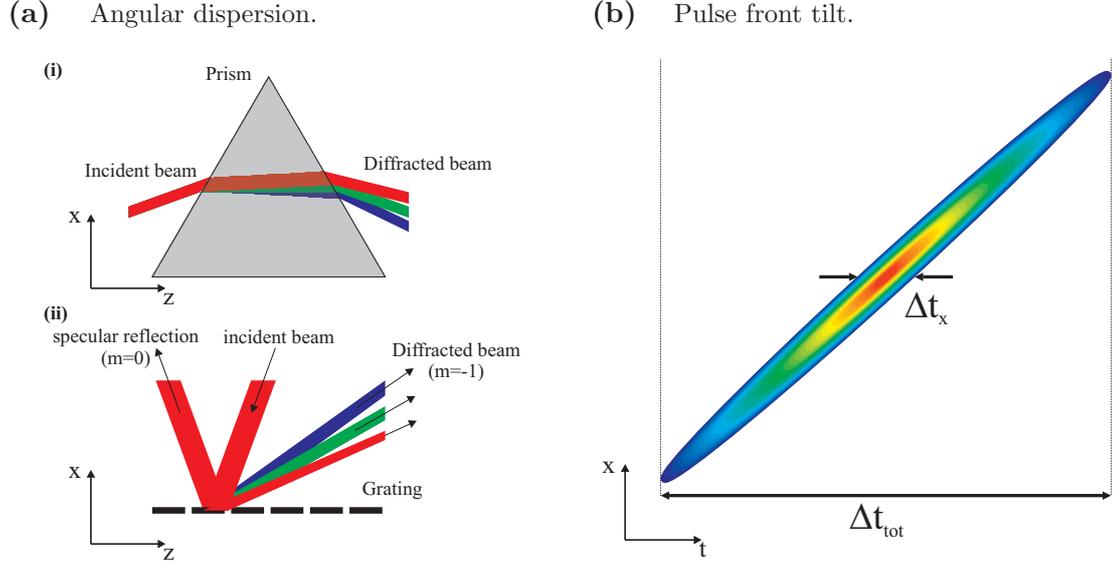


Figure 2.7: Schematic of angular dispersion / pulse front tilt. (a) An input beam is diffracted by (i) a prism or (ii) a grating such that the propagation direction is frequency dependent. (b) Angular dispersion is equivalent to pulse front tilt, which is when the arrival time of the pulse depends on the spatial position. The pulse duration at any point (Δt_x) is less than the spatially integrated pulse duration (Δt_{tot}) and can reduce the temporal resolution in spatially integrated measurements, even though the spectrum is less at a given spatial position than when integrated across the whole beam.

Performing an inverse Fourier transform over k_x and using the shift theorem results in

$$\begin{aligned}\tilde{E}_{\omega x} &= \mathfrak{F}^{-1} \left\{ \tilde{E}_{\text{AD}}; k_x \rightarrow x \right\} \\ &= \tilde{E}'(\omega, x) e^{i\alpha\omega x},\end{aligned}\quad (2.46)$$

where $\tilde{E}'(\omega, x) = \mathfrak{F}^{-1} \left\{ \tilde{E}(\omega, k_x); k_x \rightarrow x \right\}$ is the space-frequency distribution of the non angularly dispersed pulse. Performing a Fourier transform over frequency and using the shift theorem again to calculate the space-time field yields a pulse with PFT:

$$\begin{aligned}E_{\text{PFT}} &= \mathfrak{F} \left\{ \tilde{E}_{\omega x}; \omega \rightarrow t \right\} \\ &= E(t - \alpha x, x),\end{aligned}\quad (2.47)$$

where $E(t, x) = \mathfrak{F} \left\{ \tilde{E}(\omega, x); \omega \rightarrow t \right\}$ is the space-time field of the pulse without angular dispersion.

Pulse front tilt is generally detrimental to time-resolved spectroscopy as it results in temporal smearing in spatially integrated measurements. This is because the arrival time of the probe at the sample is less well defined if the measurement is averaged over the spatial profile of the beam. Angular dispersion normally arises from improperly aligned prism/grating stretchers/compressors or from transmission through a wedged optic. Later in this thesis, it will be shown that angular dispersion, and thus pulse front tilt, can also arise in high harmonic generation. It is therefore quite possible that the temporal smearing introduced from averaging over the complete spatial profile of the harmonics is comparable to the pulse duration, and thus can smear out any dynamical effects being observed. One can minimise the problem by spatially filtering the beam, but this will significantly reduce the signal-to-noise ratio (SNR). A better solution is to optimise the generation to minimise the space-time coupling at the sample.

2.3.4 Spatial chirp

Spatial chirp occurs when the central wavelength changes as a function of position and is commonly due to improper alignment of prism/grating stretchers/compressors. It can be shown, via the same analysis as for angular dispersion and pulse front tilt, that spatial chirp in a collimated beam results in pulse front tilt when the beam is focused to a spot (figure 2.8). This is evident from the fact that each frequency will have a different mean tilt due to passing the focusing optic at different spatial positions, and thus the pulse will exhibit angular dispersion at the focus. Spatial chirp is also present in high harmonic generation and needs to be measured and reduced to obtain the optimal pulse for time-resolved spectroscopy.

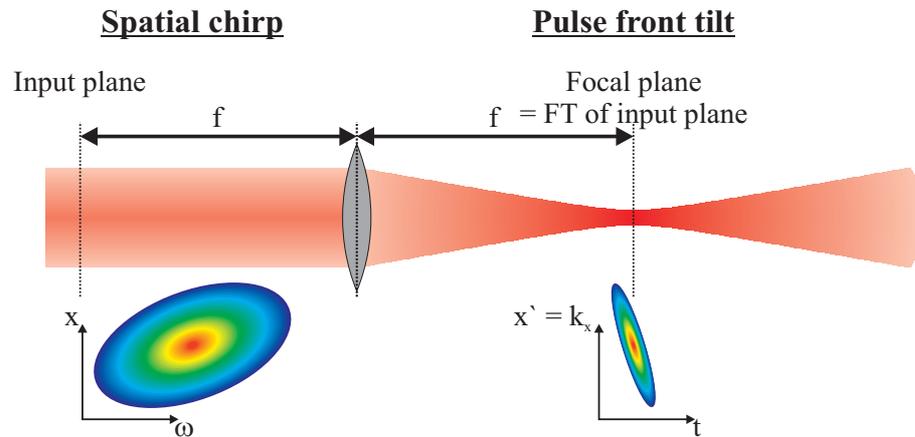


Figure 2.8: A collimated beam that exhibits spatial chirp (i.e. a spatially dependent carrier frequency) will exhibit pulse front tilt at the focal plane of a lens, as the angular distribution is mapped to space due to the Fourier transform property of the lens.

2.4 Measuring the wavefront of optical beams

Due to the Fourier transform property of a lens, it is possible to switch between the spatial and angular distributions via linear optics. It is simple to perform spatial-phase operations, and hence perform complete self-referencing spatial characterisation of optical beams via linear optics only. Thus I outline some of the most common examples for characterising the spatial profile and phase of optical pulses. Although it is not possible to perform the same operations in the time domain using linear optics, the ideas still form the basis of temporal characterisation techniques as well as forming important characterisation tools in their own right.

2.4.1 Shack-Hartmann wavefront sensor

The Shack-Hartmann wavefront sensor^[45,46] (SHWS) is an array of microlenses (lenslets) of equal focal length. Each lenslet will focus the local wavefront to a spot on a CCD. The principle of the SHWS is shown in figure 2.9. For a well collimated beam, the spots will form a grid with a predefined spacing. Wavefront aberrations can be modelled as an independent tilt at each position in the beam. This is then measured by the displacement of the spots on the CCD from their nominal position. Such a device is not practical for

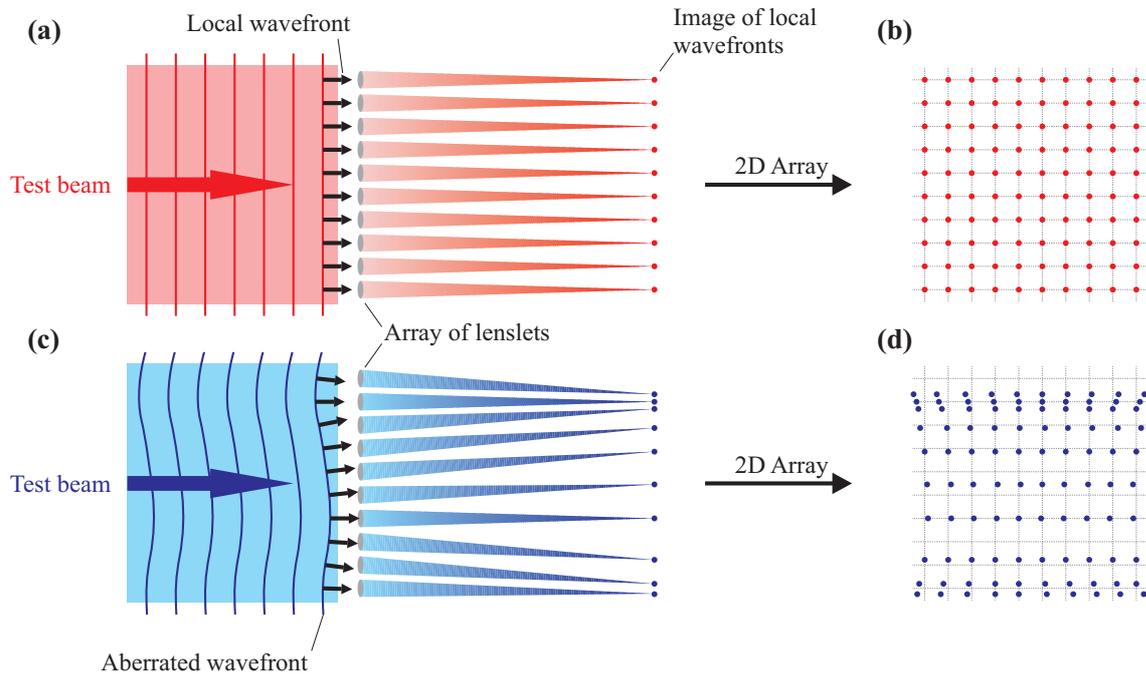


Figure 2.9: Principle of the Shack-Hartmann wavefront sensor. (a) A wavefront without any aberrations will form an array of focused spots on the screen (CCD) with an even spacing. (b) In two dimensions, this appears as a grid with a spacing equal to the lenslet spacing. (c) An aberrated wavefront will form an array of unevenly spaced spots. (d) In two dimensions, the spots will have moved away from the reference positions by an amount proportional to and in the same directions as the local tilt on the wavefront at the lenslet.

ultrashort pulses due to the chromatic aberration of the lenses and it can not measure STC because the grid needs to be measured for each frequency. However, the device can be extremely compact and is alignment insensitive.

2.4.2 Point diffraction interferometry

Point diffraction interferometry^[47] (PDI) is one of the simplest forms of a common path interferometer. The test beam wavefront is measured by the interference of itself with a spherical wavefront derived from the test beam as illustrated in figure 2.10 (a). The spherical wave is formed by passing the test beam through a semi-transparent film with a very small aperture in the centre. The aperture causes the central part of the beam to diffract, forming a near-perfect spherical wave (perfect in the limit of a point aperture) to act as a well defined reference wavefront. The rest of the beam is attenuated as it passes

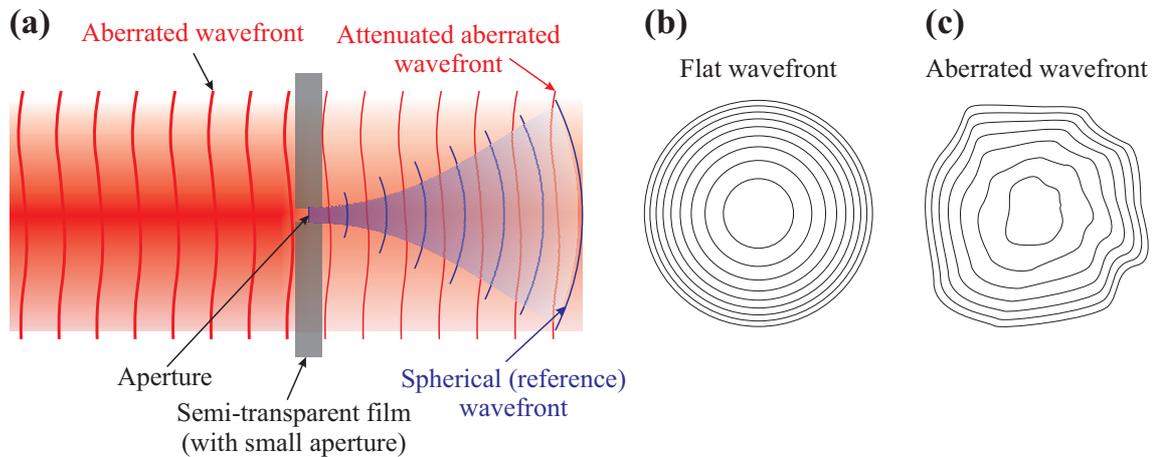


Figure 2.10: Principle of point diffraction interferometry. (a) A small aperture in a semi-transparent screen causes a small section of the test beam to diffract with a near-perfect spherical wavefront (blue beam), forming the reference beam. The rest of the beam passes through the film without any phase alteration but with some attenuation (attenuated red beam), forming the test beam. The attenuation is to balance the intensity of the reference and test beams at the detection screen to improve the fringe contrast. (b) Circular fringes form for a test beam with a flat wavefront. (c) Distorted circular fringes form for a test beam with an aberrated wavefront.

through the film to balance the intensities and thus increase the fringe contrast.

The test beam can be collimated, which results in circular fringes for a perfectly flat test wavefront, as shown in figure 2.10 (b). Any aberration results in structure on the rings, as shown in figure 2.10 (c). The phase can be extracted via a Fourier filtering algorithm, but the spherical wavefront of the reference beam needs to be subtracted in order to recover the aberrated wavefront.

An alternative arrangement allows measurement of focused beams, in which case the film and aperture are placed at the centre of the focus. This allows one to test the aberrations of lenses and other optical systems. In this case, the aperture size, a , needs to be much smaller than the size of the theoretical Airy disc of the lens: $a \ll 1.22\lambda F$, where F is the Fresnel number of the lens. Typically, the wavefront of the spherical reference beam and the aberrated test beam will have similar diffraction properties and thus similar radii of curvature. In this case, any fringes are the direct result of aberrations on the test beam. However, the size of the aperture typically needs to be much smaller than for collimated beams. It is also much harder to resolve small phase variations because it is

not possible to perform a Fourier filtering routine to extract the interference pattern from the intensity background of the two beams.

PDI has the advantage of being simple to implement and is a very stable interferometer, owing to the common path. It has proved robust in measuring aberrations of optical components and in other applications.^[48,49] However, this technique is very sensitive to the spatial coherence of the test beam and is not well suited to ultra-broadband pulses without assumptions about the space-time coupling.^[50] Both these effects reduce the fringe contrast away from the centre of the beam. However, it should be possible to spectrally resolve the fringes by imaging a slice of the interference pattern in a spectrometer.

2.4.3 Lateral shearing interferometry

In lateral shearing interferometry^[51,52] (LSI), two identical beams are interfered with a small spatial separation (lateral shear). Typically, a tilt is also introduced between the beams, resulting in fringes along the tilt direction that aids in reconstructing the spatial phase.

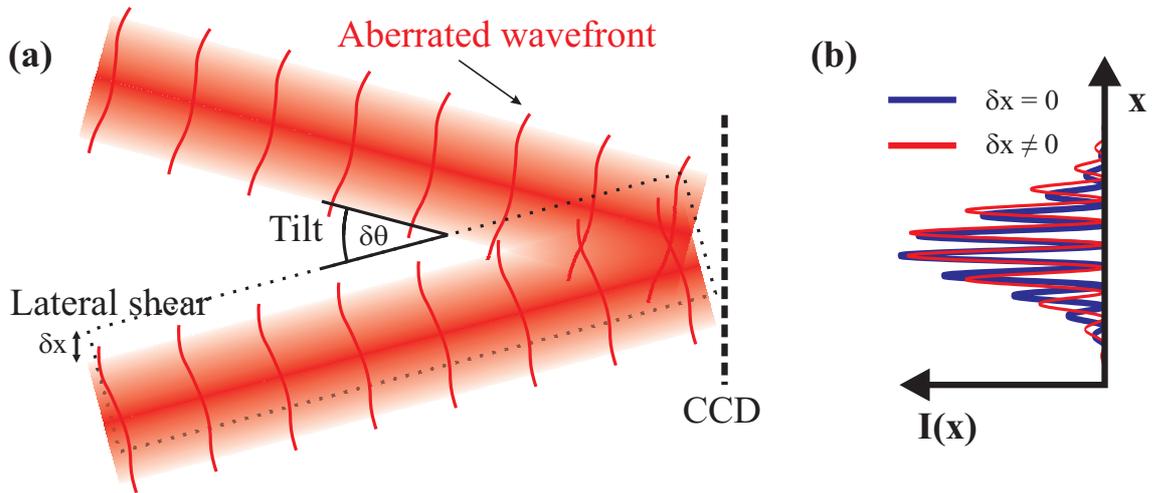


Figure 2.11: Principle of lateral shearing interferometry. (a) Two replica monochromatic wavefronts are interfered onto a CCD with a tilt between them. (b) The tilt produces equidistant fringes with a periodicity of $2\pi/(k \sin \delta\theta)$ provided there is no lateral shear (blue line). Introducing a lateral shear, δx , in addition to the tilt modulates the fringe periodicity. The modulation is directly related to the spatial gradient of the wavefront (scaled by δx); the tilt allows one to extract this via a simple Fourier transform filtering algorithm.

The lateral shear allows the phase at one spatial position to be referenced to the phase at a different position, separated by the lateral shear. If a tilt exists between the two beams, then fringes occur with a periodicity equal to the inverse of the lateral shear. Any wavefront aberrations are encoded as a displacement of the fringes from their nominal position. In the absence of a tilt, then any aberration manifests itself as a change in the intensity distribution relative to a single beam. Reconstruction of the spatial phase requires integration/concatenation of the extracted phase difference (see 2.5.5), and thus cannot measure the spatial phase across a beam with well separated regions in a single shot. However, the spatial phase across each individual region can be measured using a small shear, then the phase between each separated region can be measured by overlapping these regions with each other. By spectrally resolving the interference pattern, it is possible to use LSI to measure the wavefront for each frequency, and hence directly measure STC.^[53]

2.4.4 Radial shearing interferometry

Radial shearing interferometry^[54,55] (RSI) is a similar method to both point diffraction interferometry and lateral shearing interferometry and can be used to measure the two dimensional spatial phase. In RSI, two identical diffracting / converging beams are interfered with a longitudinal displacement between each path relative to the focus. The path difference results in slightly different radii of curvature of the two beams, and allows the phase to be measured between two radially separated points. The difference in curvature of the beams leads to circular fringes. The spatial phase gradient is encoded in a displacement of these fringes, similar to LSI. Spectrally resolving the interference pattern allows one to measure STC.

2.5 Complete temporal characterisation of ultrashort optical pulses

Although it is not possible to measure the temporal field of an ultrashort pulse directly because the pulse duration is much shorter than any current detector response time, equation (2.6) implies that it is possible to measure the pulse completely by measuring the spectrum and spectral phase. The pulse spectrum can easily be measured using a slow detector (i.e one whose response time is much longer than the duration of the pulse) via a spectrometer. It is much more problematic to measure the spectral phase because slow detectors only measure the time-integrated power, and hence are insensitive to phase variations. If a known reference pulse is available, then it is possible to characterise the unknown pulse via a linear measurement method. However, this thesis is concerned with the measurement of pulses when no known reference pulse exists, which is the case for the shortest pulses that can be produced. Such a method is called self-referencing because the pulse being measured is essentially referenced with itself.

In order to perform complete characterisation of an ultrashort pulse, it is necessary to perform some spectral and temporal filtering on the pulse. The filtering essentially allows one to obtain information on the frequency makeup of a pulse as a function of time. Walmsley and Wong formally identified the filters required to perform complete temporal characterisation.^[17,56] Based on the type and order of the filters, Walmsley describe three main classes of techniques: *spectrographic*, *tomographic* and *interferometric*. Spectrographic techniques are based on measurements of filtered time-frequency representations of the pulse. The pulse can be reconstructed via a deconvolution method because the form of the filters are known. Tomographic techniques measure projections of the Wigner representation of the pulse and reconstruct the pulse via a direct algorithm. Interferometric techniques measure a one dimensional slice in the two-frequency correlation function, which completely defines a coherent field, and reconstructs the pulse via a direct algorithm. A general overview of ultrashort pulse characterisation methods is given

by Dorrer and Joffre.^[16] The most common methods for complete characterisation are described below.

2.5.1 Field autocorrelation

The *field autocorrelation*, $A(\tau)$, of a pulse is the time-integrated product of a pulse with a time-delayed version of itself as described by equation (2.48).^[16] From the Wiener-Khinchin theory^[57], it can be shown that the Fourier transform of the field autocorrelation is equal to the pulse spectral intensity, $|\tilde{E}(\omega)|^2$:

$$\begin{aligned} A(\tau) &= \int_{-\infty}^{\infty} dt E(t) E^*(t - \tau) \\ &= \mathfrak{F}^{-1} \left\{ |\tilde{E}(\omega)|^2; \omega \rightarrow \tau \right\}. \end{aligned} \quad (2.48)$$

The field autocorrelation can be measured by placing a slow detector at the output of a Michelson interferometer as shown in figure 2.12. The measured signal, $I_M(\tau)$, is described by

$$\begin{aligned} I_M(\tau) &= \int_{-\infty}^{\infty} dt |E(t) + E(t - \tau)|^2 \\ &= 2 \int_{-\infty}^{\infty} dt I(t) + 2\Re \left\{ \int_{-\infty}^{\infty} dt E(t) E^*(t - \tau) \right\}, \end{aligned} \quad (2.49)$$

where $I(t)$ is the temporal intensity of the pulse. Clearly this is equal to the field autocorrelation plus a background signal. This method forms the basis of *Fourier transform spectroscopy* (FTS). FTS provides much higher spectral resolution than grating spectrometers, limited only by the interferometric stability and available range of the delay. However, it is also clear that such a measurement does not contain any phase information and therefore cannot be used to reconstruct the temporal field. Measurement of the temporal field can be achieved by applying a temporal modulation on the order of the pulse duration to one of the pulses. For ultrashort pulses in the femtosecond and attosecond regimes, this modulation requires a nonlinear interaction. For longer pulses however, typically in the

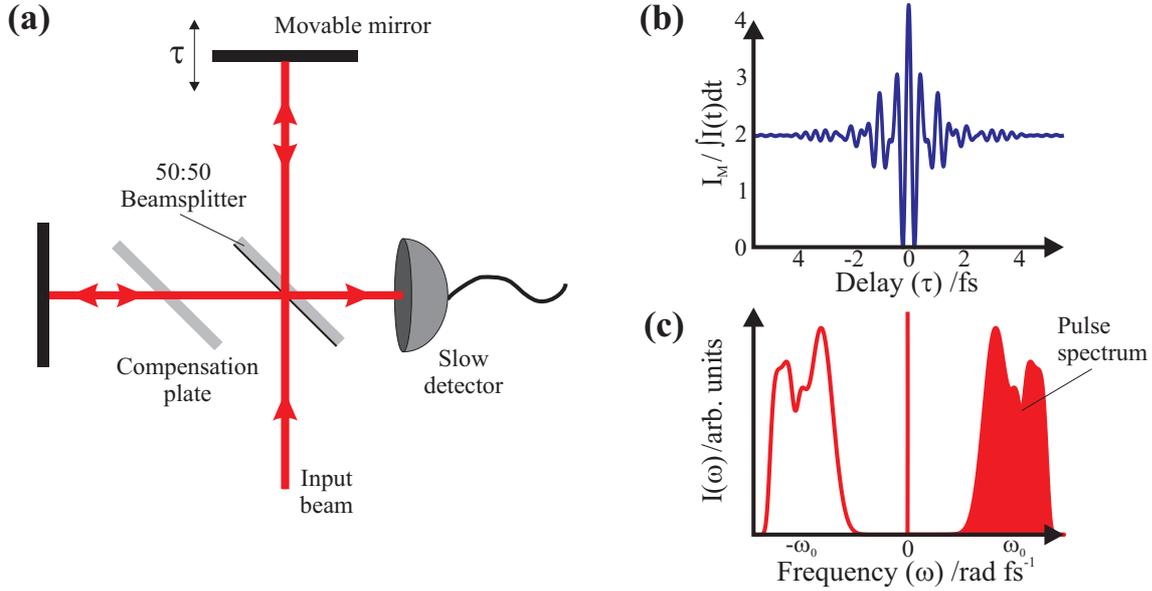


Figure 2.12: The field autocorrelation detected via a Michelson interferometer: (a) setup, (b) field autocorrelation trace and (c) Fourier transform of the field autocorrelation trace. The pulse spectrum is equal to the positive sidelobe, shown by the shaded region shown in (c).

picosecond regime, it is possible to perform self-referenced characterisation using linear optics via an electro-optic modulator (EOM).^[18,58–64]

2.5.2 Intensity and interferometric autocorrelation

By placing a second-order nonlinear crystal after the Michelson interferometer used above, it is possible to measure the second order interferometric autocorrelation, given by

$$\begin{aligned}
 I_2(\tau) &= \int_{-\infty}^{\infty} dt | [E(t) + E(t - \tau)]^2 |^2 \\
 &= 2 \int_{-\infty}^{\infty} dt I^2(t) + 2\Re \left\{ \int_{-\infty}^{\infty} dt [E(t) E^*(t - \tau)]^2 \right\} \\
 &\quad + 4\Re \left\{ \int_{-\infty}^{\infty} dt I(t) E(t) [E(t - \tau) + E(t + \tau)]^* \right\} \\
 &\quad + 4 \int_{-\infty}^{\infty} dt I(t) I(t - \tau).
 \end{aligned} \tag{2.50}$$

The last term in equation (2.50) corresponds to the *intensity autocorrelation*^[65–69] and can be measured directly by spatially overlapping the two beams in a non-collinear geom-

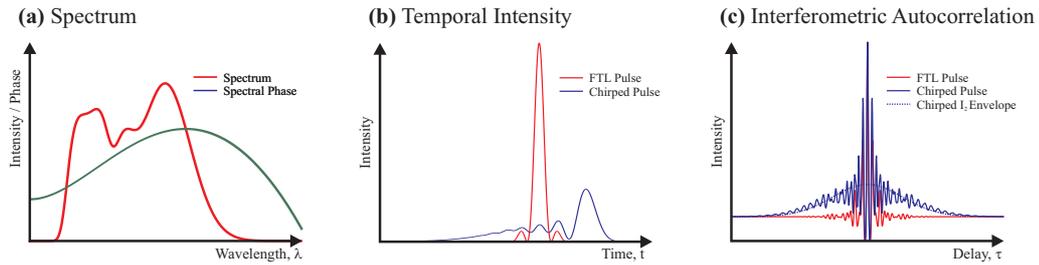


Figure 2.13: Simulated interferometric autocorrelation data. (a) An example pulse spectrum and phase. (b) The temporal intensity at its FTL pulse. (c) The interferometric autocorrelations for both the chirped and FTL pulse, the dashed line represents the Gaussian envelope that is used to give the approximate pulse duration after rescaling the time axis.

etry. However, this does not contain any phase information, thus it can only be used to estimate the pulse duration via assuming some pulse shape and fitting the theoretical autocorrelation to the measured one.^[69] In practise, one typically assumes either a Gaussian or hyperbolic-secant-squared pulse shape and divides the width of the measured autocorrelation trace by a scaling factor to attain the width of the measured pulse. Although such a measurement only gives approximate pulse durations, the simplicity, compactness and cost efficiency of the device make it a useful tool for general alignment and optimisation of pulses and optical devices. The autocorrelation trace is often used to check the reconstruction of more complex, complete characterisation devices.

If a collinear geometry is used, then all the terms in equation (2.50) are present, resulting in an *interferometric autocorrelation*.^[70–78] As this contains products of the electric field, the trace contains some phase information. For a FTL pulse (i.e. one that contains only a constant or linear spectral phase), then fringes will exist over the full extent of the autocorrelation trace (excluding the background). The fringes have the same periodicity as the carrier frequency of the pulse, thus allowing a calibration of the time axis. For a chirped pulse, the fringes wash out in the wings, allowing one to determine the spectral phase via an iterative deconvolution algorithm when combined with a measurement of the pulse spectrum.^[72] The intensity autocorrelation traces for two different pulses with the same spectral intensity are shown in figure 2.13. However, such an inversion routine typically requires an extremely high SNR to be accurate, and can struggle for even mod-

estly shaped pulses. Alignment of an interferometric autocorrelator is checked by ensuring that the ratio of the peak of the trace to the background is 8 : 1, which is evident when considering equation (2.50) for $\tau \rightarrow \infty$ and $\tau = 0$.

2.5.3 FROG

Frequency resolved optical gating (FROG) is a technique based on gating the pulse with itself, and is essentially a frequency resolved autocorrelation trace.^[16,79–82] In optical gating, a short pulse is used to gate a long pulse in a nonlinear crystal. By spectrally resolving this, it is possible to build up a spectrogram of the pulse: a chart of the frequency makeup at each time interval defined by the temporal overlap between the pump and probe pulses. The measured trace can be written as:

$$I_{\text{FROG}}(\omega, \tau) = |\mathfrak{F}^{-1}\{P(t)G(t-\tau); t \rightarrow \omega\}|^2 \quad (2.51)$$

where $P(t) = E(t)$ is the probe pulse being measured, $G(t)$ is the gate function (which in this case is the time delayed pulse, but may not necessarily be the case) and τ is the delay between the probe and gate.

In FROG however, there does not exist a shorter pulse and thus it is the pulse itself that is used as the gate. There exist a variety of versions of FROG with varying nonlinear orders and geometries.^[80,83–89] The most common version is the second harmonic generation (SHG-) FROG — a frequency resolved second order intensity autocorrelation.^[90–92] Due to the symmetry of the process, SHG-FROG suffers from the time ambiguity of SHG autocorrelators: it is not possible to distinguish between $E(t)$ and $E(-t)$. It is possible to determine the sign of the phase experimentally by adding extra known dispersion, which can be done by introducing a piece of glass into the beam before the SHG-FROG.^[93] SHG-FROG offers the highest sensitivity of all FROG methods due to using the lowest order nonlinearity. In addition, it is possible to perform single-shot operation by using the width of the beam to map the delay to spatial position by overlapping the beams in a

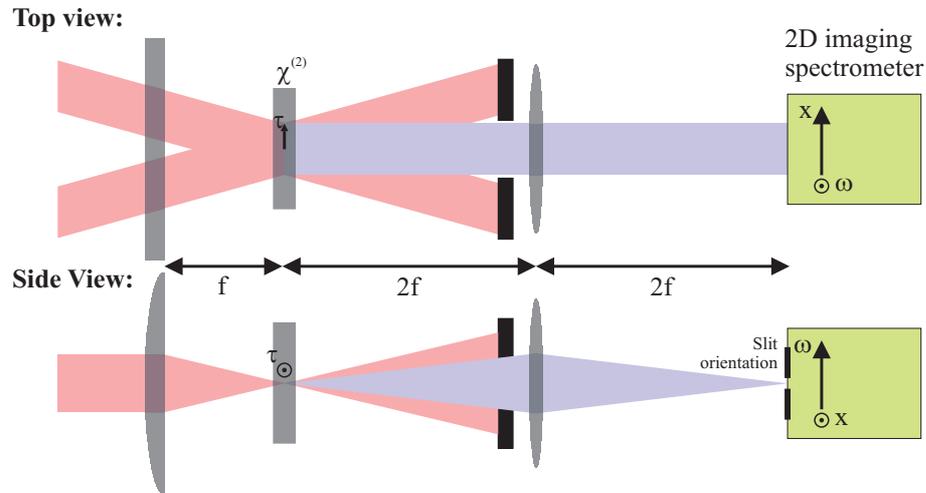


Figure 2.14: Schematic of a single-shot SHG-FROG setup. Two identical collimated beams cross in a non-collinear geometry in a thin $\chi^{(2)}$ nonlinear crystal. The upconverted signal is imaged onto the entrance of a 2D imaging spectrometer. The width of the beams map delay, τ , to space, x . The slit of the spectrometer is aligned parallel to the space-time mapping. The beams may be focused into the crystal in the direction orthogonal to the space-time mapping by a cylindrical lens to improve efficiency.

non-collinear geometry, as shown in figure 2.14.^[92,94] The two beams may be generated from a ‘misaligned’ Michelson interferometer, or by overlapping two halves of the same beam. The latter method has practical advantages (i.e. simpler design and fewer optical components). FROG traces can become distorted due to STC of the probe pulse. In SHG-FROG where the two beams are derived from different halves of the probe pulse, simple forms of STC (e.g. spatial chirp and pulse-front tilt) can be measured from these distortions.^[95,96]

The inversion routine for FROG is a two-dimensional phase retrieval problem, which is known to have a unique solution, in comparison to the one-dimensional phase retrieval problem of the intensity autocorrelation.^[97,98] However, FROG requires an iterative algorithm to reconstruct one dimensional information (of length $2N$) from a two dimensional data set (of length N^2), where N is the number of sample points in one dimension of the FROG trace. Thus it is not the most efficient method, but contains redundancy, which can be used to check the alignment of the device.

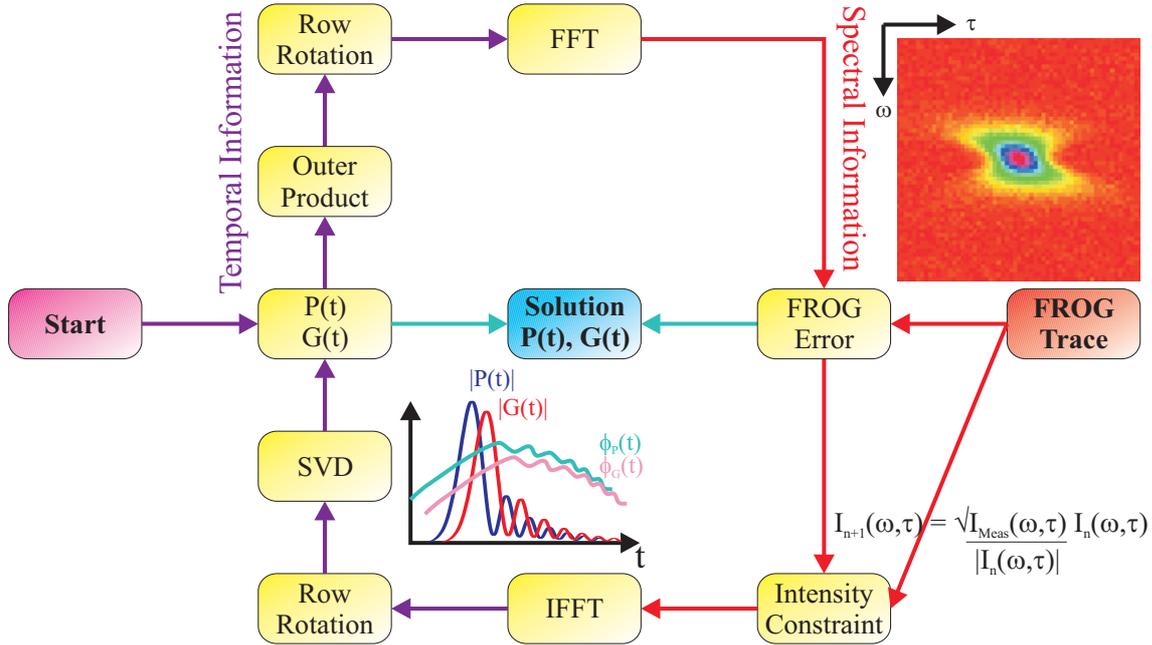


Figure 2.15: Schematic of the iterative FROG PCGPA algorithm. The algorithm is seeded with an initial guess for the pulse, $P(t)$ and gate, $G(t)$ fields, usually derived from noise or a previous reconstruction. The algorithm switches between the frequency and time domains, using generalised projections to minimise the FROG error. In the spectral domain, the magnitude of the FROG trace calculated from equation (2.51) is replaced by the measured FROG trace. As it is possible to represent the FROG trace in the time domain as the outer product of two vectors, the probe and gate pulses can be found from a singular-valued decomposition (SVD) of the calculated FROG trace in the time domain.

One deconvolution method for reconstructing the electric field from a FROG trace is called *principle component generalised projection algorithm* (PCGPA).^[99–101] Although other methods exist,^[102–104] the PCGPA tends to be widely used because it is independent of the type of FROG used as it reconstructs both the gate and probe pulse simultaneously and allows rapid convergence for real-time reconstruction. The algorithm is described schematically in figure 2.15, and relies on the fact that the FROG trace can be generated by the outer product of two vectors that represent the electric field of the probe, $P(t)$, and gate, $G(t)$, pulses. An important issue in measuring the spectrogram is that the temporal and spectral resolutions are constrained with respect to each other by the sampling theorem. This is typically not satisfied experimentally and it is necessary to perform interpolation on the measured data before processing it. In addition, it may also be necessary to filter the data (e.g. remove noise, background signal etc.).

The algorithm is seeded with initial guesses for the probe and gate fields. These can either be derived from uniform noise, a noisy Gaussian or even a previous reconstruction. The latter case improves the speed of the convergence provided the pulse being reconstructed only differs slightly from the seed pulse. However, it can be difficult for the routine to move out of a local minima if the pulse differs significantly and thus it can be better to use noise as this allows the algorithm to be re-executed for different seeds to help prevent stagnation.

The performance of the inversion routine is measured by the FROG error. This is defined as

$$\sigma_n^{FROG} = \frac{1}{N} \sqrt{\sum_{i=1}^N \sum_{j=1}^N [\tilde{I}_{Meas}(\omega_i, \tau_j) - \tilde{I}_n(\omega_i, \tau_j)]^2}. \quad (2.52)$$

The measured FROG trace is given by $\tilde{I}_{Meas}(\omega_i, \tau_j)$ and the n^{th} iteration of the reconstructed FROG trace is given by $\tilde{I}_n(\omega_i, \tau_j)$, where the subscripts i, j represent the pixel number. Typical values are 0.1 or lower. However, the definition of the FROG error used above depends on the sampling rate and not the span of the time or frequency domains. It is therefore difficult to compare the accuracy of the FROG routine to other methods (or even different FROG setups) unless both the sampling grid and the span of the time and frequency axes are known. Low values for the FROG error can only be obtained if the FROG trace corresponds to one calculated from a physical field, i.e. any systematic errors result in a high FROG error. However, the opposite is not true, that is to say that a high FROG error does not imply some systematic error. High FROG errors can also be a result of noise, or from stagnation at a local minimum of the iterative algorithm. As the reconstruction is a global routine, it can be difficult to determine the source of large FROG errors.

2.5.4 Chronocyclic tomography

Chronocyclic tomography can reconstruct the complete complex temporal or spectral electric field from a two-dimensional data set using a direct algorithm. The process derives its

name from the fact that it utilises the Wigner representation of pulses in the chronocyclic space. In fact, it is this distribution that is reconstructed; the temporal and spectral distributions are found from the marginals of the Wigner distribution.

The principle of the technique is as follows. Projections of the Wigner distribution at different angles are recorded by measuring the pulse spectrum for different rotations of the pulse in chronocyclic space. This is shown schematically in figure 2.16. Rotation of the Wigner distribution is achieved by application of both quadratic spectral phase and quadratic temporal phase to the probe pulse. If a sufficient number of rotations are made, corresponding to different amounts of quadratic phase, then it is possible to fully reconstruct the Wigner distribution via an inverse Radon transform.^[58]

The quadratic phase can be implemented via the use of a dispersive delay line. The temporal phase is applied via an electro-optic phase modulator driven sinusoidally at radio frequencies (RF). If the pulse is short compared to the RF period, and aligned with the peak (or trough) of the modulation, then the phase can be approximated as a quadratic modulation. However, it is clear that such an implementation is limited by the device

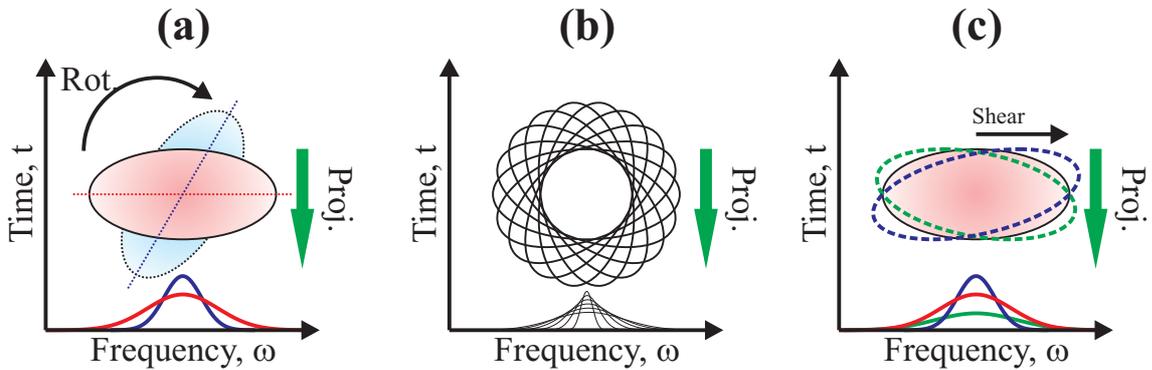


Figure 2.16: Schematic of chronocyclic tomography via rotations of the Wigner distribution of the pulse. (a) The original Wigner distribution in time-frequency space is shown as a red oval, with its projection on to the frequency axis is shown in red beneath. The rotated distribution, achieved by applying a quadratic temporal and spectral phase, is shown as a blue oval; its projection onto the frequency axis is shown in blue beneath. (b) Many rotations of the Wigner distribution and their projections onto the frequency axis: a complete and unambiguous reconstruction of the pulse is possible via an inverse Radon transform. (c) Two small shears of the Wigner distribution along the frequency axis, achieved with quadratic temporal phase only; the projections onto the frequency axis is shown in blue and green beneath, plus the original spectrum in red.

performances. For sub-picosecond pulses, the quadratic chirp would stretch the pulse to a duration outside the usable window of the phase modulator. Longer pulses are limited by the spectral resolution of the spectrometer. However, such a method has the advantage of being able to reconstruct fairly complicated pulses that the other techniques find difficult to measure (e.g. a double pulse pair) .

Simplified chronocyclic tomography

Chronocyclic tomography requires many rotations in order to reconstruct the Wigner function of the pulse in question. In addition, it is necessary to apply relatively large quadratic phases in both the temporal and spectral domains in order to achieve a large rotation. However, it is possible to simplify the process such that only two rotations, in addition to the original spectrum corresponding to no rotation, are needed.^[62] In fact, it is possible to fully reconstruct a coherent field by performing two small shears of the Wigner distribution along the spectral dimension. The concept is shown schematically in figure 2.16 (c). If the projection of a sheared distribution onto the spectral axis is given by

$$I_{\psi}(\omega) = \int_{-\infty}^{\infty} dt W(t, \omega - \psi t), \quad (2.53)$$

where ψ is the size of the temporal phase modulation such that $E(t) \rightarrow E(t) \exp(i\psi/2t^2)$, then it is possible to reconstruct the electric field by integration of

$$\begin{aligned} \frac{\partial I_{\psi}}{\partial \psi} &= -\frac{\partial}{\partial \omega} \left(I \frac{\partial \phi}{\partial \omega} \right) \\ &\simeq \frac{I_{+\psi} - I_{-\psi}}{2\psi}. \end{aligned} \quad (2.54)$$

Here $I_{+\psi}$ and $I_{-\psi}$ correspond to the projections of the Wigner distributions onto the spectral axis for a small shear, $\psi \ll 1$ in the positive or negative direction respectively. The original spectrum of the pulse with no modulation present is given by I .

This method has several advantages over other tomographic techniques: the method

only requires three spectral measurements; a small temporal phase modulation only is required; and the pulse can be sampled at the sampling limit. However, it is expected that this technique is not as robust to noise compared to interferometric or full tomographic techniques because the information here is encoded in small amplitude variations.

2.5.5 SPIDER

Spectral phase interferometry for direct electric-field reconstruction (SPIDER) is an interferometric, self-referencing technique for measuring the complete temporal characterisation of optical pulses.^[16,81,105–108] The concept is simple and allows extraction of the one-dimensional spectral phase of a pulse from a one dimensional data set and is based on *spectral shearing interferometry*, in analogy to lateral shearing interferometry (section 2.4.3) in the spatial domain.

In its conventional form, SPIDER samples the spectral phase of a test pulse by the interference in a spectrometer of two spectrally sheared replicas of the pulse: $\tilde{E}(\omega)$ and $\tilde{E}(\omega - \Omega)$, where Ω is the spectral shift (or spectral shear) between the two pulses. The spectral shear enables the spectral phase at one frequency to be referenced to the phase at a different frequency, separated by the shear. The spectrally sheared replicas are generated via upconversion of the test pulse with two different quasi-monochromatic pulses derived from a highly chirped copy of the test pulse. In order to extract the spectral phase information, the replicas are delayed in time with respect to each other by an amount τ . This delay imparts fringes with a periodicity of $2\pi/\tau$ onto the spectrum in the case of zero shear (i.e. $\Omega = 0$). If a shear between the two replicas exist, then the spectral phase between two frequencies separated by the shear is encoded as a variation in the fringe periodicity from its nominal spacing. The delay also allows the selection of two different frequencies from the same chirped pulse during upconversion (see section 2.5.5).

The SPIDER interferogram, $\widetilde{I}_{SP}(\omega)$, can be written as

$$\begin{aligned}\widetilde{I}_{SP}(\omega) &= |\widetilde{E}(\omega) + \widetilde{E}(\omega - \Omega) e^{i\omega\tau}|^2 \\ &= \widetilde{I}(\omega) + \widetilde{I}(\omega - \Omega) \\ &\quad + 2\sqrt{\widetilde{I}(\omega)\widetilde{I}(\omega - \Omega)} \cos\left[\widetilde{\phi}(\omega) - \widetilde{\phi}(\omega - \Omega) + \omega\tau\right]\end{aligned}\quad (2.55)$$

$$= \widetilde{I}^{\text{dc}}(\omega) + \widetilde{I}^{+\text{ac}}(\omega) e^{+i\omega\tau} + \widetilde{I}^{-\text{ac}}(\omega) e^{-i\omega\tau}, \quad (2.56)$$

where $\widetilde{I}(\omega)$ and $\widetilde{\phi}(\omega)$ are the spectral intensity and spectral phase of the test pulse. Equation (2.55) can be re-written in the form of background, $\widetilde{I}^{\text{dc}}(\omega) = \widetilde{I}(\omega) + \widetilde{I}(\omega - \Omega)$, and phase dependent, $\widetilde{I}^{\pm\text{ac}}(\omega) = \sqrt{\widetilde{I}(\omega)\widetilde{I}(\omega - \Omega)} \exp\left\{\pm i\left[\widetilde{\phi}(\omega) - \widetilde{\phi}(\omega - \Omega)\right]\right\}$, terms to give equation (2.56). The spectral phase can then be extracted via a Fourier transform filtering algorithm followed by integration or concatenation. The reconstruction process is shown by the flow chart in figure 2.17 and is described in detail below.

Provided the delay is large enough, typically several times the pulse duration, isolated sidebands are formed upon performing a Fourier transform on the interferogram to convert from frequency to the pseudo time domain, \tilde{t} :

$$\begin{aligned}I(\tilde{t}) &= \mathfrak{F}\left\{\widetilde{I}_{SP}(\omega); \omega \rightarrow \tilde{t}\right\} \\ &= I^{\text{dc}}(\tilde{t}) + I^{-\text{ac}}(\tilde{t} + \tau) + I^{+\text{ac}}(\tilde{t} - \tau),\end{aligned}\quad (2.57)$$

where $I^n(\tilde{t}) = \mathfrak{F}\left\{\widetilde{I}^n(\omega); \omega \rightarrow \tilde{t}\right\}$ is the Fourier transform of the individual terms in equation (2.56). Isolating one of the sidebands, performing an inverse Fourier transform and taking the argument allows one to extract the phase difference, $\widetilde{\theta}(\omega) = \widetilde{\phi}(\omega) - \widetilde{\phi}(\omega - \Omega)$ plus the delay phase, $\omega\tau$. In order to reconstruct the original phase, the delay phase needs to be calibrated and removed, which is achieved by performing the same procedure at zero shear.

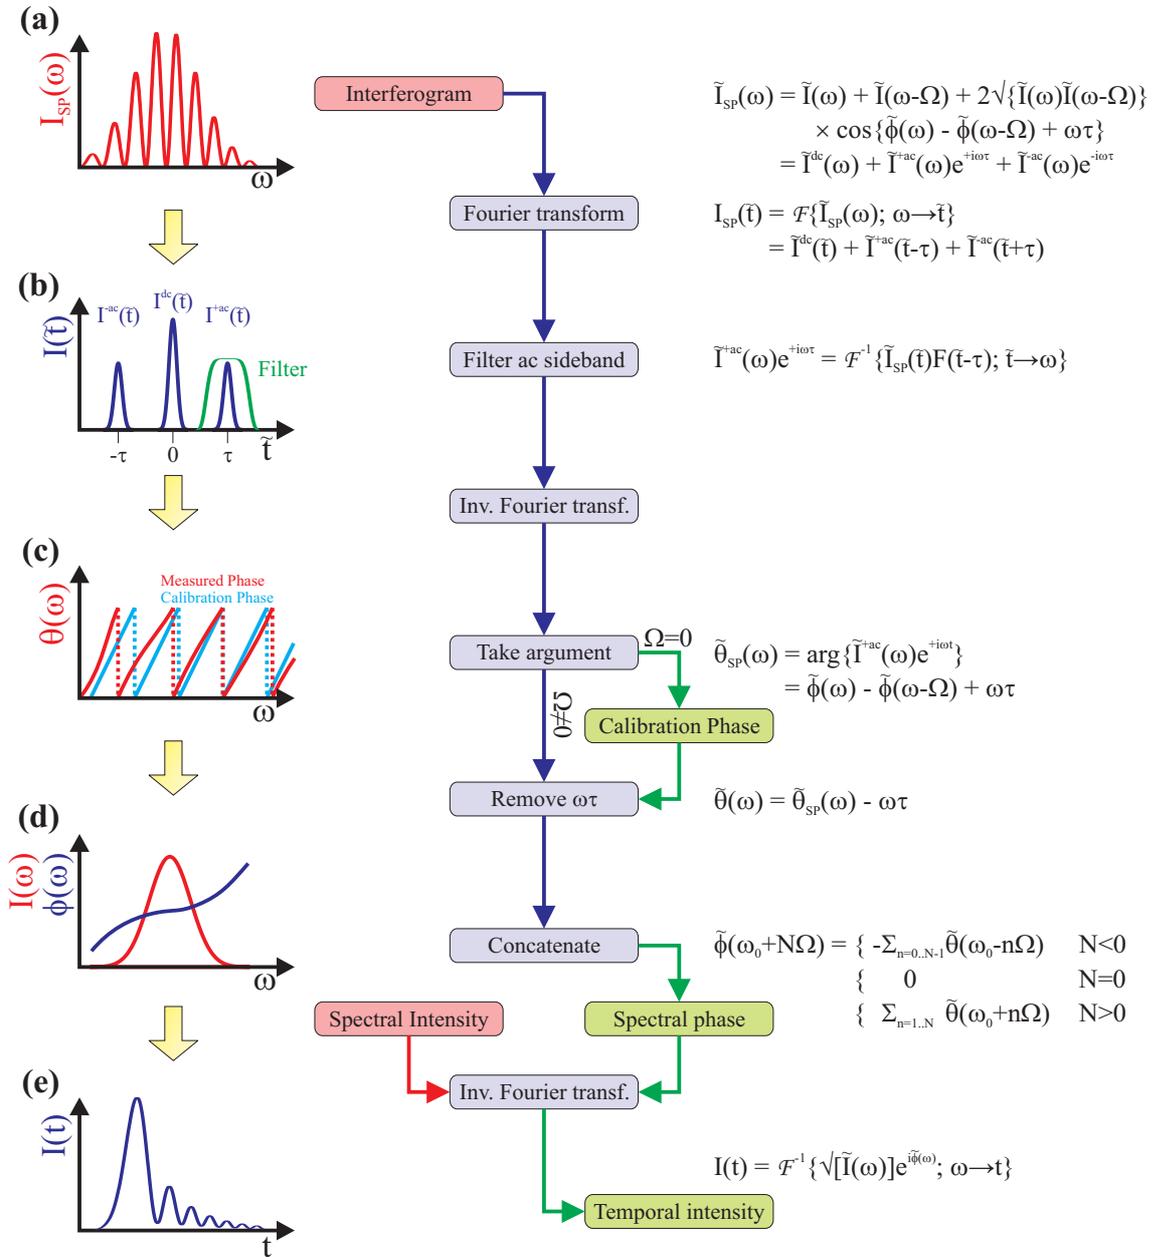


Figure 2.17: Flow chart of the SPIDER reconstruction algorithm. Left column: schematic of the main steps; central column: SPIDER flow chart (red = measured data, blue = operations on data, green = reconstructed data); and right column: mathematical algorithm. (a) The interferogram is Fourier transformed into the pseudo time domain. (b) One sideband, containing the phase information, is filtered and inverse Fourier transformed back. (c) The argument is taken, giving the phase difference plus the delay term. A calibration trace at zero shear is used to remove the delay term. (d) Concatenation results in the spectral phase of the probe pulse. (e) The temporal field is calculated from the reconstructed spectral phase and independently measured spectral intensity.

For a small shear, the spectral phase difference is approximately equal to the spectral phase gradient, or equivalently the chirp of the pulse:

$$\begin{aligned}\tilde{\theta}(\omega) &= \Omega \frac{\tilde{\phi}(\omega) - \tilde{\phi}(\omega - \Omega)}{\Omega} \\ &\simeq \Omega \frac{\partial \tilde{\phi}(\omega)}{\partial \omega}.\end{aligned}\tag{2.58}$$

Integration of the scaled phase difference, $\tilde{\theta}(\omega)/\Omega$, results in an approximation of the test-pulse phase, accurate for polynomial phase up to an order that is dependent on the integration method used. For example, trapezium integration is accurate for polynomial phase up to and including cubic order. It is useful to consider the phase difference as the spectral gradient of the test phase as this is equivalent to the chirp of the pulse. This is shown schematically by the dashed arrows in figure 2.18.

Alternatively, the phase can be reconstructed exactly via concatenation: the phase extracted from the SPIDER routine, after removal of the calibration phase, is equal to the phase difference between two frequencies separated by the spectral shear. Given the

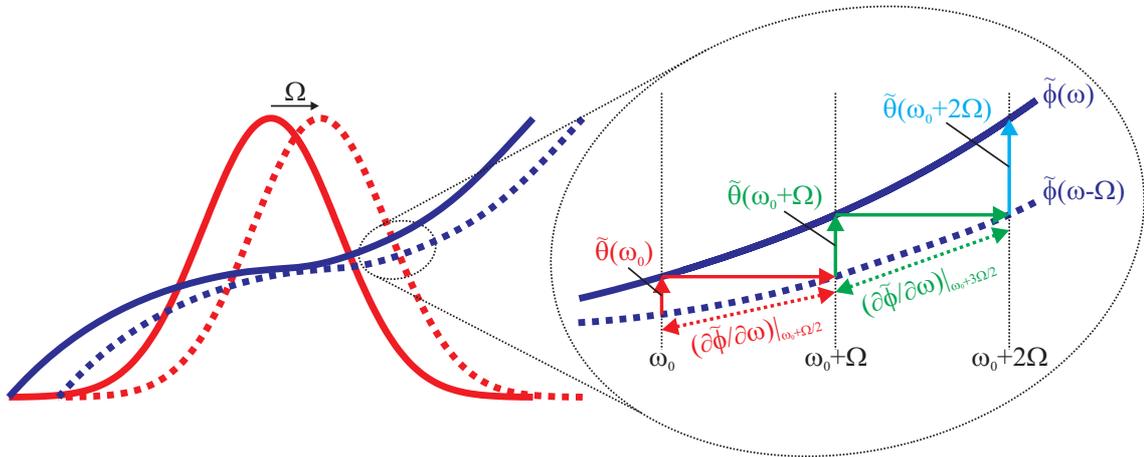


Figure 2.18: Concept of integration or concatenation of the phase difference extracted from a SPIDER interferogram. A measurement of the phase difference, $\theta(\omega) = \phi(\omega) - \phi(\omega - \Omega)$ is approximately equal to the phase gradient scaled by the spectral shear, $\tilde{\theta}(\omega) \simeq \Omega \left[\partial \tilde{\phi}(\omega_0 + \Omega/2) / \partial \omega \right]$ (dashed arrows). Alternatively, $\tilde{\phi}(\omega + \Omega) = \tilde{\phi}(\omega) + \tilde{\theta}(\omega + \Omega)$. Thus, so long as the phase is known at some value, ω_0 , it is possible to calculate the phase at any point along the sampling grid $\omega = \omega_0 + n\Omega$ (solid arrows).

phase at a reference frequency ω_0 , the phase at the frequency $(\omega_0 + \Omega)$ can be calculated according to:

$$\begin{aligned}\tilde{\phi}(\omega_0 + \Omega) &= \tilde{\phi}(\omega_0) + [\tilde{\phi}(\omega_0 + \Omega) - \tilde{\phi}(\omega_0)] \\ &= \tilde{\phi}(\omega_0) + \tilde{\theta}(\omega_0 + \Omega).\end{aligned}\tag{2.59}$$

The phase at $(\omega_0 + 2\Omega)$ can then be calculated from the phase at $(\omega_0 + \Omega)$:

$$\begin{aligned}\tilde{\phi}(\omega_0 + 2\Omega) &= \tilde{\phi}(\omega_0 + \Omega) + [\tilde{\phi}(\omega_0 + 2\Omega) - \tilde{\phi}(\omega_0 + \Omega)] \\ &= \tilde{\phi}(\omega_0) + \tilde{\theta}(\omega_0 + \Omega) + \tilde{\theta}(\omega_0 + 2\Omega),\end{aligned}\tag{2.60}$$

and so on, calculating the phase along the grid $\omega = \omega_0 + n\Omega$ (as illustrated by the solid arrows in figure 2.18). The same process can also be repeated for frequencies smaller than ω_0 . It is necessary to set the phase to some particular value at ω_0 ; typically this is set to zero. This means SPIDER cannot measure the absolute or linear spectral phase. Using the concatenation method, it is possible to reconstruct the spectral phase according to the following:

$$\tilde{\phi}(\omega_0 + N\Omega) = \tilde{\phi}(\omega_0) + \begin{cases} \sum_{n=1}^N \tilde{\theta}(\omega_0 + n\Omega) & \text{if } N \geq 1 \\ 0 & \text{if } N = 0 \\ \sum_{n=0}^{|N|-1} \tilde{\theta}(\omega_0 - n\Omega) & \text{if } N \leq -1. \end{cases}\tag{2.61}$$

From equation (2.61), it is clear that the SPIDER routine samples the spectral phase at intervals separated by the spectral shear. Thus the concatenation routine is exact provided the shear satisfies the Whittaker–Shannon sampling theorem: for a pulse with compact temporal support, T (i.e. one with its energy completely contained within a time window T), the spectrum needs to be sampled with a maximum angular frequency $2\pi/T$.

Thus the spectral shear must satisfy

$$\Omega \leq \frac{2\pi}{T}. \quad (2.62)$$

This only places an upper limit on the size of the spectral shear, or equivalently: given a spectral shear, the maximum pulse duration that can be measured is given by $T_{\max} = 2\pi/\Omega$. However, oversampling the spectral phase, i.e. using a smaller shear than is required according to equation (2.62), reduces the accuracy of the reconstruction due to noise and shot-to-shot instabilities (see section 3.1.1). As a conventional SPIDER device is set up for a fixed shear, there is a trade-off between the range of pulse durations that can be measured and the accuracy of the reconstruction.

Experimental implementation of SPIDER

A conventional implementation of SPIDER for visible to NIR pulses is shown in figure 2.19. The two time-delayed replicas of the probe pulse are generated either with an etalon or a dispersion balanced Michelson interferometer. It is necessary to ensure interferometric stability in the delay between the two replicas. Taking the front and back reflections from

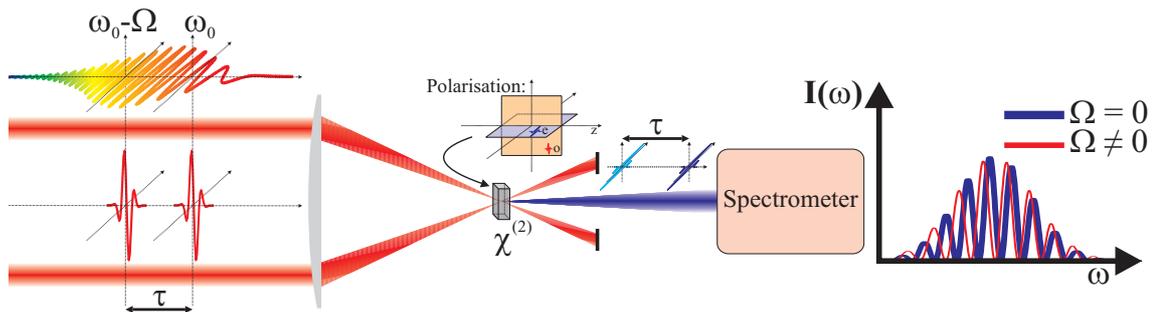


Figure 2.19: Schematic of a conventional SPIDER concept. Two time delayed replicas of the test pulse upconvert with a strongly chirped ancillary beam in a nonlinear crystal cut for $\chi^{(2)}$ nonlinear crystal cut for type II sum frequency generation. The test pulse replicas are aligned with the o-axis of the crystal and the chirped pulse with the e-axis in order to maximise the phase-matching bandwidth. During the interaction, the replicas each see a different quasi-monochromatic frequency from the chirped pulse, resulting in two frequency shifted replicas of the test pulse centred near twice the fundamental frequency. The two sheared replicas are interfered in a spectrometer. The time delay induces fringes onto the spectrum with a nominal periodicity of $2\pi/\tau$ (blue fringes). The spectral shear modulates the fringe spacing from the nominal periodic fringes (red fringes).

an etalon provides high stability.^[109] However, the pulse taken from the back reflection will experience some material dispersion compared to the front reflected pulse. This effect is small for many cycle pulses with a small bandwidth, so an etalon is sufficient. For few cycle pulses, the dispersion due to the etalon becomes significant; hence a dispersion balanced interferometer is often used.^[105] However, it is still possible to use an etalon for few cycle pulses, provided the difference in dispersion between the two ‘replicas’ is accurately calibrated and removed during the reconstruction.^[109] An interferometer also provides for an adjustable delay between the two replicas, providing the ability to vary the shear between the two signal pulses and therefore measure a range of different pulse durations accurately. However, the range of shears available will be limited because of the coupling between the shear and delay. The maximum shear possible is set by the spectral resolution, which limits the maximum delay possible. The ability to separate the sidebands in the Fourier domain determines the minimum delay possible and hence sets the minimum shear.

The chirped ancillary pulse can either be derived from the test pulse itself, or from an alternative laser. Deriving the chirped pulse from another laser has the advantage of allowing it to be at a different frequency to perform homodyne (H-) SPIDER,^[110] or higher intensity for modified (M-) SPIDER^[111]. Such a pulse could be derived from a CPA system before the compressor, hence will be both high power and highly chirped. If the pulse is not chirped, or is derived from the test pulse, then it is necessary to induce a chirp by temporally stretching it in a dispersive delay line. For few-cycle pulses, this can be achieved via a thick piece of dispersive glass (e.g. 5–10 cm of SF10, giving 10,000–20,000 fs²). For lower bandwidth pulses, much higher dispersion is required (several hundred thousand fs²) and so a ‘grating compressor’ is used to add a large amount of negative dispersion. The grating compressor is used because it is relatively easy to build, although other dispersive delay lines can be used. In either case, it is necessary to ensure that the dispersion is enough to stretch the pulse by a factor of 10–100 so that each test pulse interacts with a

quasi-monochromatic frequency in the nonlinear crystal. It is important to note that the shear is set by both the dispersion of the ancillary pulse and the delay between the two replicas according to

$$\Omega = -\frac{\tau}{\tilde{\phi}''}, \quad (2.63)$$

where $\tilde{\phi}''$ is the second order Taylor co-efficient of 2.37, representing the GDD of the phase.

The chirped and time delayed replicas are then focused into a $\chi^{(2)}$ nonlinear crystal, typically β -barium borate (BBO), such that they spatially and temporally overlap. The crystal is cut for type-II sum-frequency generation (SFG). This geometry maximises the bandwidth of the phase-matching function for the SFG process. This is because it is necessary to upconvert two closely separated frequencies from the chirped pulse, and all the bandwidth of the replica pulses. Thus the chirped pulse is aligned with the extraordinary (e) axis of the crystal, and the replicas with the ordinary (o) axis.

During SFG, each time-delayed replica will interact with a different quasi-monochromatic frequency from the chirped pulse. The test pulse replicas are thus shifted in frequency by different amounts, resulting in two time-delayed and spectrally-sheared replicas of the test pulse, with a central frequency close to twice that of the fundamental. The calibration can be performed by detecting the two time-delayed replicas at the fundamental frequency in the second order diffraction of a grating spectrometer. As there will not be any shear between the two replicas, the phase extracted from the Fourier filtering routine will be equal to $\omega\tau$. Measuring in the second order allows the measured calibration phase to be subtracted directly from the SPIDER phase, eliminating the need for direct spectral calibration.^[112]

2.6 Accuracy, precision and consistency

The ability of any metrology device, and the various reconstruction algorithms, to reconstruct the pulse being measured can be quantified by three parameters: accuracy, precision and consistency.^[113,114] Each term has a specific meaning; their definitions are listed below.

Accuracy is a measure of how closely the reconstructed field matches the actual *physical* field. Quantifying accuracy is a theoretical task based on simulations and modelling of the apparatus as best as possible.^[115,116] Loss of accuracy typically results due to distortions applied to the pulse within the apparatus, for example dispersion due to optics, as well as noise on the measured data.^[113]

The *precision* of a measurement is defined as the deviation of several estimates of the same quantity obtained from the same device. Such a measurement can only be attained experimentally and quantifies the ability of a device to reproduce the same data. In SPIDER, for example, fluctuations in the delay between the two pulses results in varying small amounts of quadratic phase that is added to the reconstructed spectral phase. It is often difficult to distinguish between the precision of a device from fluctuations of the test pulse itself. For inherently multi-shot devices, such as scanning FROG setups, fluctuations in the test pulse will decrease the precision of the device because the precision relates to differences in complete measurements.^[114]

Consistency is a measure of the correlation between the reconstructed field and the measured data. One can envisage this as the error between the measured data and data simulated from the reconstructed field, although it can be measured via different routes. Iterative routines tend to maximise the consistency by minimising the error between measured and simulated data. Many devices contain more data than there is information contained within the pulse. For example, FROG and tomography use two-dimensional data sets to reconstruct the one-dimensional field; SPIDER spectrally oversamples the pulse. This built in redundancy, allowing the possibility to generate multiple reconstructions of the same pulse from one data set, can also be used to estimate the consistency of the method. Consistency can only be calculated on a single measurement trace.^[114,116,117]

It is very important to note the difference in each definition. Often a device or its reconstruction is misquoted to have a particular accuracy when it is actually the consistency that has been determined. For example, the FROG error is a measure of the best fit

between a physically realisable field to the measured data. It does not, however, give any measure on how well that reconstructed field matches the input test pulse. In addition, one must also be aware of possible ambiguities in the reconstructed pulse. For example, it may not be possible to relate the phase between well separated spectral or temporal regions.^[118–120]

2.6.1 RMS field error

There exist a variety of different measures for quantifying the difference between two fields. Often this is the error measure used to define the convergence in an iterative algorithm. However, other than as a test for convergence, an error measure is useless if it is specific to a particular measurement technique and does not allow a comparison between methods (e.g. the FROG error). One universal error measure is based on the Wigner distribution.^[121] This is more useful than the FROG error as it allows a direct comparison between different techniques, since it should be possible to reconstruct the Wigner distribution from any complete characterisation device. However, routines to evaluate this can be complex and computationally inefficient. Another universal error measure is the root-mean-square (RMS) field error^[113]. The RMS field error is defined as the norm of the difference between two normalised fields, E_1 and E_2 :

$$\begin{aligned} \sigma_{\text{RMS}} &= \|E_1 - E_2\| \\ &= \left\{ \int_{-\infty}^{\infty} dt |E_1(t) - E_2(t)|^2 \right\}^{\frac{1}{2}} \\ &= \left\{ 2\pi \int_{-\infty}^{\infty} d\omega |\tilde{E}_1(\omega) - \tilde{E}_2(\omega)|^2 \right\}^{\frac{1}{2}}. \end{aligned} \quad (2.64)$$

As a consequence of Parseval's theorem, both representations result in equal values. In addition, it is possible to calculate the Wigner error from the RMS error.^[113]

2.7 Summary

The aim of ultrashort pulse metrology is to find the electric field of propagating electromagnetic pulses. The electric field can be represented as a continuous superposition of plane waves of different frequencies, wave vectors and their relative phases. Characterisation techniques strive to measure the amplitude and phase of these Fourier components in order to reconstruct the real electric field as a function of position and time.

There are a variety of approaches to pulse characterisation, each with very different methods of extracting the spectral or spatial phase; each has its own set of advantages and disadvantages, as well as different parameter spaces of operation. Thus it is necessary to consider the class of pulses one wishes to measure and then select the apparatus that best suites their requirements.

The two most common methods of complete temporal characterisation are second harmonic generation frequency resolved optical gating (SHG-FROG) and spectral phase interferometry for direct electric-field reconstruction (SPIDER). Both techniques have similar parameter ranges and have been proven to reliably reconstruct pulses from the NIR, visible and UV spectral regions. Having two completely different yet complimentary methods allows one to compare both techniques and gives confidence to the results when they agree, or to give indications of where the limitations are when they do not agree.

This chapter only offers a brief overview of the most common methods. Many varieties of each technique exist, allowing a wide assortment of different pulse types to be measured. Typical factors that are needed to be considered are the central wavelength, bandwidth, pulse duration, pulse-shape complexity and amount of space-time coupling. Often it is desirable to measure in a single shot, or to have online characterisation to allow real-time adjustment of parameters. Other times this is not necessary. Thus it is clear that one needs to consider the variety of pulses that one would like to measure and to use the most appropriate device. Often there may be more than one capable device allowing an independent test of the validity of their measurements. Even when this is not possible, both FROG and SPIDER have redundancy in their data that allows one to test the consistency of their measurements.

Complete characterisation of near infrared few-cycle pulses

Few-cycle pulses (e.g. a sub-10 fs pulse contains less than 3.75 cycles at 800 nm) have become an important tool in ultrafast science, with applications in pump-probe spectroscopy^[122], high harmonic generation (HHG)^[123] and other intense laser-matter interactions^[124,125]. Few-cycle pulses provide the shortest temporal resolution for pump-probe experiments for a given carrier frequency of the pulse, and achieve higher peak powers for a given pulse energy, which is useful for applications such as HHG. In addition, the peak electric field strength is reached with a minimal effect on the medium due to the leading edge of the pulse. Take HHG as an example, ionisation of the medium will cause defocusing and thus will reduce the peak intensity that can be achieved. This in turn will reduce the maximum harmonic that can be produced. In addition, depletion of the ground state from ionisation will also reduce the efficiency of the process. These effects can be minimised by preventing ionisation due to the leading edge of the pulse if the HHG process is driven by a few-cycle pulse, this is highlighted in figure 3.1. As ionisation effects in HHG occur at a relatively low ionisation fraction (of a few percent), it is clear that the peak intensity of the many cycle pulse will not be reached due to defocusing in the gas.

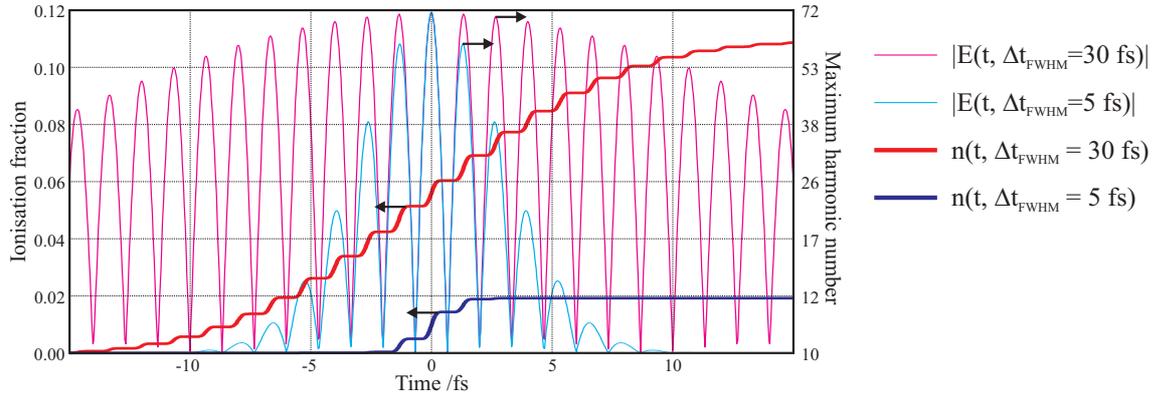


Figure 3.1: Fractional ionisation of Argon from two ultrashort laser pulses of different durations. Both pulses have a Gaussian envelope function but with a FWHM intensity duration of 30 and 5 fs — the electric field magnitude is drawn in magenta for the 30 fs pulse and cyan for the 5 fs pulse, the fractional ionisation, n , is drawn in red and blue for the 30 fs and 5 fs pulses respectively. The peak intensity for both pulses is $5 \times 10^{14} \text{ Wcm}^{-2}$, corresponding to a cut-off at harmonic 71 for Argon (11.2 nm). However, as ionisation of only a few percent can reduce the harmonic emission and peak intensity, due to ionisation induced defocusing, the longer pulse will not be able to produce the higher frequency photons.

Current short pulse technology is based around Ti:Sapph or fibre laser systems, resulting in a carrier frequency in the optical to near-infrared. As the pulse duration cannot be shorter than a single period of the carrier frequency, these pulses are limited to a duration of several femtoseconds. This also means that the bandwidth required to support such pulse durations is on the same order of magnitude as the actual carrier frequency (i.e. they have bandwidths spanning over an octave). In this chapter, I therefore concentrate on pulses with a carrier frequency near 2.35 rad/fs (800 nm) and total bandwidths greater than 1 rad/fs (300 nm), and thus sub-10 fs Fourier transform limited full width at half maximum pulse durations. However, many of the ideas presented in this chapter are also valid for few-cycle pulses with lower or higher carrier frequencies, and the techniques presented can often be transferred into different spectral ranges (e.g. into the ultraviolet or mid-infrared) by changing the optics and the nonlinear process.^[70,83,126,127]

The availability of few-cycle near-infrared pulses has increased dramatically over the past few years due to the rapid development in the sources used to generate them. These pulses may be generated directly out of an oscillator, aided by the development of chirped mirrors used for intracavity dispersion compensation^[128,129]. Alternatively, intense few-

cycle pulses can be generated via hollow fibre pulse compression (HFPC)^[130,131] or filamentation^[132,133] from a chirped pulse amplifier (CPA)^[134,135]; or directly from a non-collinear optical parametric amplifier (NOPA)^[136,137] or an optical parametric chirped pulse amplifier (OPCPA)^[138,139]. The choice of generation method depends on the parameters of a pulse that one requires, such as wavelength tunability, shot-to-shot stability, long term stability, peak power, average power, contrast ratio or spectral range, of which each method has its own distinct advantages and disadvantages, often resulting in a compromise. However, all these methods are constrained by the problems of dealing with extremely short, ultrabroadband pulses and the problem of measuring them.

The difficulty in measuring few-cycle pulses compared to many cycle pulses is that they often exhibit one or more of the following properties: (1) an octave spanning spectrum, (2) a highly modulated spectrum and (3) space-time coupling. Having such a large bandwidth causes problems with propagation, since dispersion due to air can cause significant temporal broadening and even propagation in vacuum results in space-time coupling (see chapter 2). These properties need to be taken into consideration when designing systems that utilise few-cycle pulses. The extremely short duration and thus relatively high peak power can cause problems such as damage to optics or nonlinear effects (e.g. self-phase modulation and Raman scattering). Metrology of few-cycle pulses becomes challenging or even impossible using conventional temporal characterisation techniques due to the properties that these pulses tend to exhibit. The desire to measure the pulse in a single shot, or to have online acquisition, complicates the matter even further. In this chapter, I will describe the problems of using conventional characterisation techniques and demonstrate a method of spectral interferometry that I have developed for measuring near-infrared few-cycle femtosecond pulses.

3.1 Metrology of few-cycle pulses

In the following, I shall only consider methods that are capable of complete electric field reconstruction, that is the pulse spectrum and phase (although the spectrum can be acquired separately if necessary). Thus I shall not consider autocorrelation as this does not recover the complete phase. Although this is theoretically possible for interferometric autocorrelation, many different pulse shapes lead to similar traces that poses difficulties for all but the most modest pulses (i.e. smooth spectra with zero phase), especially in the wings of the trace where the SNR is small. Chronocyclic tomography requires both a linear temporal and spectral chirp to rotate the Wigner distribution. Currently, the size of the temporal chirp required is too large to perform chronocyclic tomography for few-cycle pulses and thus I shall not consider this method. Simplified chronocyclic tomography has the advantage that it only requires a small quadratic temporal phase, resulting in a sheared Wigner distribution. This small chirp can be performed via an electro-optic modulator. However, such a technique is relatively new and beyond the scope of this thesis. Due to measuring small changes in the spectrum, this is also likely to suffer from poor signal to noise, especially in the wings of the spectrum, and thus may not prove suitable to few-cycle pulses. In the following, I shall therefore only concentrate on the FROG and SPIDER characterisation techniques.

All characterisation techniques require measurements of the pulse spectrum that satisfy the Whittaker–Shannon (WS) sampling limit. In this context, this means that a pulse that has compact support in a time window T (i.e. no intensity outside this window) requires its spectrum to be sampled at a rate that is less than the WS frequency, $\delta\omega_{\text{ws}} = 2\pi/T$. Thus it may not be possible to acquire the whole spectrum in a single shot if the method requires significant oversampling of the spectrum (e.g. SPIDER requires sampling of fringes in the interferogram at the Nyquist limit and thus typically oversamples for the WS limit by a factor of 5–10 or more). The spectral response of the technique is also an issue; neither the spectral response of current detectors or the phase-matching properties of the nonlinear

crystals are flat over such a bandwidth. Some techniques (e.g. SPIDER) are insensitive to the spectral response, whereas others (e.g. FROG) require this to be calibrated in order to reconstruct the spectral phase accurately. However, all methods require a measurement of the pulse spectrum (which is either contained within the measurement data or taken independently) in order to perform a temporal reconstruction. The spectrum must be intensity calibrated. This intensity calibration results in a non-uniform SNR across the spectrum, which can cause problems in global reconstruction methods such as FROG.

A pulse that exhibits a modulated spectrum — generally true in the case of HCFPC or filamentation schemes — can become very problematic to measure, depending on the characterisation method used. In SPIDER, it is necessary to be able to isolate the ‘ac’ sideband in the Fourier domain, which requires the period of the fringes in the interferogram to be much smaller than any spectral features. Thus it is almost impossible to obtain the complete interferogram in a single shot. For example, measuring a 2.8 fs pulse with conventional SPIDER required multiple shots of the interferogram and then ‘stitching’ it together.^[140] As FROG does not need to oversample the spectrum, this is not normally an issue.

Alternatively, there may be regions of low spectral intensity. However, the method still needs to be able to recover the phase over these regions. This can be a problem in methods that are sensitive to noise or that have a low dynamic range (i.e. the range of intensities that can be detected). If the SNR is too small to measure the spectral phase over any regions with low spectral intensity, then it is not possible to reconstruct the temporal pulse accurately. Since the spectral intensity is low within these regions, it is expected that this should have negligible effect in the temporal domain. Examples of inaccuracies include the inability to measure sub-pulses, fine temporal structure or long pedestals. In these cases, it is normally necessary to use alternative methods (e.g. third harmonic autocorrelation is very useful for identifying low intensity sub-pulses or pedestals due to its high dynamic range)^[141] or even a combination of methods (e.g. FROG or SPIDER combined with spectral interferometry)^[142].

There are two different scenarios that can arise when the spectrum contains regions of low spectral intensity: (1) the frequency span of the low spectral intensity region is less than the maximum allowed sampling frequency and (2) when the frequency span is larger than the maximum possible sampling frequency. These scenarios are shown in figure 3.2.

If the region of low spectral intensity has a frequency span that is less than $\delta\omega_{\text{WS}}$, then it is not necessary to sample the phase within those regions. In SPIDER, this is achieved by setting the spectral shear equal to or close to the maximum sampling frequency, $\delta\omega_{\text{WS}}$ such that the sampling grid does not fall within the region of low spectral intensity, as illustrated in figure 3.2 (a). FROG will recover the phase in this case with no difficulty or ambiguity.

If there exists well separated spectral regions, such that the spectral gap is greater than the maximum allowed sampling frequency, then it is necessary to measure the spectral phase of each individual region *and* to be able to relate the phase of each individual region

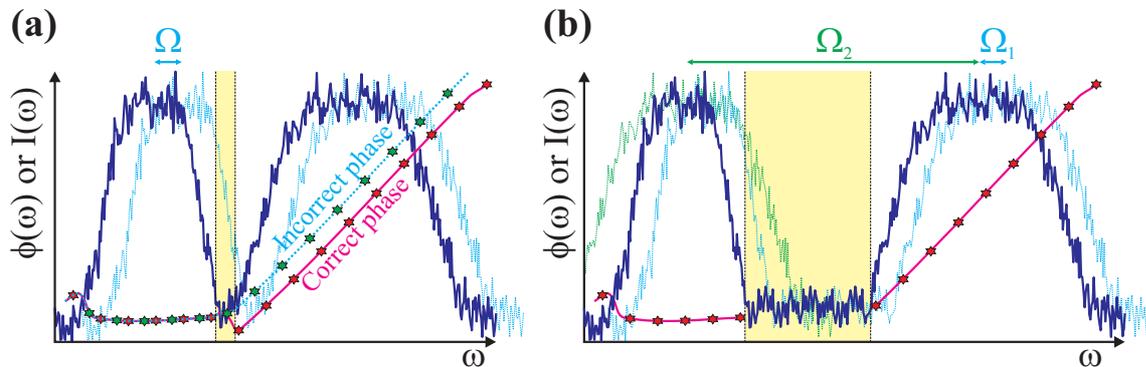


Figure 3.2: Reconstructing the spectral phase over regions of low spectral intensity with SPIDER. (a) The region of low spectral intensity (shaded yellow region) spans a frequency range that is less than the maximum shear allowed by the WS sampling limit. It is only necessary to ensure that the sampling grid (shown by the stars) steps over the shaded region. The reconstructed spectral phase is shown for two scenarios: (1) the sampling grid falls within the shaded region (green) resulting in an incorrect phase difference between the separated regions; (2) the sampling grid steps over the shaded region, ensuring the correct phase difference between the two regions is measured. (b) The region of low spectral intensity (shaded yellow region) is larger than the maximum possible spectral shear. Each separated region is measured with a shear that satisfies the WS sampling limit for each spectral region (Ω_1). The phase difference between the two regions is measured separately with a large shear (Ω_2). The reconstructed phase is undefined within the shaded region if the SNR is too low to allow spectral phase reconstruction, resulting in an inaccuracy in the reconstructed temporal profile. If the energy contained within this spectral region is small, then the inaccuracy of the reconstructed pulse shape will be negligible.

to each other. In SPIDER, this requires two measurements. The first measurement is obtained with a shear equal to or less than the WS limit (i.e. $\Omega \leq 2\pi/T$) so that the spectral phase of each individual region is effectively measured independently. The second measurement is obtained with a shear that spans the spectral gap. This then relates the phase between the two regions, as shown in figure 3.2 (b). However, it is not possible to relate the spectral phase between well separated spectral regions in conventional FROG. This scenario would require a cross-correlation FROG (X-FROG).^[84]

Few-cycle pulses are more prone to space-time coupling compared to multicycle pulses — the octave spanning bandwidth can result in a spatially dependent carrier frequency and pulse duration.^[143] In addition, it becomes very difficult to align pulse-shaping devices that utilise space-time coupling (e.g. prism/grating stretchers/compressors and zero-dispersion lines) to eliminate any space-time coupling completely. This is especially true for certain generation methods (e.g. OPCPA) that require spatio-temporal shaping over the complete spectrum of the pulse for their operation.

The effect that STC has on the characterisation device depends on the specific implementation used. For example, single-shot versions of FROG utilise the spatial profile of the pulse to map the time delay to spatial position. If the pulse shape varies as a function of position, then this will clearly distort the measured trace. In some versions, it is possible to use this distortion to measure simple forms of STC (e.g. linear spatial chirp or linear angular dispersion/pulse front tilt).^[95,96] However, this is only true if it is possible to assume some functional form of the STC. Due to the nonlinear nature of generating few-cycle pulses, it is likely that the STC may take on non-simple forms, which would ideally require a method that can measure arbitrary STC.

SPIDER is more resilient to STC because it encodes the 1D spectral phase information in a 1D experimental trace, ensuring that the temporal profile is measured at one spatial position only. If one would like to characterise the STC, then this is possible by spatially resolving the interferogram, thus giving the temporal profile for every position across one

spatial dimension. As the phase between different spatial positions is not measured, it is not possible to measure the spatial phase (i.e. it is possible to add a phase that is an arbitrary function of space only without changing the SPIDER interferogram). In order to measure the complete space-time pulse, it is necessary to measure both the spectral and spatial phase gradient. This is possible by performing both spatial and spectral shearing interferometry simultaneously.^[144]

3.1.1 Errors in SPIDER pulse reconstructions

Due to the fidelity of few-cycle pulses, extra care and attention is required when measuring them. As SPIDER is an interferometric technique, it clearly requires good stability to enable accurate reconstruction. Thus I shall formulate how the reconstructed spectral phase is modified in the presence of shot-to-shot fluctuations in a pulse train. These fluctuations manifest themselves in different forms, either as a change in the delay between the two signal pulses used to generate the SPIDER interferogram, or as a change in the pulse being measured. I shall describe these situations separately.

Calibration of the spectral shear

An error in the calibration of the spectral shear, often as a result of an error in the value of the delay or the amount of chirp in the ancillary pulses, effectively results in a scaling of the magnitude of the reconstructed phase. This can easily be checked by adding known phase to the pulse being measured and comparing the spectral phase of the pulse with and without the addition of the known phase. Any scaling of the measured phase compared to the known phase is likely to be from an improper calibration of the shear. Using this data, it is possible to obtain an accurate measurement of the shear.

Calibration of the SPIDER delay

It is well known that, due to the effective integration of the extracted phase from the Fourier filtering routine to obtain the pulse's spectral phase, any error in the calibra-

tion of the delay introduced in SPIDER leads to an extra erroneous chirp on the reconstructed pulse. The exact value of this chirp depends on the SPIDER parameters and the reconstruction algorithm used (i.e. whether by integration or concatenation) but can be found approximately by considering the integration of the extracted phase difference, $\tilde{\theta}(\omega) = \tilde{\phi}(\omega) - \tilde{\phi}(\omega - \Omega)$:

$$\begin{aligned}\delta\tilde{\phi}(\omega) &= \int_{\omega_0}^{\omega} d\omega' \frac{\delta\tau\omega'}{\Omega} \\ &= \frac{\delta\tau}{2\Omega} (\omega - \omega_0)^2,\end{aligned}\tag{3.1}$$

where $\delta\tilde{\phi}$ is the error in the phase and $\delta\tau$ is the error in the delay, such that $\tau_{\text{true}} = \tau_{\text{meas}} + \delta\tau$. It is important to note that the best method of removing the delay term, $\omega\tau$, is not to measure the value of the delay, τ , and then remove a linear spectral phase. It is much better to subtract the measured calibration phase directly, i.e.:

$$\tilde{\theta}(\omega) = \arg \left[I_{\text{SPIDER}}^f \times \left(I_{\text{Cal}}^f \right)^* \right],\tag{3.2}$$

where I^f is the filtered sideband in frequency space, i.e. $\tilde{I}^{\pm ac}(\omega) e^{\pm i\omega\tau}$ (equation (2.56)), and the subscripts ‘SPIDER’ and ‘Cal’ refer to the interferograms with and without a spectral shear respectively. A common technique is to perform the calibration with the two time delayed replicas (i.e. at the fundamental wavelength) and measure the SPIDER signal in the second order diffraction of the spectrometer, thus ensuring a one to one correspondence between the frequency axis of the SPIDER and calibration phases. This then avoids any issues arising from an improper calibration of the spectral axis of the spectrometer.^[112,145] However, if the value of the delay used for the calibration differs from the delay used for the SPIDER measurement, either from adjustment of the experimental setup or from any other sources of error, then there will be an error in the reconstructed pulse. It is therefore critical that any change in the value of the delay requires the device to be re-calibrated. It is possible to record the calibration and SPIDER interferograms simultaneously, thus simplifying the acquisition.^[146]

The effect that this miscalibration has on the pulse depends on the pulse shape and how far from Fourier transform limited (FTL) it is. For example, a pulse with a large time bandwidth product (TBP) must already have significant spectral phase and hence a small error in the reconstructed phase will have an insignificant effect on the pulse spectrum. In contrast, a few-cycle FTL pulse may only require a small phase error to temporally stretch the pulse by a significant proportion due to the extremely large bandwidth of the pulse. As the size of the temporal stretch resulting from a chirp of the pulse is dependent on the pulse shape, it is not possible to give a general result. However, the approximate size of the error in the pulse duration resulting from an error in the calibration of the delay can be estimated via Gaussian pulse analysis. In the following, the temporal and spectral intensities will be given by

$$I(t) = Ae^{-\alpha\left(\frac{t-t_0}{\Delta t}\right)^2} \quad (3.3)$$

$$\tilde{I}(\omega) = \tilde{A}e^{-\alpha\left(\frac{\omega-\omega_0}{\Delta\omega}\right)^2}, \quad (3.4)$$

where Δt and $\Delta\omega$ are the temporal and spectral widths and α is a scaling factor that specifies which width is being used. For example, if $\alpha = 4 \ln(2)$, then the widths specified in equations (3.3) and (3.4) correspond to the full width at half maximum intensity (FWHM). If $\alpha = 2$ then the widths correspond to the half-width at e^{-2} of the maximum intensity (HW).

The electric field is given by $E = \sqrt{I}e^{i\phi}$ where ϕ is a temporal/spectral phase. For a spectral intensity given by equation (3.4), the Fourier transform limited (FTL) temporal intensity is given by

$$\begin{aligned} I_{\text{FTL}}(t) &= \left[\mathfrak{F} \left\{ \sqrt{\tilde{I}(\omega)}; \omega \rightarrow t \right\} \right]^2 \\ &= Ae^{-\alpha\left(\frac{t-t_0}{\alpha/\Delta\omega}\right)^2}. \end{aligned} \quad (3.5)$$

Thus, by comparison of equation (3.5) and equation (3.3), the FTL time bandwidth product for Gaussian pulses is given by:

$$\Delta\omega\Delta t_{\text{FTL}} = \alpha. \quad (3.6)$$

If the pulse given by equation (3.4) has a quadratic phase (i.e. linear chirp) given by

$$\tilde{\phi}_{\text{chirp}}(\omega) = \frac{\phi''}{2}(\omega - \omega_0)^2 \quad (3.7)$$

then the temporal field will be given by

$$\begin{aligned} E_{\text{chirp}}(t) &= \mathfrak{F} \left\{ \sqrt{\tilde{I}(\omega)} e^{i\tilde{\phi}_{\text{chirp}}(\omega)}; \omega \rightarrow t \right\} \\ &= A' \exp \left[-\frac{\alpha}{2} \left(\frac{t - t_0}{\alpha\beta/\Delta\omega} \right)^2 \right] \exp \left[i\omega_0 t - i\frac{\phi''}{2} \left(\frac{t - t_0}{\alpha\beta/\Delta\omega^2} \right)^2 + i\Phi_0 \right], \end{aligned} \quad (3.8)$$

where A' is a scaling factor, $\Phi_0 = [\tan^{-1}(\phi''\Delta\omega^2/\alpha)]/2$ is an absolute phase and

$$\begin{aligned} \beta^2 &= 1 + \left(\frac{\phi''\Delta\omega^2}{\alpha} \right)^2 \\ &= 1 + \left(\frac{\Delta t'_{\text{chirp}}}{\Delta t_{\text{FTL}}} \right)^2. \end{aligned} \quad (3.9)$$

$\Delta t'_{\text{chirp}} = \phi''\Delta\omega$ is thus the pulse duration for a large chirp. In general, the chirped pulse duration is given by

$$\Delta t_{\text{chirp}} = \beta\Delta t_{\text{FTL}} \quad (3.10)$$

as is evident from equations (3.3), (3.8) and (3.9). This can be considered as a general result, in the sense that the chirped pulse duration will be given by (3.10) where the value of β depends on both the chirp, spectral shape and bandwidth and the scaling is given by equation (3.9)

A pulse with a GDD ϕ'' has a pulse duration Δt_{true} . However, the value of the GDD measured from a SPIDER device with an error in the calibration is given by $\phi'' + \delta\tau/\Omega$, resulting in a pulse duration $\Delta t_{\text{meas}} = \Delta t_{\text{true}} + \delta t$. Substitution for the phase into equation (3.10), taking the difference of the squares of the corresponding pulse durations, solving for the error in the estimated pulse duration, δt , and dividing by the ‘true’ pulse duration yields the relative error in the pulse duration:

$$\begin{aligned}\epsilon_t &= \frac{\delta t}{\Delta t_{\text{true}}} \\ &= \sqrt{1 + \left(\frac{N\epsilon_\tau}{\beta}\right)^2 \left(1 + \frac{2}{\epsilon_\phi}\right)} - 1.\end{aligned}\quad (3.11)$$

The value of β is defined in equation (3.9) for the ‘true’ phase.

$$N = \frac{\Delta\omega}{\Omega} \quad (3.12)$$

is the number of sampling points across the bandwidth, $\Delta\omega$, and ϵ_τ is the error in the delay relative to the *Fourier transform limited* pulse duration:

$$\epsilon_\tau = \frac{\delta\tau}{\Delta t_{\text{FTL}}} \quad (3.13)$$

The relative phase error (i.e. the ratio of the chirp due to the miscalibration and the intrinsic chirp on the pulse) is defined as

$$\epsilon_\phi = \frac{\delta\phi''}{\phi''}, \quad (3.14)$$

where $\delta\phi'' = \delta\tau/\Omega$. Since it is difficult to determine directly how the fractional error in the estimated pulse duration scales relative to the bandwidth of the pulse being measured from equation (3.11), I shall make some approximations to (3.11) for a few different cases.

Highly chirped pulses: if the pulse being measured is highly chirped, such that $\phi'' \gg \alpha\Delta\omega^2$, then the chirped pulse duration is given by

$$\Delta t_{\text{chirp}} \simeq \phi'' \Delta\omega. \quad (3.15)$$

Note that the limit to the approximation is equivalent to when the approximated pulse duration (3.15) is much larger than the FTL pulse duration. Provided that the extra chirp introduced due to the miscalibration remains small compared to the chirp of the pulse, then the fractional error in the estimated pulse duration is equal to the extra chirp relative to the actual chirp:

$$\epsilon_t \simeq \epsilon_\phi. \quad (3.16)$$

Near Fourier transform limited pulses: in the case when the pulse duration is close to its FTL, then the extra chirp introduced to the reconstructed spectral phase can be large enough to introduce an error in the estimated pulse duration that is as large as the pulse being measured. In the worst case scenario, which occurs when the laser system is optimised to produce the shortest pulse duration possible based on SPIDER measurements, the extra chirp introduced in the pulse reconstruction will be of opposite sign and of the same magnitude as the actual chirp on the pulse, i.e. $\phi'' = -\delta\tau/\Omega$. The measured pulse duration will be equal to the FTL pulse duration, or at least very close to it. The ‘true’ and measured pulse durations will be given by

$$\Delta t_{\text{true}} = \frac{\alpha}{\Delta\omega} \sqrt{1 + \left(\frac{\delta\tau\Delta\omega^2}{2\Omega}\right)^2} \quad (3.17)$$

$$\Delta t_{\text{meas}} \simeq \frac{\alpha}{\Delta\omega} \quad (3.18)$$

Substituting for the true and measured pulse duration gives $\beta = 1 + (N\epsilon_\tau)^2$ and $\epsilon_\phi = -1$, therefore

$$\epsilon_t = 1 - \frac{1}{\sqrt{1 + (N\epsilon_\tau)^2}}. \quad (3.19)$$

Equation (3.19) would imply that as the error becomes exceptionally large, the relative error approaches unity. Although correct, this does not highlight the gross error that occurred in the measurement. As one has the ability to vary the pulse duration as one wishes, it is expected that the true pulse duration should be as close to the FTL pulse duration as possible. It can therefore make much more sense to measure the error in the pulse duration relative to the FTL pulse duration. This is then a measure of the error between the true and measured pulse durations, relative to the desired pulse duration. In this case, the relative pulse error is

$$\begin{aligned}\epsilon'_t &= \frac{\delta t}{\Delta t_{\text{FTL}}} \\ &= \sqrt{1 + (N\epsilon_\tau)^2} - 1.\end{aligned}\tag{3.20}$$

In the case when $\epsilon_\tau \ll 2N^{-1}$:

$$\epsilon_t \simeq \epsilon'_t \simeq \frac{(N\epsilon_\tau)^2}{2}.\tag{3.21}$$

However, if the error is large, such that $\epsilon_\tau \gg 2N^{-1}$ then the relative error approaches

$$\epsilon'_t \simeq N\epsilon_\tau.\tag{3.22}$$

The exact value of the relative error and the different approximations are plotted in figure 3.3.

From figure 3.3, it is clear that it is necessary to use a shear as large as possible to lower the constraints on the error in the calibration. However, few-cycle pulses tend not to have a Gaussian spectrum or temporal intensity, and the phase is typically oscillatory. The net result is that in the temporal domain, the oscillations in the wings of the pulse can extend to tens or even hundreds of times the FWHM. If it is necessary to measure these low intensity oscillations, then a much finer sampling grid is needed, leading to a sampling number of $N = 30\text{--}50$. However, it is expected that aliasing of the low intensity wings should not affect the FWHM value significantly. Therefore, if one is

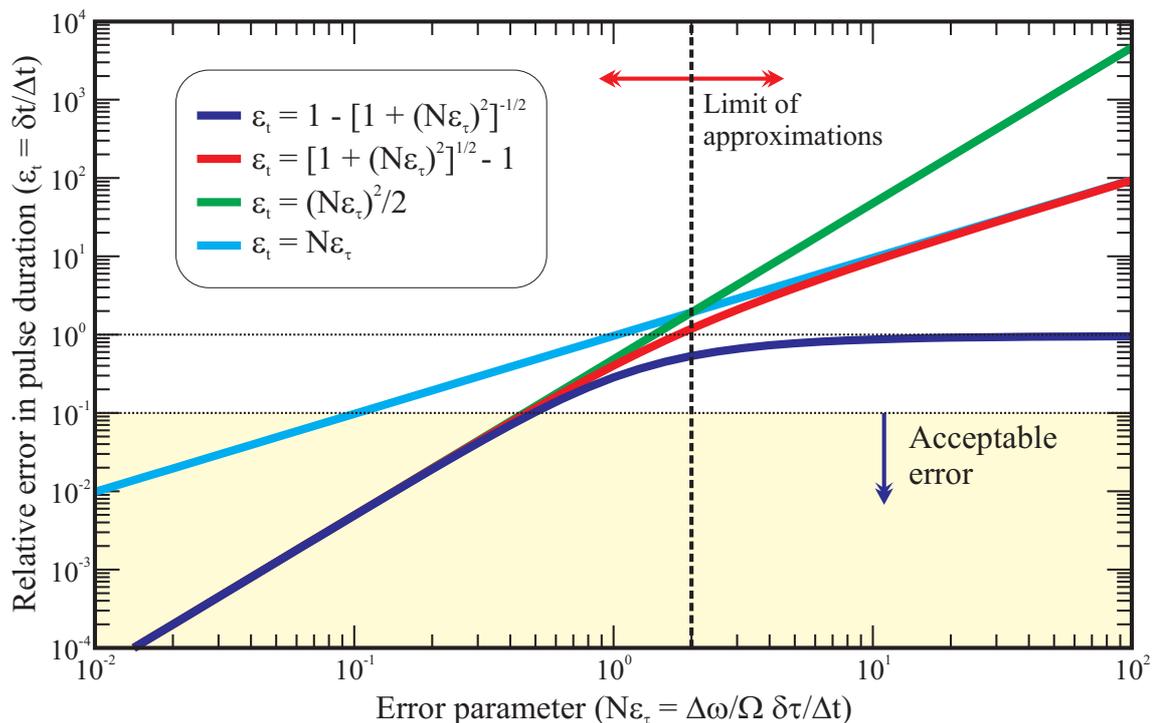


Figure 3.3: Error in the reconstructed pulse duration due to an error in the calibration of the SPIDER delay. The measured pulse is assumed to be FTL (i.e. the laser setup is optimised to give the shortest pulse according to the SPIDER measurement). The error, ϵ_t , in the measured duration relative to the FTL (red) or true (blue) duration is plotted on a logarithmic scale in x and y as a function of $N\epsilon_\tau = \delta\tau\Delta\omega^2/\alpha\Omega$. The shaded region corresponds to an acceptable error. Approximations to the FTL result are also plotted for $\epsilon_\tau \ll 2N^{-1}$ (green) and $\epsilon_\tau \gg 2N^{-1}$ (cyan).

only interested in ensuring minimal GDD, it is sufficient to measure with a small time window that encompasses the majority of the pulse energy. If, however, one is interested in the relatively low intensity oscillations, a finer sampling grid is required. In this case, it is necessary to perform the calibration accurately and either immediately before the SPIDER measurement, or simultaneously to minimise any effects due to long term drift. Alternatively, one may perform two measurements, with a small and large shear, to verify the GDD if this is possible. It is also necessary to note that sampling on such a fine grid over a large bandwidth will cause problems for any device. For spectrographic techniques, this will translate to sampling with fine temporal and spectral resolutions over large temporal and spectral windows. Not only does this result in an excessively large amount of data (scaling as N^2) that can cause problems in the retrieval algorithm, it is difficult to ensure

linearity of the delay and the required accuracy in the calibration of the frequency axis over that range.

As an example, I consider two particular conventional SPIDER measurements. Gallmann *et al.* first adapted conventional SPIDER for measuring few-cycle pulses.^[105] In this setup, sub-6 fs pulses were measured using a shear of $\Omega = 30$ mrad/fs ($T = 210$ fs). Using a bandwidth of $\Delta\omega \simeq 0.55$ rad/fs FWHM for a 5 fs pulse width (FWHM) gives $N \sim 20$. Therefore for 10% accuracy in the reconstruction, this requires an error in the delay of $\epsilon_\tau \lesssim 2\%$ or $\delta\tau \lesssim 100$ as. For an interferometer, this is not particularly special as interferometric stability requires

$$\delta\tau \ll \frac{\lambda_0}{2c}. \quad (3.23)$$

Assuming an interferometer has a stability of one tenth of the above yields $\delta\tau \sim 130$ as at 800 nm. Therefore the shot-to-shot stability of the interferometer will ensure accurate reconstruction. However, drift in the interferometer would require recalibration over time, or active stabilisation. Yamane *et al.* have a much more stringent requirement on the stability constraints, as they measured a pulse with a duration of 2.6 fs (in the single cycle regime).^[40] In this case, a shear of $\Omega = 26$ mrad/fs ($T = 240$ fs) was used to cover a bandwidth of $\Delta\omega \simeq 1.06$ rad/fs, giving $N \simeq 40$. This leads to a relative error of $\epsilon \simeq 1\%$ ($\delta\tau \simeq 26$ as). This is still relatively simple to obtain, but is near the limit of standard laboratory interferometers. In both cases, the time window used to plot the pulse in the temporal domain was much smaller than the time window given by the shear, which was approximately 200 fs. However Gallmann^[105] plotted the pulse over 180 fs (although 100 fs may have been sufficient) and Matsubara^[40] plotted their pulse over 80 fs. This implies a much coarser sampling grid was possible. However, it is likely that the device may have been set up to measure a wide range of pulse durations.

Shot-to-shot fluctuations and space-time coupling

A SPIDER interferogram that has been averaged over a large spatial region of a pulse that exhibits space-time coupling is equivalent to an average of multiple interferograms, each generated with different pulses. This is because each spatial position can be considered as giving rise to a single SPIDER interferogram that the integrating detector merely sums over. As will be shown, the net effect is that the reconstructed phase is the average of the ensemble of spectral phases and is independent of any spectral intensity changes.

Phase fluctuations: Considering an ensemble of pulses, whether each individual pulse is a different instance in time or space, in all cases the spectral phase can be written as

$$\tilde{\phi}_n(\omega) = \langle\langle \tilde{\phi}(\omega) \rangle\rangle + \delta\tilde{\phi}_n(\omega), \quad (3.24)$$

where $\langle\langle \tilde{\phi} \rangle\rangle$ is the average phase for the ensemble and $\delta\tilde{\phi}_n(\omega)$ is the difference between the spectral phase for pulse n and the ensemble average. Thus each pulse can be written in the form of $\tilde{E}_n(\omega) = |\tilde{E}(\omega)| \exp \left\{ i \left[\langle\langle \tilde{\phi}(\omega) \rangle\rangle + \delta\tilde{\phi}_n(\omega) \right] \right\}$. Note that equation (3.24) does not imply any relationship between each pulse within the ensemble (i.e. we are not restricted to a statistical distribution). Substituting equation (3.24) into and summing over each interferogram given by equation (2.55) results in

$$\begin{aligned} \langle\langle \tilde{I}_{\text{SP}}(\omega) \rangle\rangle &= \tilde{I}(\omega) + \tilde{I}(\omega - \Omega) + \left\{ \sqrt{\tilde{I}(\omega) \tilde{I}(\omega - \Omega)} e^{i[\langle\langle \tilde{\phi}(\omega) \rangle\rangle - \langle\langle \tilde{\phi}(\omega - \Omega) \rangle\rangle - \omega\tau]} \right. \\ &\quad \left. \langle\langle e^{i[\delta\tilde{\phi}_n(\omega) - \delta\tilde{\phi}_n(\omega - \Omega)]} \rangle\rangle + \text{c.c.} \right\}. \end{aligned} \quad (3.25)$$

Provided $\delta\tilde{\phi}_n(\omega) \ll 1$, the exponential inside the ensemble average can be expanded as

$$\begin{aligned} \langle\langle e^{i[\delta\tilde{\phi}_n(\omega) - \delta\tilde{\phi}_n(\omega - \Omega)]} \rangle\rangle &= 1 + i\langle\langle \delta\tilde{\phi}_n(\omega) \rangle\rangle - i\langle\langle \delta\tilde{\phi}_n(\omega - \Omega) \rangle\rangle + \langle\langle \text{O}(\delta\tilde{\phi}_n^2) \rangle\rangle \\ &\simeq 1. \end{aligned} \quad (3.26)$$

In calculating the summation, I have used the fact that the summation over all phase deviations is zero by the definition given in equation (3.24) and that the product of the phase deviations is much less than unity. Thus equation (3.25) reduces to the interferogram given by equation (2.55) for the average phase. Figure 3.4 shows the reconstruction for the average of 256 interferograms, for GDD and TOD with a mean value of 50 fs^2 and -300 fs^3 respectively. Both multiplicative and additive Gaussian noise was added to the values of the GDD and TOD with a standard deviation of 20%. The phase reconstructed from the measured SPIDER signal (i.e. the signal averaged over the individual interferograms) matches the average of the spectral phase from each individual pulse.

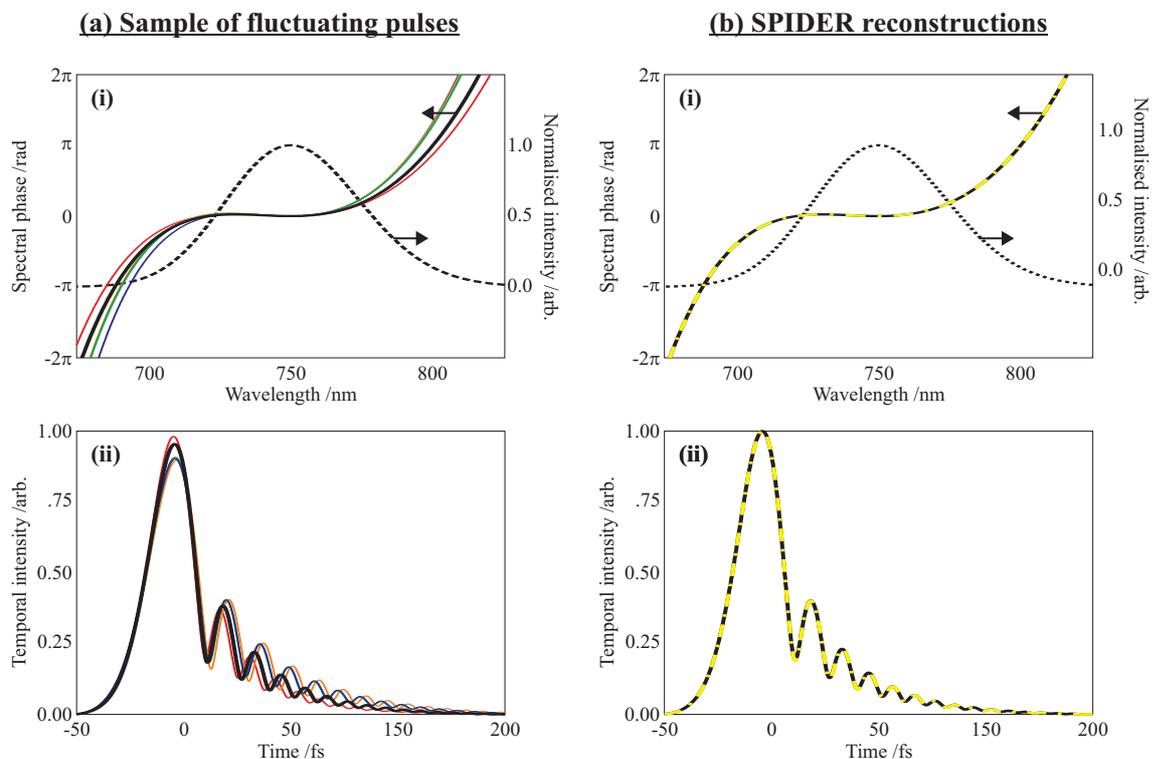


Figure 3.4: SPIDER pulse reconstruction averaged over many fluctuating pulses. (a) A sample from 256 pulses with fluctuating phase. (i) Spectral intensity (dashed black), spectral phase of ensemble average (solid black) and a sample from the ensemble of independent spectral phases (coloured). (ii) Temporal intensity of the ensemble average (solid black) and a sample of the temporal intensities from each individual pulse in the ensemble (coloured). (b) Comparison of the ensemble average to the reconstructed spectral phase and temporal intensity from the interferogram averaged over all individual interferograms. (i) The spectral intensity (dashed black), spectral phase of the ensemble average (solid black) and the reconstructed spectral phase from the averaged interferogram (yellow). (ii) Temporal intensity corresponding to the ensemble average (black) and the reconstructed temporal intensity from the averaged interferogram.

Spectrum only: It can be shown, using the same analysis as for the phase, that fluctuations in the pulse spectrum between different shots or spatial positions do not change the phase that is reconstructed. For example, consider a set of pulses with varying spectral intensity:

$$\tilde{E}_n(\omega) = \left[\langle\langle |\tilde{E}(\omega)| \rangle\rangle + \delta\tilde{E}_n(\omega) \right] e^{i\tilde{\phi}(\omega)}, \quad (3.27)$$

where $\delta\tilde{E}_n$ is a real function representing the change in the pulse spectrum for each pulse, and the triangular brackets represents the spectrum of the ensemble average. Then the SPIDER interferogram is given by

$$\tilde{I}_{\text{SP}}(\omega) = \sum_{n=1}^N \left\{ \tilde{I}_n(\omega) + \tilde{I}_n(\omega - \Omega) + 2\sqrt{\tilde{I}_n(\omega)\tilde{I}_n(\omega - \Omega)} \cos \left[\tilde{\phi}(\omega) - \tilde{\phi}(\omega - \Omega) - \omega\tau \right] \right\}, \quad (3.28)$$

where $\tilde{I}_n = |\tilde{E}_n|^2$. It is clear from equation (3.28) that averaging over a changing spectrum does not affect the reconstructed phase whatsoever. Thus, it is only changes in the spectral phase that may change the reconstructed pulse shape. This is because the spectral phase is encoded via a change in the fringe modulation, which is independent of the spectrum. The spectrum only causes problems in reconstructing the spectral phase if (1) the spectral intensity falls below the noise floor, resulting in an undefined phase or (2) the spectrum is highly modulated, preventing isolation of the AC and DC sidebands in the Fourier domain.

3.1.2 Carrier-envelope phase measurement in few-cycle pulses

It is not possible to discuss the metrology of few-cycle pulses without reference to the measurement of the absolute phase of the pulse, as this can play a crucial role in many experiments which utilise few-cycle pulses, such as attosecond pulse generation via HHG. Often, it is required to be able to stabilise the CEP on a shot-to-shot basis and thus schemes to do so are incorporated directly into the few-cycle laser independent of the spectral phase of the pulse. In addition, it is not possible to measure the absolute or

carrier-envelope phase (CEP) of a pulse directly using self-referencing techniques, but it is possible to measure variations of the absolute phase. Thus the following discussion on CEP measurement and control is split into two sections: (1) stabilisation of the absolute phase and (2) measurement of the absolute value of the CEP.

It is necessary to point out that measurement and control of CEP is not possible in any complete characterisation device discussed in this thesis. This is not a problem as the CEP is normally measured and controlled independently to the spectral chirp. The f -to- $2f$ interferometer discussed below could be incorporated into a SPIDER device, however this is generally not practical or beneficial. In addition, the CEP value needs to be known at the location of the experiment, i.e. *in situ*, which would require careful calibration of all the optics and path lengths between the measurement location and the experiment location. The accuracy in these measurements required to obtain an accurate value of the absolute phase is not usually possible. It therefore makes sense to utilise the experiment itself, which is sensitive to CEP, to actually measure its absolute value.

Carrier-envelope phase stabilisation

The most common method for CEP stabilisation is to measure variations in the CEP using an f -to- $2f$ interferometer and then feedback to either the power of the oscillator pump^[147], or to adjust the cavity length and group delay of the oscillator^[148]. The former method is usually used for fast-feedback using an AOM. The latter method is better suited to a slow-loop feedback scheme where it might be necessary to compensate large variations in CEP. It is usually necessary to perform feedback with both a fast and slow loop to ensure long term stability. The fast-loop is used to eliminate rapid changes in CEP phase due to fluctuations in, for example, the pump power, cavity alignment and cavity length. However, longer term drift, e.g. from changes in ambient temperature, are controlled by the slow-loop feedback. This is crucial for amplifier systems, for example.

In an f -to- $2f$ interferometer, an octave spanning pulse is upconverted via second harmonic generation and interfered with the fundamental. The pulse to be measured can be

spectrally broadened in a photonic-crystal fibre or sapphire plate if its frequency spectrum does not already span an octave. Fluctuations in the CEP can be measured by monitoring the shift in the position of the fringes resulting from the interference of the second harmonic of the long wavelengths with the short wavelengths of the fundamental pulse in a spectrometer.^[149] This method is usually used for a slow-loop feedback scheme on lasers with a repetition rate in the kilohertz range or lower. Alternatively, the two frequencies can be heterodyned in a photodiode from a multi-megahertz repetition rate laser.^[148,150,151] As the photodiode will measure a pulse train, the frequency spectrum consists of a comb of evenly spaced frequency components. The n^{th} frequency component is given by $f_n = nf_{\text{rep}} + \delta$, where f_{rep} is the repetition rate of the laser and δ is the frequency offset of the zeroth frequency component from DC. A frequency at $2n$ is given by $f_{2n} = 2nf_{\text{rep}} + \delta$. Thus a beat frequency equal to the absolute phase offset, $2f_n - f_{2n} = \delta$, between the second harmonic and fundamental pulses will be detected. By referencing this beat frequency to a master clock, usually derived from the pulse train itself, this can be used to directly generate a feedback signal electronically. As it is possible to perform the processing of the signal electronically, this scheme is used in combination with an AOM in fast-loop feedback schemes.

Carrier-envelope phase measurement

The schemes described above only measure the absolute value of the spectral phase. As the CEP depends on the pulse shape, i.e. complete spectral phase and spectral intensity, this is insufficient information to obtain the value of the CEP itself. In addition, the absolute phase is typically measured at a location which does not coincide with the experiment. The value of CEP can be measured via nonlinear interactions, as these depend on the electric field, rather than the field envelope. The first such experiments measured the total photoelectron emission from a gold surface as a function of the CEP.^[152] More recent experiments tend to be based on terahertz emission,^[153] the asymmetry of photoelectron emission along the direction of laser polarisation from above threshold ionisation

(ATI),^[154,155] or the spectral dependence of half-cycle cut-offs in high harmonic generation (HHG)^[149,156]. Since most experiments which make use of intense few-cycle pulses are based on HHG and photoionisation of electrons, these latter two methods prove extremely useful as they measure the CEP at the location of the experiment, thus eliminating any ambiguity due to propagation. The value of the CEP is typically varied by rotation of a thin glass plate inserted into the beam path.

3.2 Spatially encoded arrangement for SPIDER

3.2.1 Concept

In the preceding sections, I highlighted the problems of measuring pulses that exhibit space-time coupling with conventional characterisation techniques. Due to the nature of the acquisition and reconstruction routine, SPIDER has the advantage over the other techniques because it enables one to measure STC directly simply by spatially resolving the interferogram. In conventional SPIDER however, it is necessary to significantly oversample the spectrum according to the WS sampling theorem in order to resolve the fringes caused by the temporal delay. This becomes more significant for few-cycle pulses where the spectrum can become highly modulated, and/or have regions of low spectral intensity. As a result, it is necessary to increase the fringe periodicity to ensure there is no overlap between the DC and AC sidebands in the Fourier domain during the reconstruction. In addition, the domain of the spectrometer is approximately linear in wavelength whereas the phase accumulated due to the delay is linear in frequency, resulting in a change in the periodicity of the fringes on the interferogram due to the delay. Thus, the sidebands become extended even further in the Fourier domain. One can either increase the delay, and hence the spectrometer resolution, or interpolate the data onto a linear frequency grid. The latter choice is more practical, but can lead to artifacts in the reconstruction due to noise.

I also highlighted in the previous section that it is necessary to ensure the delay be-

tween the two pulses is calibrated accurately. However, due to long term drift in the setup (e.g. ambient temperature swings and beam pointing), it may be necessary to calibrate the delay for every measurement for few-cycle pulses. There exist a variety of adaptations to SPIDER that make this possible. For example, the calibration spectrum can be acquired with the SPIDER interferogram simultaneously^[146], or the delay can be calibrated out via H-SPIDER^[127]. Alternatively, the delay can be removed completely, as in spatially encoded arrangement (SEA-) for SPIDER^[157] or two dimensional spectral shearing interferometry (2DSI)^[158].

SEA-SPIDER modulates the spectrum in the spatial domain by introducing a tilt between the two signal pulses at zero delay, instead of introducing a delay between two collinear signal pulses. Thus the SEA-SPIDER interferogram can be represented by

$$\begin{aligned} \tilde{I}_{\text{SEA}}(\omega, x) = & |\tilde{E}(\omega, x)|^2 + |\tilde{E}(\omega - \Omega, x)|^2 \\ & + 2|\tilde{E}(\omega, x)||\tilde{E}(\omega - \Omega, x)| \cos \left[\tilde{\phi}(\omega, x) - \tilde{\phi}(\omega - \Omega, x) + Kx \right], \end{aligned} \quad (3.29)$$

where x is the spatial position across the beam and K is the difference in the mean wave vector between the two signal arms. The final term inside the cosine is the tilt between the two pulses and, strictly speaking, is frequency dependent: $Kx = (\mathbf{k}_1 - \mathbf{k}_2) \cdot \mathbf{x}$. For extremely large bandwidths, the frequency dependence of K becomes significant and needs to be taken into account via the removal of a calibration phase recorded with zero shear. The phase matching properties of the nonlinear crystal used for SFG can result in a slight additional frequency dependent tilt, which is also removed via the calibration. However, this term is unlikely to be significant as I use thin crystals (sub-50 μm thickness).

The phase of the test pulse is encoded in the spectral domain via the difference between two spectrally sheared replicas as in conventional SPIDER, except the fringes are now in the spatial dimension. As a result, it is possible to measure the spectrum at the sampling limit, i.e. $\Omega \simeq \delta\omega$, where $\delta\omega$ is the resolution of the spectrometer. Thus it is possible to characterise single-cycle pulses using a low resolution spectrometer and still

be able to capture the whole spectrum in a single shot. In practice, the spectrum is slightly oversampled to allow redundancy in the data to be used for consistency and precision checks,^[113,114] as well as allowing a wider range of pulses to be measured with little adjustment of the device.

The two spectrally sheared, tilted signal pulses (SP1 & SP2) are generated via sum frequency generation (SFG) of the test pulse with the two spatially tilted, time delayed and highly chirped ancillary pulses (figure 3.5). During the SFG, the test pulse will interact with two different quasi-monochromatic frequencies (CP1 & CP2), resulting in two spectrally sheared replicas centred near twice the central frequency of the test pulse. As both upconverted pulses are derived from the same test pulse, they will have zero delay between them. Due to conservation of momentum, the two sheared pulses will travel in different directions, allowing them to be spatially separated and then imaged onto the entrance slit of a 2D imaging spectrometer (2D-IS). SEA-SPIDER was first demonstrated by Kosik *et al*^[157] for many cycle pulses. This demonstration was only a proof-of-principle experiment and did not fully demonstrate the advantages of the technique and the apparatus used was not suitable for characterising few cycle pulses. In the following sections, I describe fully how I have further developed the idea to measure few-cycle pulses and detail many features of the device which were not previously demonstrated.

3.2.2 Intuitive interferograms

A common problem with conventional SPIDER is that the measured trace does not show any intuitive information directly, thus it can be difficult to determine where there are difficulties that cause errors in the reconstruction. Although methods exist which can show the precision and consistency of the measurement,^[113,114] and thus show where errors occur, further information and systematic errors are easier to spot and solve if they can be identified directly in the measured data. In SHG-FROG, for example, any asymmetry in the measured trace can usually be attributed to misalignment of the device or because the pulse exhibits STC. As a FROG trace is essentially a spectrogram, it shows

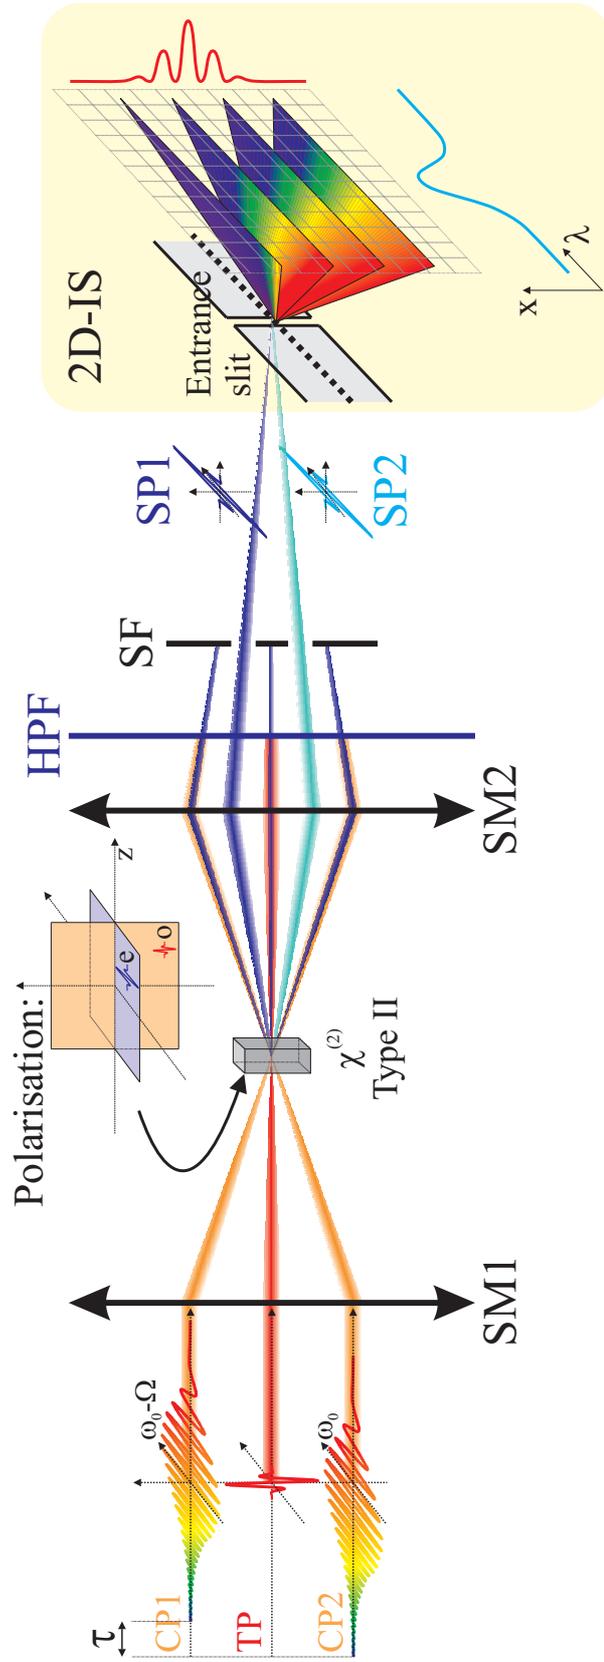


Figure 3.5: SEA-SPIDER concept. CP1 & CP2: chirped ancillary pulses 1 & 2; TP: test pulse; τ : delay between CP1 & CP2; SM1 & SM2: spherical mirror 1 & 2; $\chi^{(2)}$: nonlinear crystal cut for type II — CP1 & CP2 are aligned along the e-axis and TP aligned along the o-axis; HPF: high pass (frequency) filter; SF: spatial filter; SP1 & SP2: signal pulse 1 & 2; 2D-IS: 2D imaging spectrometer — the entrance slit is used to select a slice of the beam orthogonal to the spectral axis; x: spatial co-ordinate; λ : spectral co-ordinate.

intuitive nature, highlighting the temporal/spectral nature of the pulse and thus allowing optimisation without the need to reconstruct the pulse.

SEA-SPIDER interferograms are quite intuitive however. Simple structure can be used to both highlight regions where errors in the reconstruction occur and to allow optimisation of pulse parameters without reconstruction. If the spectral phase is independent of position (i.e. the pulse does not exhibit any space-time coupling in its phase), then the phase difference will be spatially independent. As a result, the contours of the fringes will directly map this phase difference, which is approximately proportional to the chirp of the test pulse. This is evident if the interferogram is written as

$$\tilde{I}(\omega, x) = \tilde{I}^{dc}(\omega, x) + 2\tilde{I}^{ac}(\omega, x) \cos \left[\tilde{\theta}(\omega) + Kx \right], \quad (3.30)$$

where the phase difference is given by

$$\tilde{\theta}(\omega) = \tilde{\phi}(\omega) - \tilde{\phi}(\omega - \Omega) \quad (3.31)$$

$$\simeq \Omega \frac{\partial \tilde{\phi}(\omega - \Omega/2)}{\partial \omega}. \quad (3.32)$$

The position of maximum fringe intensity is then given by $\tilde{\theta}(\omega) + Kx = 2\pi n$, where n is an integer. Taking $n = 0$ gives the position of the zeroth fringe maximum:

$$x_0 = -\frac{\Omega}{K} \frac{\partial \tilde{\phi}(\omega - \Omega/2)}{\partial \omega}. \quad (3.33)$$

Thus the position of the spatial fringes will be shifted by an amount that is proportional to the chirp of the pulse. This behaviour is demonstrated by the four images shown in figure 3.6, each corresponding to a different chirp on the test pulse. In each interferogram, the contour of the fringes, calculated via a 1D phase extraction routine (section 3.2.3), is plotted across the centre of the interferogram as a dashed black line. The shear is set to zero in (a), resulting in perfectly flat fringes. The interferogram in (b) corresponds to

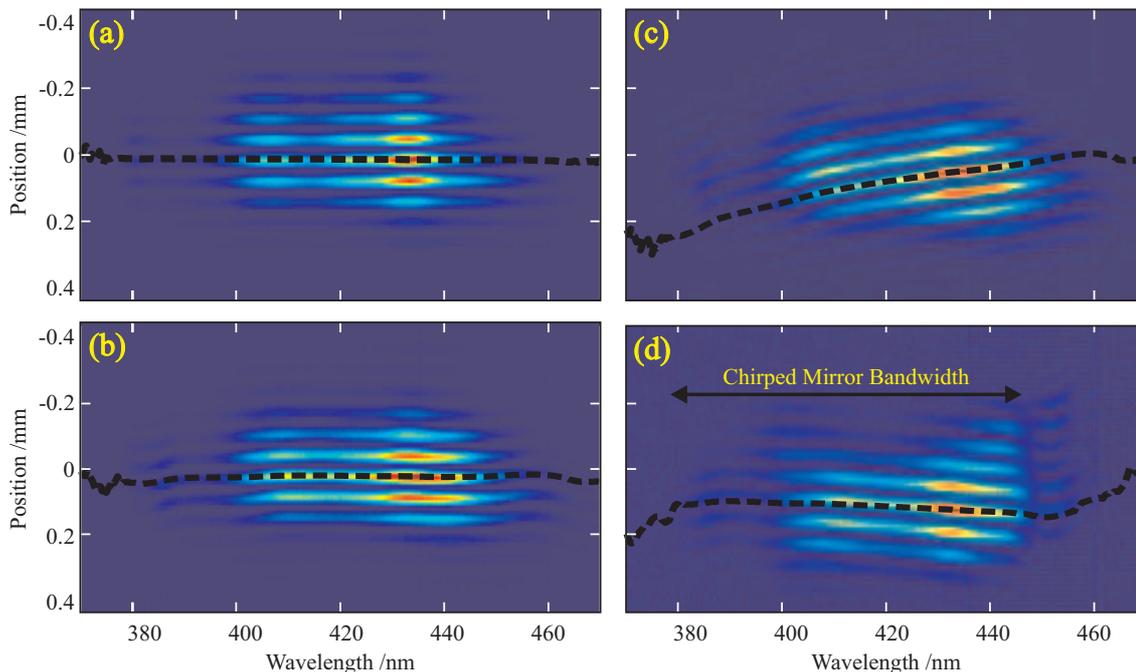


Figure 3.6: A set of SEA-SPIDER interferograms showing their intuitive nature. In each case, the chirp (i.e. the spectral phase gradient) is shown as a dashed black line across the centre, calculated using a 1D FFT extraction routine. (a) Calibration trace (i.e. $\Omega = 0$), resulting in perfectly flat fringes. (b) Measurement of a near FTL few-cycle pulse. The fringes are nearly flat, corresponding to no net GDD. The slight modulation is due to higher order phase. (c) Pulse with a large GDD ($+400 \text{ fs}^2$) resulting in positive tilt. The fringes are curved due to TOD and FOD from the fused silica used to introduce the dispersion. (d) Pulse with a large negative GDD (-120 fs^2) resulting in fringes with a negative tilt. There is also a discontinuity in the fringes at 450 nm due to the limited bandwidth of the chirped mirrors used to generate the negative dispersion.

a shear of 70 mrad/fs , suitable for measuring a time window of 90 fs . The fringes remain approximately flat, showing that there is no significant chirp on the pulse. The slight undulation of the fringes is due to higher order phase. The interferograms in (c) and (d) correspond to a large positive ($+400 \text{ fs}^2$) and negative (-120 fs^2) GDD respectively. The fringes clearly tilt in proportion to the chirp of the pulses, the tilt in (c) being about three times that in (d). As the positive GDD was introduced via 1 cm of fused silica, there is a slight curvature to the fringes due to the third and fourth order phase, in addition to the phase of the laser pulse itself.

Not only do the interferograms directly show the chirp of the pulse, but they can also highlight regions where the reconstruction may not be accurate. For example, the negative

GDD applied in figure 3.6 (d) was via a pair of chirped mirrors. However, the bandwidth of these mirrors was less than the spectrum of the pulse, which can clearly be seen by the discontinuity in the fringes at 450 nm. It may therefore be possible that there is a phase jump in this region that would normally not be taken into account. In order to check this, it is possible to scan the shear and observe the change in the fringes as a function of the shear to ensure that any phase jumps are taken into account. Such phase jumps are normally the cause of problems in conventional SPIDER.

It is also possible to gauge information on the amount of STC directly from the interferogram. In the absence of STC, and only a minor dependence of the tilt, K , on wavelength, then the fringes will have a frequency independent periodicity. However, any phase dependent STC will result in a frequency dependent fringe periodicity. Thus it is possible to decide whether it is necessary to perform the full 2D reconstruction algorithm, or if the quick 1D analysis is sufficient.

3.2.3 SEA-SPIDER reconstruction algorithms

2D FFT filtering routine

The SEA-SPIDER reconstruction algorithms are slightly modified versions of the conventional SPIDER algorithms, owing to the two dimensional data set. Instead of performing a one dimensional fast Fourier transform (FFT) filtering routine, one performs a 2D-FFT and 2D filter. The sidebands are formed in the pseudo- (\tilde{k}, \tilde{x}) space, with the AC sidebands being shifted to $\tilde{k} \pm K$. Hence, provided the spatial fringe periodicity is sufficiently large, it is possible to isolate the sideband regardless of the pulse spectrum being measured. Compare this for conventional SPIDER, where it may not be possible to isolate the sidebands for modulated spectra with significant phase, resulting in an overlap between the AC and DC terms. Note however, that in all versions of SPIDER it is necessary to ensure that the test pulse upconverts with a quasi-monochromatic frequency from the chirped pulse. After performing a 2D filter in the Fourier domain, an inverse 2D-FFT is

performed on the filtered data, and the phase extracted:

$$\tilde{\theta}(\omega) = \arg \left\{ \mathfrak{F}_2^{-1} \left[\mathfrak{F}_2 \left[\tilde{I}_{\text{SEA}}(\omega, x) \right] f(\tilde{t}, \tilde{k}) \right] \right\} \quad (3.34)$$

The calibration phase can be removed by multiplication of the filtered SPIDER data by the conjugate of the filtered calibration data before taking the argument, similar to equation (3.2). The extracted phase difference is then integrated or concatenated along the frequency domain as for conventional SPIDER, except that this is performed for every spatial position. Note that it is necessary to preserve the absolute phase difference between spatial positions as this encodes the PFT. In conventional SPIDER, a constant is removed from the extracted phase difference to ensure the reconstructed pulse in time is centred at zero. In SEA-SPIDER, it is also possible to remove a constant from the whole 2D phase difference but it is not possible to remove a different constant for each spatial position otherwise the PFT will be removed.

The 2D FFT routine has a lower SNR requirement compared to the 1D FFT routine used in conventional SPIDER because the 2D filtering has better noise discrimination and redundancy in the measured data. Even though it is necessary to acquire much more data than for conventional SPIDER, it is still possible to acquire and reconstruct the data in real time and even acquire data at kHz repetition rates with a suitable camera (e.g. 512×128 pixels at 8-bits is sufficient for a wide range of pulse shapes).

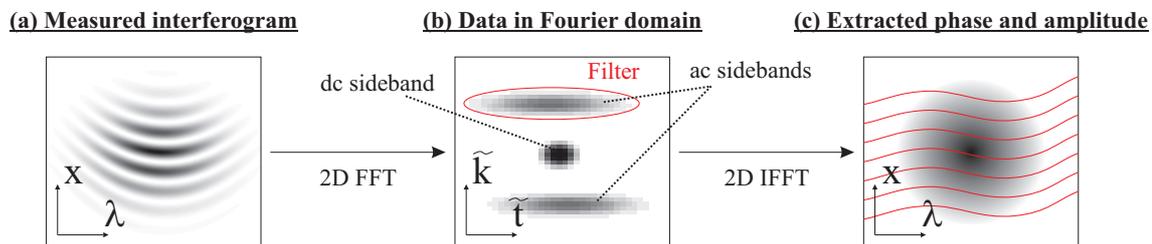


Figure 3.7: 2D FFT SEA-SPIDER inversion routine. (a) The two dimensional interferogram is converted into the pseudo wave vector/time domain (b) via a 2D fast Fourier transform (2D FFT); note that only the amplitude is plotted, but the complex value is actually retained in the reconstruction. One of the AC sidebands is extracted via a 2D filter and (c) the filtered data is inverse Fourier transformed (2D IFFT) back into the real space/frequency domain, where the phase is extracted and concatenated for each spatial position (red lines).

1D direct routine

The 2D reconstruction routine above can be implemented with a repetition rate of several Hertz on current computers. If one would like to have faster reconstruction rates, e.g. video rates (20 Hz and above), then this can be achieved using a 1D Fourier transform routine that does not require any filtering. Such a routine exists because of how the data is encoded in the fringes and is illustrated graphically in figure 3.8. In the absence of STC, the contours of the fringes directly map out the chirp of the pulse (i.e. the gradient of the spectral phase), as described above (section 3.2.2). It is possible to extract the position of the zeroth order fringe maximum as a function of frequency by performing a Fourier transform along the spatial co-ordinate for each frequency and extracting the phase along the ($\tilde{k} = K$) co-ordinate. However, it is not necessary to even perform the full Fourier transform (which is of order $N \log N$ operations for each frequency) as one knows *a priori* the value of K . Thus it is possible to calculate the phase difference directly from the formula:

$$\tilde{\theta}(\omega) = \arg \left[\sum_{n=0}^{N-1} \tilde{I}_{SEA}(\omega, x_n) e^{\frac{i2\pi nK}{N}} \right], \quad (3.35)$$

where x_n is the n^{th} spatial pixel ($n = 0 \dots (N - 1)$) and K is the tilt between the two beams in units of pixels. Removal of a calibration phase is possible by multiplication by the complex conjugate of the processed calibration data before taking the argument. This formula is of order N for each frequency, hence using this formula it is possible to reconstruct the pulse field at video rates. This can be useful for real-time adjustment of a pulse shaping device, such as a compressor for example. However, more accurate information can be obtained with the full 2D reconstruction, including information on STC.

If the test pulse exhibits STC, then the fringes will no longer have a constant periodicity in space for all frequencies. If the phase has a chirp that is linearly dependent on position such that $\tilde{\phi}(\omega, x) = \alpha x f(\omega)$, then the fringe periodicity becomes spectrally dependent, i.e. $K \rightarrow K(\omega)$. If the chirp is nonlinearly dependent on spatial position, then the fringes

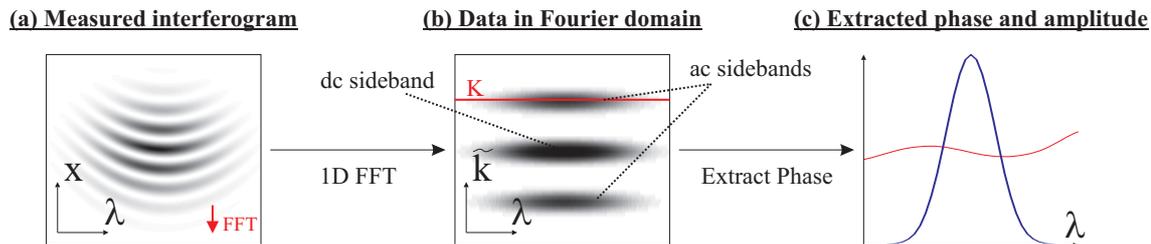


Figure 3.8: 1D FFT SEA-SPIDER inversion routine. (a) The two dimensional interferogram is converted into the pseudo wave vector domain (b) via a 1D fast Fourier transform (1D FFT) along the spatial dimension for each frequency. (c) The phase is extracted along the co-ordinate $\tilde{k} = K$ and then concatenated to obtain the average spectral phase.

will change periodicity along both co-ordinates. Such phases can be retrieved with no problems using the 2D FFT routine, provided the sidebands remain separated in the Fourier domain. In the case of the 1D FFT routine, in the presence of STC, then the phase that is extracted will equal the average of the spectral phases for all positions (see section 3.1.1). However, it may be necessary to extract the phase along a frequency dependent tilt value, $K(\omega)$.

3.2.4 Experimental setup for few-cycle pulses

SEA-SPIDER was first demonstrated by Kosik *et al.* for many cycle pulses, (100 fs at 800 nm).^[157] They also demonstrated that it is possible to measure STC with such a device by measuring the pulses after passing through a misaligned grating stretcher, where there was some spatial chirp. I have made some important modifications to their device that need to be made in order to measure few-cycle pulses.^[19] The experimental setup used for measuring few-cycle pulses is show in figure 3.9.

The ancillary chirped pulses (CP1 & CP2) are generated from the test pulse by taking the transmitted beam from a beamsplitter (BS1), stretching and then replicating it in an interferometer. Due to the large bandwidth, a thick piece of dispersive glass can be used to perform the stretching; in this case it was 10 cm of SF10 giving a group delay dispersion of approximately 20,000 fs². As with conventional SPIDER, these ancillary pulses could be derived from an external source, which can be much more intense for improved sensitivity,

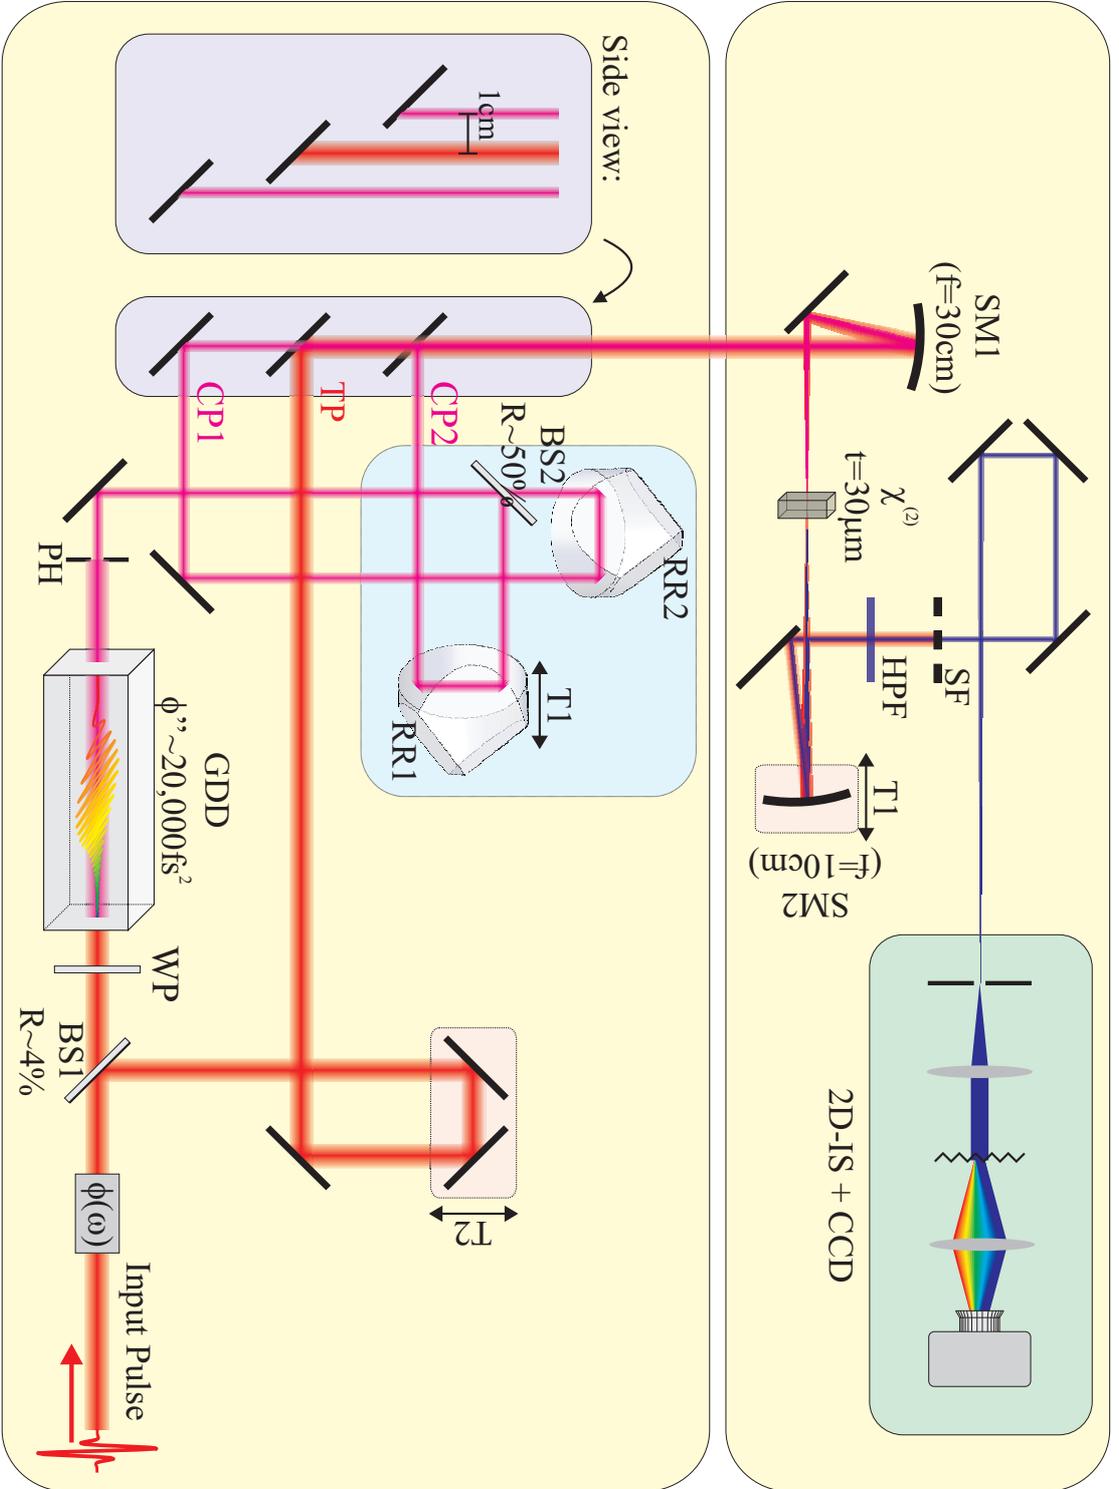


Figure 3.9: SEA-SPIDER setup for few-cycle pulses. $\phi(\omega)$: optional additional known phase; BS1 & BS2: beamsplitter 1 & 2, R = reflection coefficient; WP: half waveplate; GDD: group velocity dispersion (10 cm of SF10 giving $\phi'' = 20,000 \text{ fs}^2$); PH: pinhole; CP1 & CP2: chirped pulse 1 & 2; TP: test pulse; T1, T2 & T3: translation stages 1, 2 & 3; RR1 & RR2: retro-reflectors 1 & 2 (note that these do not need to be broadband, i.e. solid RRs are sufficient); SM1 & SM2: spherical mirror 1 & 2; $\chi^{(2)}$: nonlinear crystal cut for type II sum frequency generation, t = thickness; SF: spatial filter; HPF: high pass (frequency) filter; 2D-IS + CCD: 2D imaging spectrometer and charge coupled device. Shaded blue region = interferometer for generating two chirped arms. Upper shaded region = same as SEA-SPIDER concept in figure 3.5.

or at a different wavelength to allow downconversion (DC-)^[127] or H-SPIDER^[110]. By varying the delay between the two chirped pulses (using T1) in the interferometer, it is possible to change the shear independently to the tilt on the beams. Compare this to the conventional SPIDER, where the shear and delay are coupled. In SEA-SPIDER, it is possible to vary the shear to the most appropriate value for the pulses being measured, enabling more complex pulses to be characterised. The chirped pulses are spatially filtered using the pinhole (PH) to ensure that they contain no STC that can be transferred to the signal pulses during the SFG process, thus ensuring any STC is directly attributed to the TP.

The test pulse (TP) is taken as the reflection from the ultrabroadband beamsplitter (BS1). It is normally best to use as little power as possible in the TP because of its short duration. Therefore the reflection from a glass wedge is sufficient. Note that I used a beamsplitter rather than the reflection from the glass block used to stretch the ancillary pulses to help aid the alignment. Only reflective optics optimised for few-cycle pulses are used in the TP arm, ensuring that the TP does not acquire any additional dispersion other than air. As a result of not requiring a replica of the test pulse, SEA-SPIDER performs a zero additional phase (ZAP) measurement that does not require extra calibration of the dispersion of the apparatus, similar to ZAP-SPIDER.^[159] However, in ZAP-SPIDER, the two signal pulses are recombined with a delay between them so that the interferogram has fringes in the spectral dimension, allowing one to use a 1D spectrometer in akin to the conventional SPIDER. Obviously this has the advantage over the conventional device for short pulses that no additional phase is introduced to the test pulse and it decouples the shear from the delay, but it still requires a high resolution spectrometer and can have difficulties in measuring pulses with a complicated spectrum. Such a device cannot measure STC since it acquires a 1D data set.

The three pulses (the test and two chirped pulses), are aligned parallel but vertically displaced such that the test pulse (TP) is in the middle of the two chirped pulses and

a separation of 1 cm between each adjacent beam. This is sufficient to allow spatial separation of beams up to 0.5 cm in diameter, and to allow separation of the AC sidebands in the Fourier processing, but not too large such that the fringes cannot be resolved in the spectrometer. Using a 30 cm focal length spherical mirror (SM1), the three pulses are focused into a thin $\chi^{(2)}$ nonlinear crystal (30 μm of BBO cut for type II SHG at 750 nm). The test pulse is orientated along the o-axis of the crystal and the chirped pulses orientated along the e-axis. The thin crystal is sufficient for pulses with over one octave bandwidth. The test pulse can be delayed with respect to the two chirped pulses (adjusting T2 in figure 3.9) to change the mean wavelength of the upconverted pulses. This allows one to select the optimal spectral region for upconversion. Such an ability is desirable for two reasons: (1) it enables selection of two quasi-monochromatic frequencies from the chirped pulse that have relatively high intensity, which can be difficult for highly modulated spectra, and (2) when measuring over octave bandwidth pulses, it is desirable to ensure there is no spectral overlap between the fundamental and the upconverted spectra. Although the two signal beams are spatially separated and thus there should be no overlap between the fundamental and upconverted beams, in practise this is hard to achieve because of the relative intensities and sensitivity of the CCD for the different wavelengths. A spectral filter (HPF) is also used to reduce the fundamental background. Depending on the beam size being measured, it may be possible, and more desirable, to utilise spatial filtering (SF) only to remove the fundamental background.

After spatial and spectral filtering, the two signal pulses (SP1 & SP2) are imaged onto the entrance slit of a 2D imaging spectrometer (2D-IS). The imaging can be changed using a spherical mirror (SM2) mounted on a translation stage (T3). The adjustment of the imaging is to ensure perfect spatial overlap of the pulses for SEA-SPIDER. As the two pulses travel separate paths, it is necessary to use focusing mirrors instead of lenses to do the imaging before and inside the spectrometer. This is because the lens can introduce different dispersion or group delay (2.3), but the mirrors will not. Focusing mirrors also

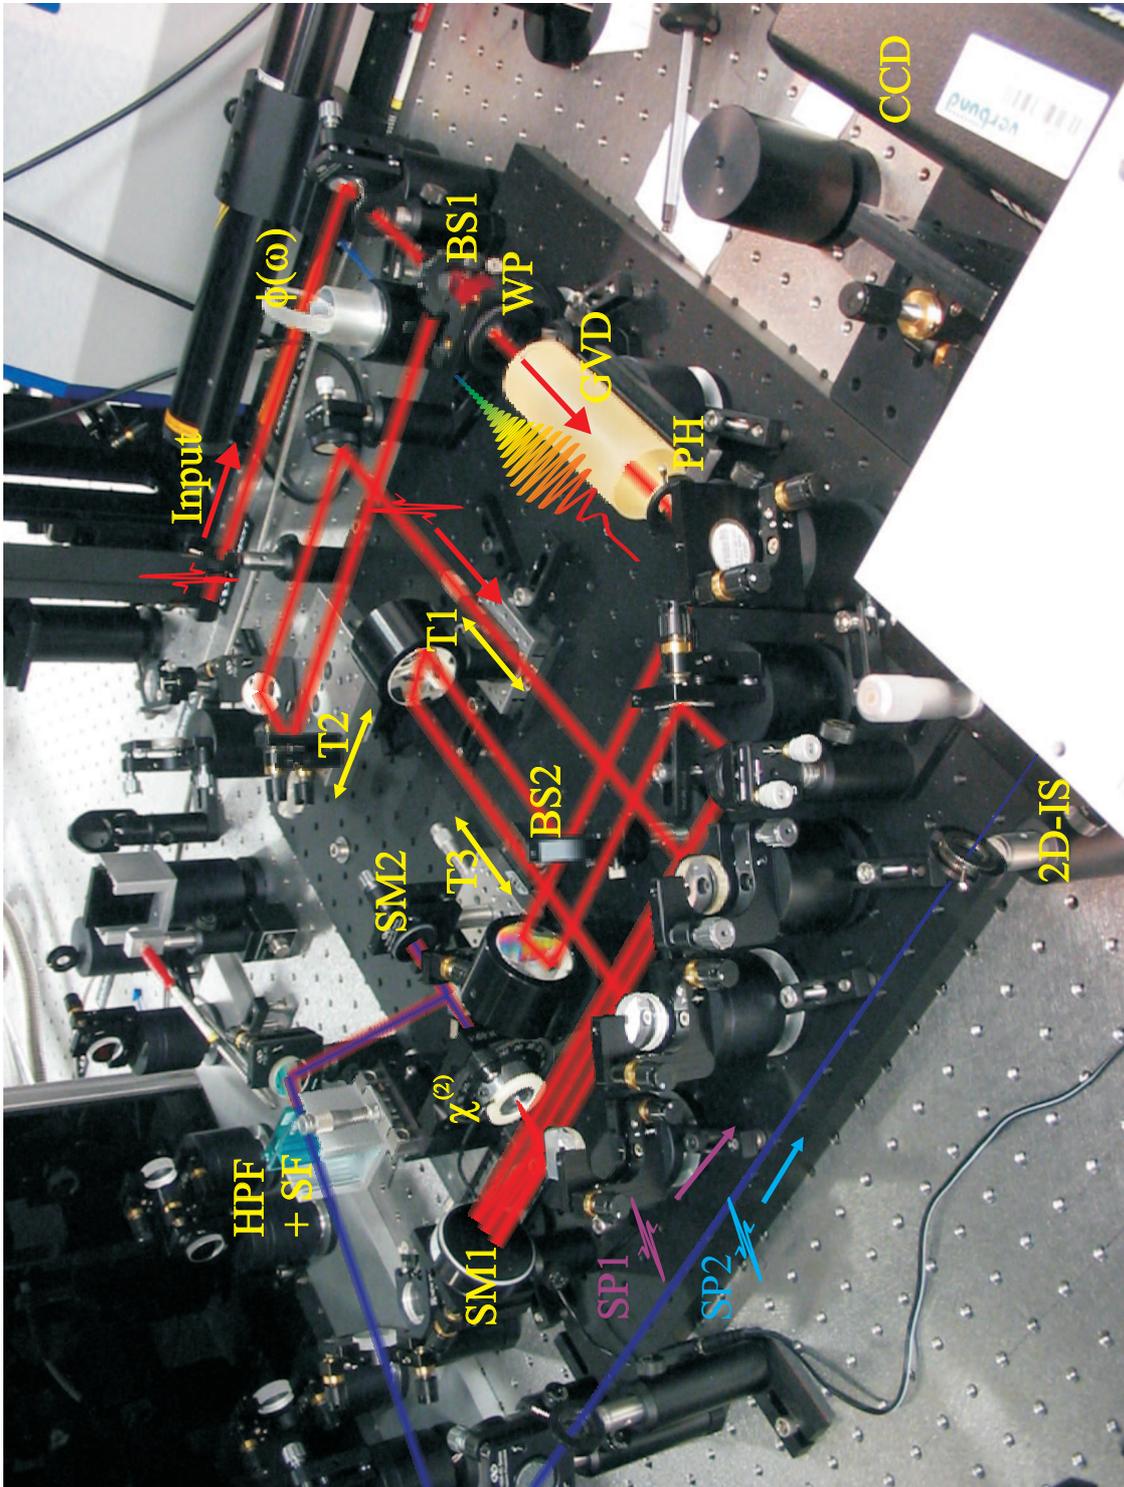


Figure 3.10: Experimental setup for SEA-SPIDER. The labels are the same for figure 3.9.

ensure that the two pulses will arrive with zero delay between them, and thus prevent any errors resulting from a timing jitter between the two signal pulses (section 3.1.1). The image is captured using a standard CCD or CMOS sensor. However, using an intensified CCD can allow single-shot operation, and thus possibly kHz acquisition. The experimental device is shown in figure 3.10.

3.2.5 SEA-SPIDER alignment

As with any device, it is necessary to ensure correct alignment in order to accurately measure the pulse. Due to acquiring a two dimensional trace and because each individual pulse (i.e. TP, CP1, CP2, SP1 & SP2) is spatially separated, there are certain procedures and alignment cross-checks that can be performed to ensure that the device is set up correctly.

Spatial overlap in the crystal

As the crystal is extremely thin (10–50 μm depending on the bandwidth being measured), it is necessary to ensure that the two CPs and the TP come to a focus at the exact same spatial position. However, spherical aberration from the focusing mirror causes the two CPs to focus at a slightly different position along the propagation axis compared to the TP. To set the alignment, I have found that it is best to scan a camera through the focus and ensure that the two CPs converge towards the TP equally as all three beams come to a focus. At the focus of the TP, it is necessary to ensure that there is perfect and equal spatial overlap of the two CPs with the TP. This should be the case if all three beams are parallel before the spherical mirror. A pick-off mirror after SM1 is used to minimise astigmatism. Spatial filtering of the CPs using the PH helps to improve the spatial overlap. This is because the CPs will have a smaller spot size on SM1 and therefore a larger focus. This will also minimise any STC being transferred from the CPs to the signal pulses.

Spatial overlap at the spectrometer

It is necessary to ensure that the two signal pulses perfectly spatially overlap at the entrance and exit of the spectrometer. Assuming the spectrometer images correctly, it should only be necessary to check the image plane of the spectrometer. The spatial overlap of the two signal pulses is optimised by adjusting the imaging of the crystal via translation of T3 in figure 3.9. It is critical to have perfect overlap to be able to measure a pulse with STC correctly.

Identical signal spectra

If the signal pulses correspond to spectrally shifted replicas of the TP, then the spatio-spectral distributions of the two individual signal pulses should be identical, and match that of the fundamental spatio-spectral distribution. This is easily verified by alternately blocking each chirped ancillary pulse and measuring the upconverted spatially resolved spectrum. The distributions should also be identical to that of the fundamental, other than a small change in the gross spectral shape due to the phase matching function of the crystal and the non-uniform response of the detector. As the delay between the CPs and TP is varied, the upconverted signal arms should shift in frequency without any change in the spectral shape, other than a change due to the phase matching and detection response functions. If a significant change in the spectral shape occurs, then this is most likely due to insufficient stretching of the ancillary arms. As a result, the TP will not interact with two quasi-monochromatic frequencies during the SFG process. If the spectrum is different to the fundamental, then this is likely to be due to self-phase modulation and other unwanted nonlinear effects in the crystal (or even the air) due to too much power whilst focusing. Using a focusing optic with a longer focal length, or reducing the input power should prevent this. As a guide, I found that it is sufficient if the signal pulses and the upconverted TP and CPs are just visible to the naked eye. This corresponds to approximately 200–300 mW input power from an oscillator at 80 MHz.

Calibration

In general, it is advisory to record a calibration interferogram. This not only helps with the alignment of the device, but can be used during the reconstruction. The calibration is acquired by setting the delay between the two CPs to zero such that there is zero spectral shear between the two signal pulses. Zero shear can be found by observing the individual spectra from the two signal arms and checking that they spectrally overlap completely. The interference pattern should consist of perfectly straight fringes (the extracted phase difference should be perfectly flat). It may be possible that, due to the frequency dependence of the tilt, the fringes converge slightly towards higher frequencies. This may result in an extracted phase that varies linearly with frequency. In this case, it is necessary to remove the calibration trace during the reconstruction.

If the fringes do not appear flat, then it is necessary to check that the two signal pulses are at the exact same wavelength, and ensure that the above procedures have been performed (the two signal arms should have identical spatio-spectral distributions that match that of the fundamental). If all these checks have been performed and the beams overlap in the crystal, then it is possible that some space-time coupling is being transferred from the chirped pulses into the signal pulses. This can be removed by spatially filtering the chirped pulse using the pinhole (PH) in figure 3.9.

3.2.6 SEA-SPIDER results - Few-cycle Oscillator

In order to test the ability of SEA-SPIDER to reconstruct few-cycle NIR (i.e. sub-10 fs) pulses, I measured the spectrum and phase of a Femtolaser Rainbow oscillator. This particular laser outputs 6 fs, 3 nJ pulses at about 80 MHz. The measurements were performed at the Max Born Institute (MBI), Berlin with G. Steinmeyer and G. Stibenz. This oscillator exhibits the three features associated with few-cycle pulses that were discussed at the start of the chapter — it exhibits an ultrabroadband, modulated spectrum with regions of low spectral intensity and also exhibits space-time coupling. In addition, the

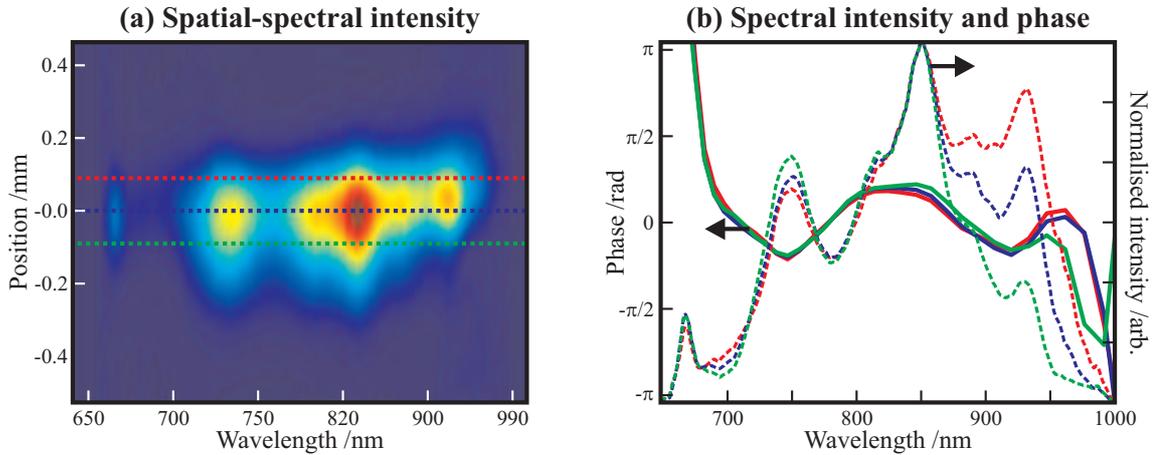


Figure 3.11: Spatially dependent spectral intensity and phase of a few-cycle oscillator, measured via SEA-SPIDER. (a) The spatially dependent spectral intensity, showing spatial chirp. The image is plotted linearly in frequency, but the scale has been converted to wavelength. (b) Spectral intensity (dashed lines) and spectral phase (solid lines) corresponding to the spatial positions indicated by the dotted lines of the same colour in (a). The intensity of each slice has its peak value normalised to one for comparison.

excellent stability of the laser enabled measurements of the pulse using several different devices, allowing a comparison between the different techniques and thus gauge their accuracy.^[19,20] The devices that were used in addition to the SEA-SPIDER were a conventional SPIDER^[109] and an interferometric (I-) FROG^[85] that had been optimised for measuring few-cycle pulses.

The calibration and SEA-SPIDER interferograms are shown in figure 3.6 (a) and (b) respectively. The spatially resolved intensity calibrated spectrum measured using the SEA-SPIDER is shown in figure 3.11. It is clear that the spectrum exhibits some spatial chirp, highlighted by the lineouts (with a peak intensity normalised to one) shown in (b).

The spectrum shown was extracted from the DC sideband of the calibration interferogram, thus removing any noise. The spectral phase was calculated via the 2D Fourier filtering routine followed by concatenation from a SEA-SPIDER interferogram with a shear of 30 mrad/fs. This shear is enough to support a 210 fs time window, enabling the low intensity wings of the pulse in the time domain to be measured. The phase functions for different spatial positions show remarkable similarity, although there is a slight increase in the group delay with position, resulting in pulse front tilt. The spatial chirp and PFT (or

equivalently angular dispersion) probably arises from transmission through wedged optics and at non-normal incidence (e.g. output coupler and Brewster-cut Ti:Sapph crystal).

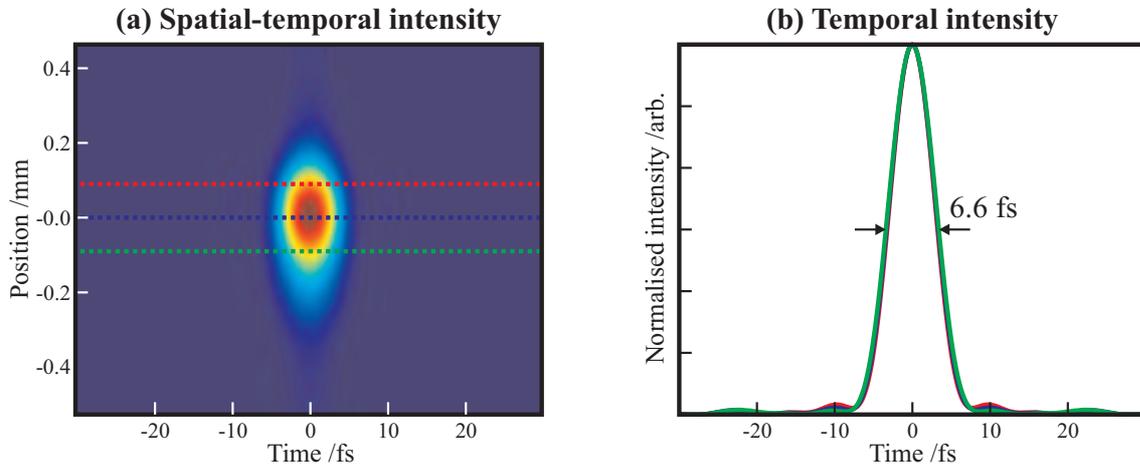


Figure 3.12: Spatially dependent FTL temporal intensity corresponding to a few-cycle oscillator. (a) The spatially dependent temporal intensity. Positive times correspond to the leading edge (i.e. pulse is travelling to the right). (b) Temporal intensity corresponding to the spatial positions indicated by the dotted lines of the same colour in (a). The FTL pulse duration is 6.6 fs FWHM intensity (2.4 optical cycles). The intensity of each slice has its peak value normalised to one for comparison. There is no PFT because the spectral phase was set to zero in calculating the temporal intensity.

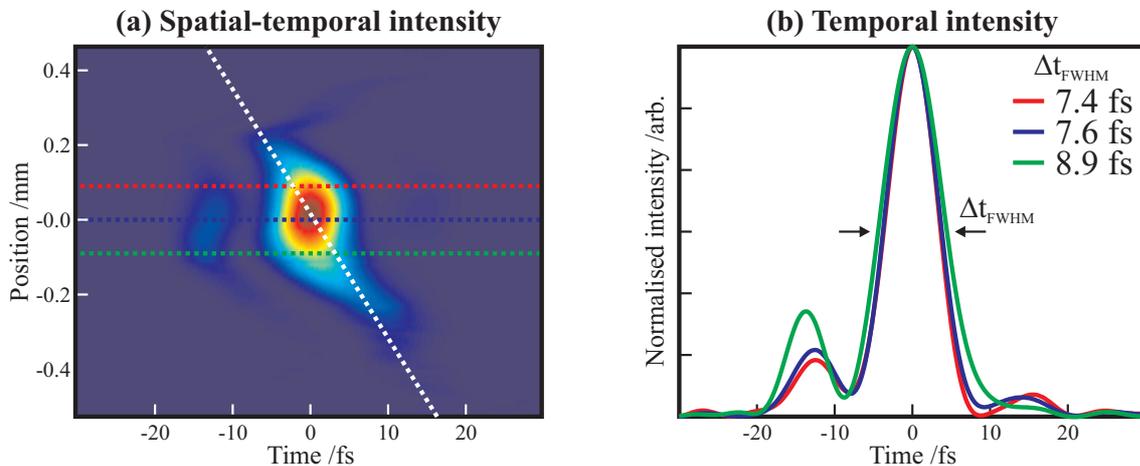


Figure 3.13: Spatially dependent temporal intensity of a few-cycle oscillator. (a) The spatially dependent temporal intensity, clearly showing PFT (indicated by the dotted white line). Positive times correspond to the leading edge (i.e. pulse is travelling to the right). (b) Temporal intensity corresponding to the spatial positions indicated by the dotted lines of the same colour in (a) and the corresponding FWHM intensity durations. The intensity of each slice has its peak value normalised to one and the group delay removed for comparison.

The Fourier transform limited (FTL) spatio-temporal pulse intensity corresponding to the spectrum shown in figure 3.11 is shown in figure 3.12. It is clear that the temporal distribution does not vary as a function of position due to the spatially dependent spectrum alone. However, the carrier frequency at each position does vary due to the spatial chirp. The FWHM intensity duration is 6.6 fs, corresponding to 2.4 cycles at this wavelength (816 nm).

The reconstructed spatially dependent temporal intensity corresponding to this pulse is shown in figure 3.13. It is clear that this pulse exhibits PFT, highlighted by the dotted white line, especially in the spatial wings of the pulse. The PFT is also observed by the spatially dependent group delay (i.e. gradient of the spectral phase as a function of position). In addition, the chirp of the pulse is spatially dependent. This is more evident in the normalised intensity slices shown in (b). Even though this particular pulse does not appear to exhibit significant nonlinear spectral phase (i.e. dispersion), it is sufficient to chirp the pulse by more than 10% of its minimum duration, due to the large bandwidth of the pulse. However, this large bandwidth also results in significant temporal broadening due to dispersion in air. This needs to be compensated for if the shortest pulse at the experiment is required. In addition, it may be necessary to compensate for the spatial chirp and PFT.

Accuracy of the reconstruction

The accuracy of the technique is hard to quantify as one cannot directly compare it to a known reference. However, there are a couple of analyses that highlight any possible inaccuracies of the data and/or reconstruction. Firstly, it is possible to gauge the accuracy and identify any systematic errors by performing a reference-based measurement. This involves reconstructing the phase of the laser with and without a piece of glass of known dispersion that is placed into the beam immediately before the device in a way that ensures no change in the beam pointing. The difference in the reconstructed spectral phases of these two measurements is then compared to the known dispersion of the glass,

calculated using the Sellmeier equations. Secondly, it is possible to change the value of the shear, to generate independent reconstructions of the pulse. Provided each value of the shear satisfies the WS sampling limit, each of these reconstructions should be identical. However, as the shear becomes large, such that the temporal intensity of the pulse extends outside the temporal window supported by the shear, resulting in aliasing, the reconstructed pulse should start to deviate from that calculated using a smaller shear. However, the phase difference between two frequencies separated by a large shear should match the phase difference of the same two frequencies resulting from a reconstruction from an interferogram generated with a small shear.

The dispersion of 1 cm of fused silica and the measured dispersion for two different values of shear is shown in figure 3.14. The values of the shears were 30 mrad/fs (i.e. is taken from the data above) and 72.5 mrad/fs corresponding to temporal windows of approximately 210 fs and 90 fs respectively. The chirp from the fused silica is enough to stretch the pulse to approximately 180 fs (FWHM). The measured phase shows very little deviation from the known value of dispersion, indicating the high accuracy of the method. At wavelengths below 725 nm and above 950 nm, the spectral intensity starts to become quite small, especially considering the detector response, often resulting in errors in conventional techniques. However, the measured phase of the glass matches that of the calculated phase in these regions. Also, deviations in the spectral phase only occur when the spectral intensity reaches the noise floor (i.e. $\lambda \leq 655$ nm and $\lambda \geq 990$ nm).

The reconstructed temporal intensity of the oscillator pulse, scaled and centred at $t = 0$ for $x = 0$ for the same two shears is shown in figure 3.15. The agreement between the two data sets is very good, showing very similar features at the same times. However, there are differences between the two data sets. Firstly, the pulse extends over about 100 fs in total, thus slightly outside the time window allowed by the larger shear. Secondly, these data correspond to a time early in the development of the SEA-SPIDER device. Thus there are certain alignment procedures that were not taken into account at that time. Thirdly, the

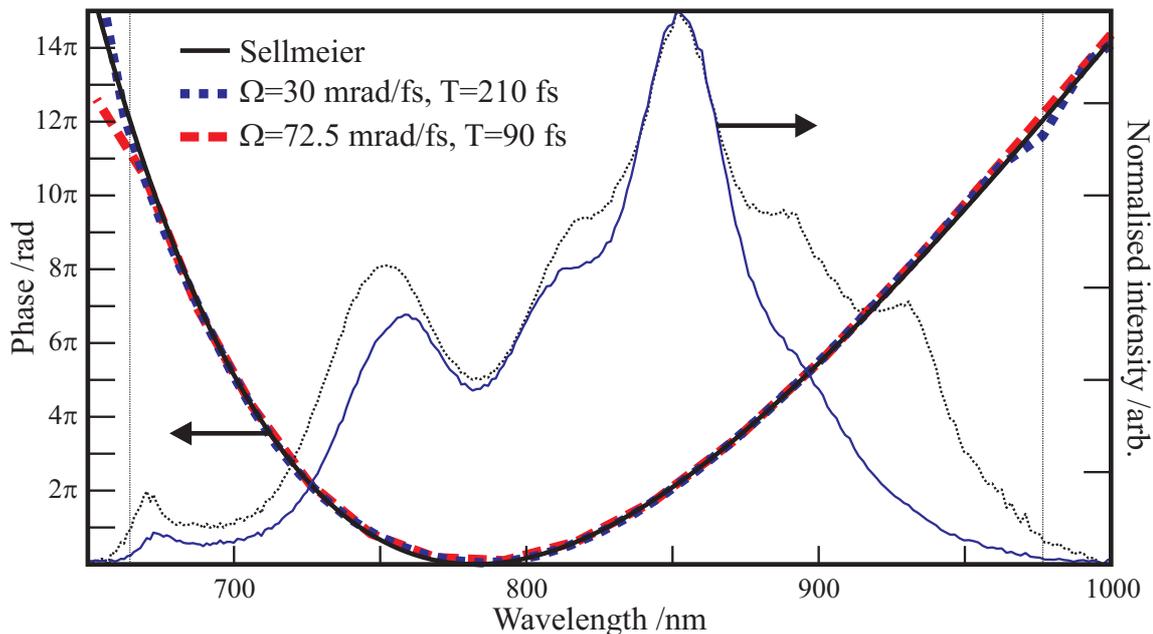


Figure 3.14: Phase of 1 cm of fused silica measured using SEA-SPIDER and calculated via the Sellmeier equation. The phase was measured with two different shears of 30 mrad/fs (short dashed blue) and 72.5 mrad/fs (long dashed red). The phase calculated from the Sellmeier equation is shown as solid black. The normalised measured upconverted spectrum (solid blue) and intensity calibrated spectrum (dotted black) are shown for reference. The limits of good agreement are shown by the dotted vertical lines.

calibration of the spectral wavelength and intensity was not as accurate as could be, due to time constraints. Thus improved alignment and calibration, coupled with a different set of shears should result in improved agreement between the different data sets. However, the FWHM pulse durations (7.6 fs and 7.8 fs) are in excellent agreement with a deviation (3%) that is less than expected due to experimental and pulse-pulse fluctuations. As zero delay between the two pulses is guaranteed due to generating from the same pulse and imaging via focusing mirrors, fluctuations in successive reconstructions is likely to be due to pulse-pulse fluctuations in the spectral phase rather than from experimental errors.

The Wigner distribution for the central position of the pulse is shown in figure 3.16, showing the small chirp across the central spectral region. It is clear that the post-pulse at $t \simeq 12$ fs arises from the peaks in the spectrum. The wings of the spectrum appear at far later times. This is evident from the spectral phase, where the GD is much more than the central region. The most probable cause for this is due to the limited bandwidth of

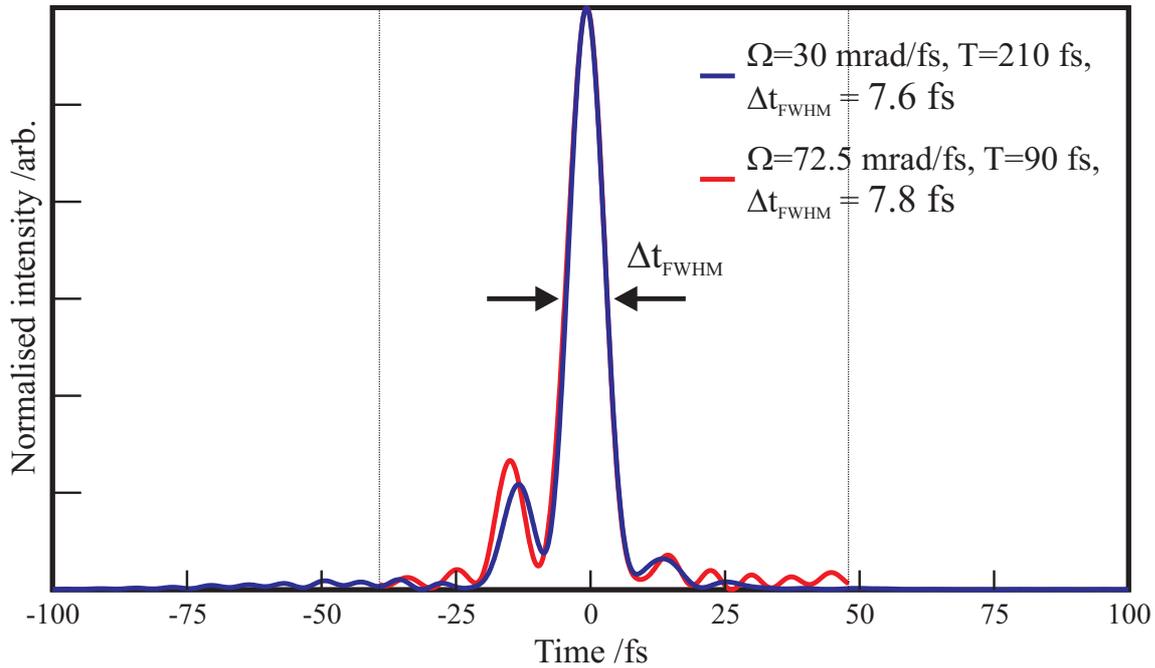


Figure 3.15: Reconstruction of a few-cycle oscillator at the centre of the pulse from two different shears. The shears were 30 mrad/fs (blue) and 72.5 mrad/fs (red), which can support a time window of 210 fs and 90 fs each. The time window for the larger shear is contained within the dotted vertical lines.

the intracavity chirped mirrors used for dispersion compensation. As a result, the wings of the spectrum will have a significantly different penetration depth compared to the bulk of the pulse.

Comparison of SEA-SPIDER with alternative techniques

Probably the best method of testing a device's ability to accurately reconstruct a pulse is to compare it to other independent complete characterisation techniques. Due to the stability of this oscillator, it acts as an extremely reliable source to perform such a test. Certain features of the pulse, such as large bandwidth, modulated spectrum, regions of low spectral intensity, sub-pulses and space-time coupling, can also be used to identify where the different techniques have problems in reconstructing few-cycle pulses and highlights the robustness of the SEA-SPIDER device. The data presented below was taken by G. Stibenz at MBI, Berlin using their own developed SPIDER^[109] and IFROG^[85] opti-

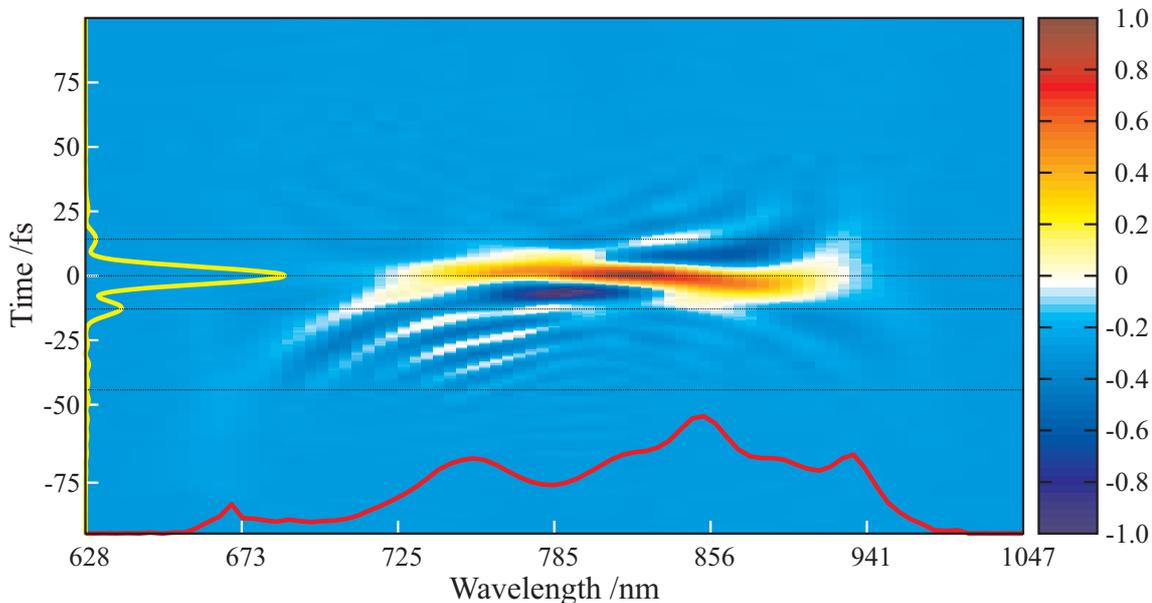


Figure 3.16: Wigner distribution of a few-cycle oscillator showing complicated structure in the wings. The temporal (yellow) and spectral (red) marginals are shown for reference.

mised for few-cycle pulses, and the SEA-SPIDER that I developed. The autocorrelation data was extracted from the I-FROG data.

The spectrum and reconstructed phase for the different techniques are shown in figure 3.17. For SEA-SPIDER, I have taken the spectral phase and intensity from the centre of the pulse. Due to the large bandwidth of these pulses, it is necessary to correct the reconstructed phase for the different path lengths inside each device; I have already included these corrections in the figure. Further details of these corrections and the comparison of the devices can be found in Stibenz^[20]. It is clear that the different techniques reconstruct the pulse with remarkable similarity, especially when one considers the fragility of these pulses with respect to space-time coupling and significant dispersion, in addition to the very different nature of the methods.

The reconstructed temporal envelope from the different techniques is shown in figure 3.18. The pulse intensity error between the different methods is shown in table 3.1 and the pulse durations in table 3.2. The intensity error for normalised fields (i.e. $\int_{-\infty}^{\infty} dt I_i(t) =$

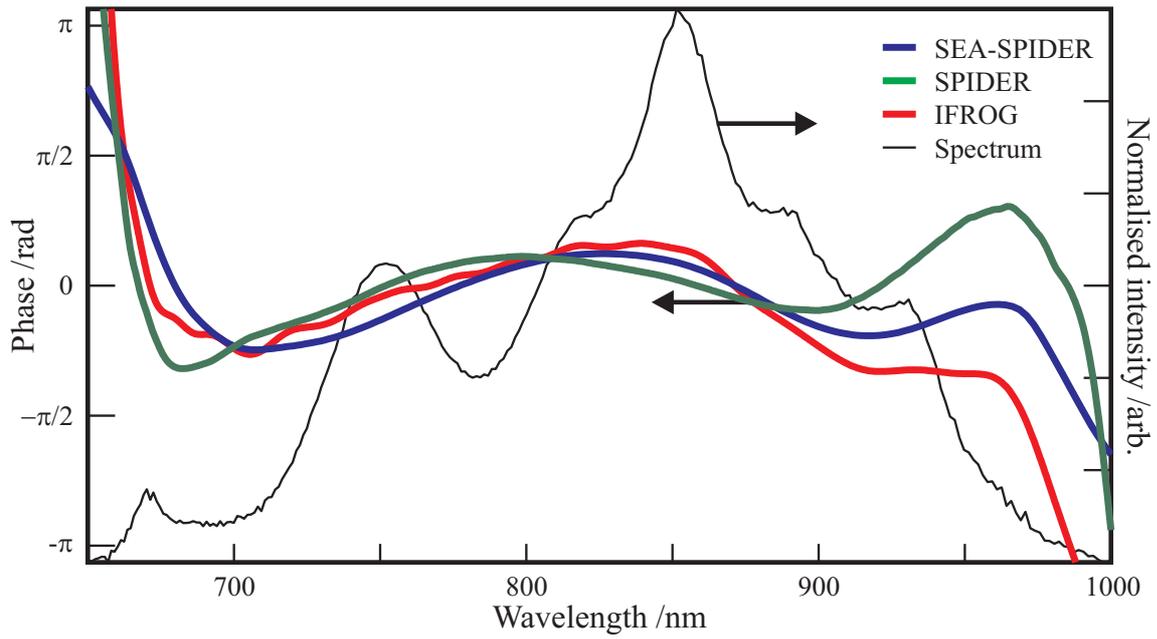


Figure 3.17: Reconstructed spectral phase of a few-cycle oscillator using different techniques: SEA-SPIDER (blue), SPIDER (green) and IFROG (red), and the spectral intensity (black).

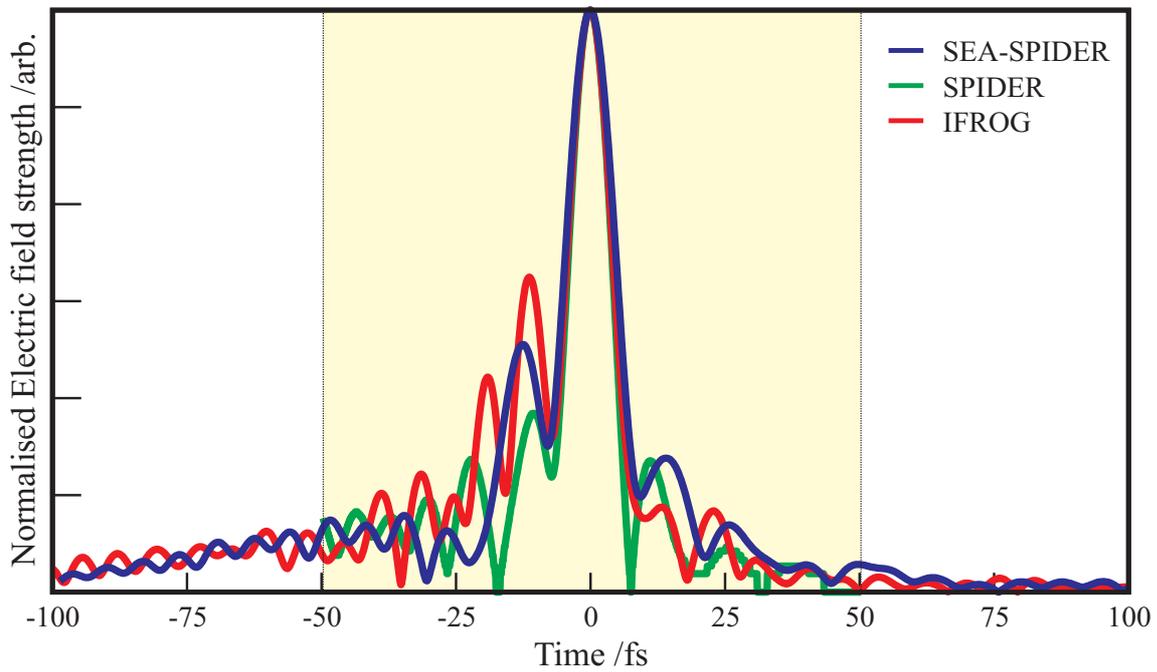


Figure 3.18: Temporal electric field envelope of a few-cycle oscillator reconstructed from different characterisation techniques: SEA-SPIDER (blue), SPIDER (green) and IFROG (red).

	SPIDER	IFROG
SEA-SPIDER	0.05	0.04
SPIDER		0.05

Table 3.1: Pulse intensity error between different reconstruction methods.

<i>Method</i>	<i>FWHM duration/fs</i>
SEA-SPIDER	7.6
SPIDER	6.6
IFROG	6.9

Table 3.2: FWHM pulse durations from different methods.

1) is defined as

$$\epsilon_I = \sqrt{\int_{-\infty}^{\infty} dt |I_1(t) - I_2(t)|^2}. \quad (3.36)$$

The electric field envelope has been shown, rather than the intensity, to highlight the regions of low temporal intensity. It is clear that the conventional SPIDER method has undersampled the pulse because the intensity does not fall to zero within the time window allowed by the shear. For that reason, it is likely that the post-pulse and main pulse is not accurate. The IFROG and SEA-SPIDER reconstructions show excellent agreement. The extremely low intensity post-pulse at $t \sim -50$ fs matches very well between these two advanced techniques. It is expected that FROG should recover satellite pulses because of the temporal information contained within the spectrogram. SEA-SPIDER manages to capture this information as well, even though it appears SPIDER did not. The reason for this may be because of the different group delay at the wings of the spectrum arising from the chirped mirrors. Due to the low spectral intensity within this region, SPIDER may struggle to recover the phase. As the 2D SEA-SPIDER algorithm results in high noise discrimination due to the two-dimensional filtering, it is able to recover this phase. This is evident by the step in the fringe contours of the SEA-SPIDER interferogram at these wavelengths. This time-frequency picture is more easily seen in the spectrogram for this pulse (figure 3.19).

3.2.7 SEA-SPIDER results - CPA

The amount of dispersion that is introduced by the glass block used to stretch the ancillary pulses in the SEA-SPIDER device used in section 3.2.6 approximately equals $16,000\text{fs}^2$.

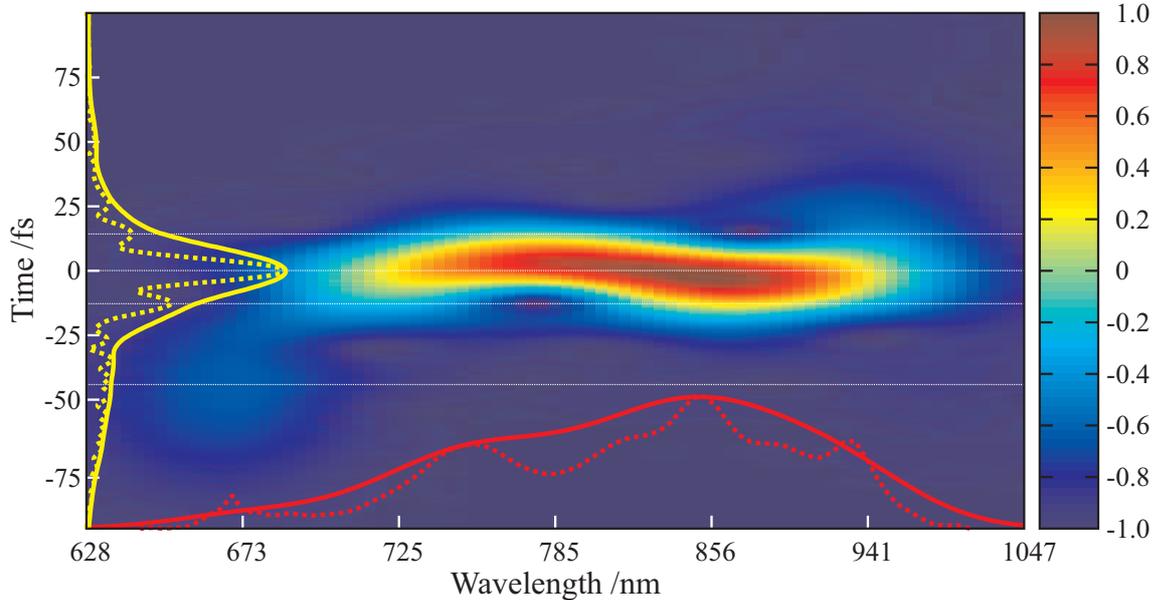


Figure 3.19: Spectrogram of a few-cycle oscillator showing a post-pulse in the wings of the spectrum. The marginals, corresponding the smoothed temporal (solid yellow) and spectral (solid red) envelopes is also shown, along with the actual temporal (dotted yellow) and spectral (dotted red) envelopes. The low intensity pre-pulse is clearly visible at $t \sim -50$ fs and $\lambda \sim 670$ nm.

As SPIDER requires the test pulse to upconvert with a quasi-monochromatic frequency from the chirped pulse, this limits the maximum duration that the device can measure. This limitation can be calculated by stating that the change in the frequency during upconversion, $\delta\omega = T/\phi''$ arising from a change in the monochromatic frequency over the duration of the test pulse, must be smaller than the sampling frequency (i.e. the shear) given by $\Omega = 2\pi/T$, where T is the pulse duration. For this device, this sets a limit of several hundred femtoseconds. Thus this particular SEA-SPIDER is capable of measuring the output of a CPA system, with a typical pulse duration on the order of 30–100 fs FWHM. However, this is at the limit for accurate reconstruction and longer pulses will result in errors in the reconstruction. It is possible to measure longer pulses using a different means of stretching the pulse (e.g. a prism/grating stretcher/compressor) or a glass block with higher dispersion.

The versatility and intuitive nature of the SEA-SPIDER can be demonstrated by the interferograms shown in figure 3.20, which corresponds to different amounts of dispersion.

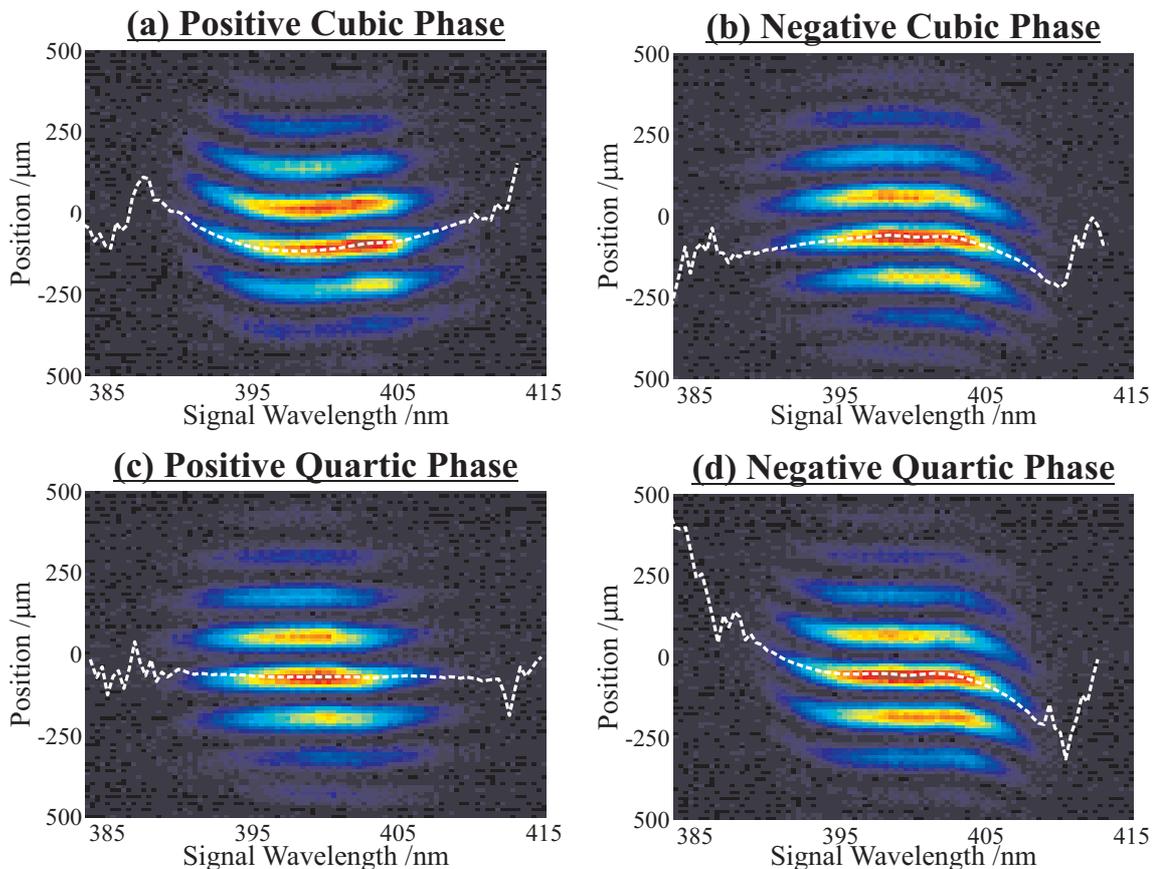


Figure 3.20: SEA-SPIDER interferograms of a CPA pulse with different amounts of chirp. The shear used was equal to 25 mrad/fs (a time window of 250 fs). The dispersion of the pulse is given by (a) $\tilde{\phi}(\omega) = 120,000/6(\omega - \omega_0)^3$, (b) $\tilde{\phi}(\omega) = -180,000/6(\omega - \omega_0)^3$, (c) $\tilde{\phi}(\omega) = 100,000/24(\omega - \omega_0)^4$, (d) $\tilde{\phi}(\omega) = -500,000/24(\omega - \omega_0)^4$. The contours of the fringes is given by the gradient of the spectral phase. The positive quartic phase shown corresponded to the maximum amount that could be applied with the Dazzler for this particular laser. The dashed white lines show the fringe contours, calculated using the 1D Fourier extraction routine.

The pulses were derived from a Femtolasers Compact Pro CPA, which employs an AOPDF (Fastlite Dazzler) before a multipass amplifier stage followed by a prism compressor, delivering 1 mJ 30 fs pulses at 1 kHz. The Dazzler is primarily used to stretch the pulse before the amplifier stage and to pre-compensate the large amount of extra fourth and third order phase introduced by the prism compressor. However, the Dazzler can also be used to add a known amount of dispersion to the pulse as the phase of the output pulse depends linearly on the phase introduced before the amplifier stage.

The interferograms shown in figure 3.20 correspond to (a) positive and (b) negative third order spectral phase, resulting in a fringe contour that varies quadratically with frequency; and (c) positive and (d) negative fourth order spectral phase, resulting in a fringe contour that varies as the cube of the frequency. The dashed white lines show the fringe contours, and was produced by the 1D extraction method equation (3.35). Using an AOPDF to perform the pulse shaping has the advantage of not introducing any significant STC and does not suffer from pixelation compared to a 4-f shaper.

3.3 Summary

Methods of characterisation of ultrashort optical pulses that provide complete spectral information (i.e. phase and intensity) have become common laboratory tools for ultrafast laser groups. However, their accuracy can be reduced for certain classes of pulses, such as those that exhibit space-time coupling. For extremely short pulses, whose duration is on the order of the carrier frequency, these conventional methods may not provide accurate results. This is because such pulses tend to exhibit one or more of the following: (1) over an octave spanning spectrum. (2) highly modulated spectrum, (3) regions of low spectral or temporal intensity and (4) space-time coupling. The accuracy and reliability criteria for measuring these pulses may be too high for certain applications such that conventional FROG and SPIDER methods are insufficient. There exists a wide array of variations on these two techniques that can circumvent certain problems, but no one device has been shown to achieve this accuracy when the pulse simultaneously exhibits all of the features stated above. Therefore I have developed a new method, based on SPIDER that encodes the spectral phase as a modulation on the periodicity of spatial fringes introduced between the two signal beams, which has been called spatially encoded arrangement for SPIDER (i.e. SEA-SPIDER).

I have shown that the accuracy of SEA-SPIDER is extremely high over the whole spectral bandwidth of the pulse provided the detected spectrum is higher than the noise

background. Even when the SNR approaches unity, SEA-SPIDER performs better than conventional SPIDER due to a higher noise reduction from utilising the extra dimension. As a comparison, IFROG also performs slightly better than conventional FROG for the same reason — the fringes allow a background free measurement to be performed and thus increasing the dynamic range of the device. Using the SEA-SPIDER and analysing the reconstruction, along with the raw interferogram and performing additional measurements, it was possible to identify features of the pulse that were not possible to measure using conventional SPIDER or FROG. For example, low intensity sub-pulses, attributed to regions of low spectral intensity that have different group delays to the rest of the pulse were clearly observed. The complete spatio-temporal profile of the pulse is also measured using SEA-SPIDER. As a result, it is clear that SEA-SPIDER outperforms the conventional techniques for measuring few-cycle pulses and has features that allow it to identify properties of the pulse being measured that other techniques might miss. The main advantages of SEA-SPIDER are that it can spectrally sample the pulse at the sampling limit, provides complete spatial-spectral phase (in a single spatial dimension), can perform single shot over a large bandwidth, and provides intuitive interferograms that allow online optimisation and highlights properties of the pulse. These features should prove useful in experiments that rely on optimisation of few-cycle pulses, and where it is necessary to characterise the pulse extremely accurately, including STC, for use in simulations. HHG, for instance, is an example where this is likely to prove very useful.

Attosecond metrology and high harmonic generation

There are two fundamental problems in generating sub-femtosecond (i.e. attosecond) pulses: (1) for propagating pulses, the minimum duration must be at least one period of the carrier frequency, and (2) the bandwidth required to support a pulse is proportional to the inverse of the pulse duration. Current ultrashort pulse technology is based around titanium sapphire (Ti:Sapph) lasers and thus have a central wavelength in the near-infrared (NIR) spectral region, limiting the period and hence shortest pulse to several femtoseconds. The bandwidth of these pulses is also limited by the gain medium used, which for Ti:Sapph is several hundred nanometres. In order to have a 100 as pulse, the central wavelength must be less than 30 nm, with a bandwidth $\Delta\lambda \simeq \lambda_0$ at such a wavelength, thus attosecond pulses must be in the extreme ultraviolet (XUV) spectral region. It is currently not possible to generate such a short wavelength and large bandwidth out of a laser directly, so it is necessary to use an optically nonlinear process to generate the extra frequencies required to synthesise these short pulses. At present, there are only a handful of realisable approaches to generating attosecond pulses. The approaches can be split into two broad categories: (1) generation of a broad frequency comb or (2) generation of a broad frequency continuum.

The first method is somewhat analogous to the generation of ultrashort pulses through mode-locked (ML) lasers. In ML lasers, the laser cavity consists of a broad frequency comb, corresponding to the longitudinal modes of the laser (i.e. standing waves), spanning up to several hundred nanometres with a spacing equal to the inverse of the cavity round trip time. This frequency comb corresponds to a train of short pulses provided the individual ‘teeth’ are phase-locked. Extending the process further, it may be possible to generate a frequency comb larger than that supported by current lasing media (e.g. Ti:Sapph) via the use of sum- and difference-frequency mixing^[160], cascaded stimulated Raman scattering^[161–163] and high harmonic generation (HHG)^[164–167].

The problem with generating a frequency comb via these methods is that the output is a train of pulses separated by the inverse of the frequency separation between successive teeth of the comb, which is on the femtosecond timescale. Thus it is unlikely that it will be possible to extract a single pulse by temporal switching in a similar manner as femtosecond pulses — for example a Pockels cell is used to select individual pulses from a Ti:Sapph oscillator. Thus the generation of a single broad continuum is deemed more practical as it can generate isolated attosecond pulses, provided the spectral phase of the continuum is sufficiently small.

One method of generating a broad continuum may be through free electron lasers (FEL) at large-scale facilities.^[168–172] Such facilities provide a high-brightness tunable source and can reach extremely short wavelengths. However, due to their large size and vast expense, there will consequently be few such sources. Current femtosecond/x-ray FELs are still in the early stages of development and as of yet have not produced isolated attosecond pulses. Smaller-scale methods could be based on molecular phase modulation via coherent Raman scattering^[161–163] or HHG^[173–175]. The former technique uses molecular vibrational (or rotational) motion as an ultrafast phase modulator: an ultrashort optical pulse is modulated by the rapidly varying refractive index created by a coherently excited vibrational (rotational) wavepacket. When correctly timed relative to the

molecular motion, an ultrashort pulse of duration shorter than the vibrational (rotational) period can be spectrally broadened, shifted, or compressed. High harmonic generation is the result of the interaction of extremely short and intense laser fields with gaseous matter. Using a few-cycle pulse and the high nonlinearity of the process, it is possible to generate a broad continuum (after spectral filtering) and thus an isolated pulse. Coherent Raman scattering has the advantage of being highly efficient, but is also in the primary stages of research and has not produced an isolated sub-femtosecond pulse. High harmonic generation, on the other hand, is an extremely inefficient process, with an efficiency on the order of 10^{-5} at best. However, HHG has been shown to produce isolated pulses as short as 130 as^[176] and is expected to be able to reach a few tens of attoseconds^[177]. HHG is therefore the most mature technology for the production of attosecond pulses, and is the primary route used in many research laboratories worldwide, hence I will concentrate on this method only for the remainder of this thesis.

As attosecond pulses necessarily have to be in the VUV to XUV spectral region, there are limited optics and methods of detection. For example, there does not exist a nonlinear crystal that is both transparent in this spectral region and can support such a large bandwidth, thus it is not possible to utilise the existing ultrashort pulse measurement techniques used in the femtosecond domain as described in their conventional form in chapter 2. Current XUV pulse characterisation techniques are based on photoelectron spectrometry (PES): due to the high energy of the photons, it is possible to ionise atomic media. A measurement of the velocity distribution of these ionised electrons (using an electron spectrometer) gives the spectrum of the XUV pulse. However, a temporal measurement also requires the measurement of the spectral phase, which is not measured directly with a spectrometer. Methods that measure the spectral phase rely on either multiphoton ionisation (e.g. auto- and cross-correlation) or streaking of the XUV ionised electron velocity distribution with an intense femtosecond pulse. These techniques can be considered as general characterisation methods for attosecond pulses in the XUV spectral

region, as they measure the temporal characteristics at some point after the generation. Each method has its own advantages and disadvantages, as well as a limited validity range. I discuss some of the methods briefly in section 4.3.

Alternatively, the metrology can be performed during the high harmonic process itself. Such techniques involve modifying the generation process in some controlled manner and then measuring the spectrum as a function of the control parameter. Using a model for the high harmonic generation process, one can then infer the temporal characteristics of the XUV pulses. Such a technique can benefit from the simplicity of the experimental apparatus because it is only necessary to use an XUV spectrometer, which can also yield a much higher signal to noise ratio. It is therefore possible to obtain data in a much shorter time frame, thus reducing the stringent long-term laser stability required for PES techniques. However, such a method is clearly dependent on the physics of the generation and relies on a well established model of the interaction. When used in combination with PES techniques to test their validity, direct detection methods can improve attoscience as they may allow faster, more reliable acquisition and a wider parameter space of measurement on a day-to-day basis.

The latter chapters of this thesis deal with the metrology of attosecond pulses and attosecond pulse trains from high harmonic generation via the use of direct detection of the XUV photons. However, it is necessary to be able to understand the process in detail before it is possible to consider how to control it in such a way that it is possible to extract information on the temporal characteristics of the generated light. The following section summarises previously developed models used to describe the interaction of an intense low frequency (i.e. NIR) field with atomic media, leading to high harmonic generation. Such a treatment only deals with the interaction of one atom with the field, thus is dubbed the single-atom response, so the following section after will deal with how to model the complete macroscopic response (i.e. the contribution of all the atoms present in the generating medium). Finally I shall discuss current methods for attosecond metrology through photoelectron spectroscopy.

4.1 HHG: Single-atom response

4.1.1 Three step semiclassical model

The underlying physical process for HHG is fairly well explained through a three step semiclassical picture first proposed by Corkum^[178] and shown schematically in figure 4.1: (i) A bound electron in an atomic gas is ionised via either tunnel or over-the-barrier ionisation due to the suppression of the atomic potential by means of an intense optical pulse (with peak intensity of 10^{13} – 10^{15} Wcm⁻²) generating an electron wave packet in the continuum; (ii) the electron wave packet is considered as a classical free electron, which is driven by and picks up energy from the optical pulse electric field. The electron may traverse many atomic radii away from the atom as it oscillates in the field; and (iii) the electron may return and collide with the parent ion, emitting high-energy photons with a frequency proportional to the energy picked up when oscillating in the optical field, plus the ionisation potential of the gas. The details of these steps will be considered below to show how this process can be used to generate attosecond pulses.

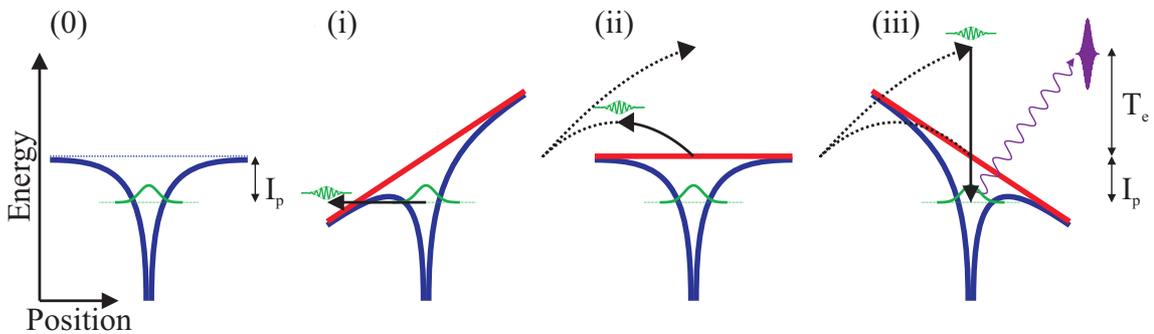


Figure 4.1: Three step model of high harmonic generation. (0) A bound electron (green) in a Coulomb force (blue). (i) The potential from an intense driving field (red) distorts the Coulomb potential to allow an electron wavepacket (green) to tunnel through the barrier (or even allow over-the-barrier ionisation if sufficiently intense). (ii) The free electron wavepacket is driven by and picks up energy from the laser field and oscillates in the field (trajectory shown in black). (iii) Approximately half a cycle later, the electron may cross the original atom and recollide, emitting high frequency radiation (purple) with a frequency equal to the electron kinetic energy plus the ionisation potential.

Noble gas	I_p (eV)	I_p (a.u.)
<i>Helium</i>	24.4	0.897
<i>Neon</i>	21.6	0.794
<i>Argon</i>	15.8	0.581
<i>Krypton</i>	13.99	0.514
<i>Xenon</i>	12.13	0.446

Table 4.1: Ionisation potential of some Noble gases used for HHG.

(i) Ionisation

The first process in the three step model is to calculate the ionisation probability $P(t) = W(t) \delta t$, during the time interval δt as a function of the electric field $\xi(t)$. This can be calculated using the quantitative quasistatic Ammosov, Devone, Krainov (ADK) model that gives the ionisation rate as:^[179]

$$W(t) = C_{n^*l^*}^2 G(l, m) I_p \sqrt{\frac{3}{\pi} \frac{\xi(t)}{\xi_p}} \left(2 \frac{\xi_p}{\xi(t)}\right)^{2n^* - |m| - 1} \exp\left(-\frac{2}{3} \frac{\xi_p}{\xi}\right) \quad (4.1)$$

$$n^* = Z (2I_p)^{-1/2}$$

$$\xi_p = (2I_p)^{3/2}$$

$$C_{n^*l^*}^2 = \frac{2^{2n^*}}{n^* \Gamma(n^* + l^* + 1) \Gamma(n^* - l^*)}$$

$$G(l, m) = \frac{(2l + 1) (l + |m|)!}{2^{|m|} |m|! (l - |m|)!}$$

where all values are in atomic units, I_p is the ionisation potential of the atom, n^* and l^* are the effective principle and orbital quantum number for the electron before ionisation, l the orbital quantum number and m its projection, and Z is the charge of the atomic residue. For $l \ll n$, then $l^* = 0$, otherwise $l^* = n^* - 1$. It can be shown that for $l = 1$ that the dominant contribution to ionisation is from electrons in the $m = 0$ state, thus only this value needs to be considered. The ionisation potential of different atoms are given in table 4.1.

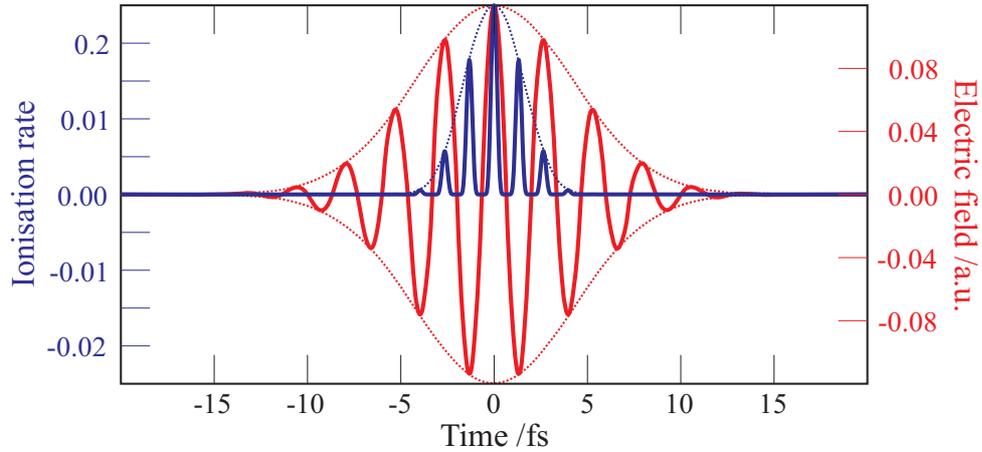


Figure 4.2: Time gating due to the high nonlinearity of the laser induced ionisation rate of Argon. The electric field (solid red) corresponds to a FWHM intensity duration of 7 fs (dotted red) and a peak intensity of $5 \times 10^{14} \text{ Wcm}^{-2}$. The FWHM duration of each electron wavepacket (solid blue) is 370 as, whilst the envelope (dotted blue) is 3.7 fs.

Immediately one can see that the ability to generate short pulses arises in the fact that ionisation is extremely nonlinear, and thus can act as a temporal gate — switching on and off faster than the period of the electric field. To illustrate this point, the electric field of a 7 fs pulse (FWHM) with a peak intensity of $5 \times 10^{14} \text{ Wcm}^{-2}$ and the ionisation rate in argon calculated using the ADK formula is plotted in figure 4.2. The FWHM duration of each individual wavepacket is approximately 370 as, whilst the FWHM duration of the envelope of the ionisation peaks is 3.7 fs. Thus the burst of electron wavepackets is confined to a timescale much smaller than the pulse used, and each individual burst is limited to a fraction of the driving laser period. It is this time-gating that provides the short timescale and attosecond resolution.

(ii) Electron motion

Once the electron wave packets are ‘born’ into the continuum, they are then treated with classical mechanics in the second stage of the model. Thus the equations of motion governing the electron in the continuum for an electric field, $\xi(t) = \xi_0 \cos(\omega t) \hat{\mathbf{x}} + a\xi_0 \sin(\omega t) \hat{\mathbf{y}}$, is given by equations 4.2 for the position and 4.3 for the velocity. For linearly polarised light, $a = 0$ and for circular polarisation $a = \pm 1$. The effects of the magnetic fields for the

pulse and the electron can be neglected for such intensities (i.e. $I_0 \ll 10^{18} \text{ Wcm}^{-2}$, where I_0 is the peak intensity so that the velocity of the electron is much less than the speed of light: $v_e \ll c$).

$$\begin{aligned} x &= \frac{e\xi_0}{m_e\omega^2} [\cos(\phi) - \cos(\phi_0) + (\phi - \phi_0) \sin(\phi_0)] \\ y &= a \frac{e\xi_0}{m_e\omega^2} [\sin(\phi) - \sin(\phi_0) + (\phi - \phi_0) \cos(\phi_0)] \end{aligned} \quad (4.2)$$

$$\begin{aligned} v_x &= -\frac{e\xi_0}{m_e\omega} [\sin(\phi) - \sin(\phi_0)] \\ v_y &= a \frac{e\xi_0}{m_e\omega} [\cos(\phi) - \cos(\phi_0)] \end{aligned} \quad (4.3)$$

where e and m_e are the charge and rest mass of the electron; $\phi = \omega t$ is the phase of the electric field and $\phi_0 = \omega t_0$ is the phase at the instance of ionisation. The electrons are assumed to be born with zero velocity and in the vicinity of the ion (i.e. $v(\phi_0) = x(\phi_0) = 0$). Such an assumption is justified for long wavelengths by the comparison of above-threshold-ionisation (ATI) experiments.^[180–182]

It is clear that the equations (4.3) for the velocity show immediately that the use of circularly polarised light will prevent the electron from returning from the ion. Thus the first result of the model is that HHG is polarisation dependent, which has been found to be true experimentally.^[183]

Solutions to equation (4.2) show that for each laser cycle, half of the electrons, corresponding to emission times $0 \leq \phi_0 < \pi/2$ and $\pi \leq \phi_0 < 3\pi/2$, that are field ionised by linearly polarised light will pass the ion at $x = y = 0$ once during the first laser period during ionisation. The other half, corresponding to emission times $\pi/2 \leq \phi_0 < \pi$ and $3\pi/2 \leq \phi_0 < 2\pi$, will never return to the ion core and take no further part in the process. Some electron trajectories are shown in figure 4.3. The kinetic energy of the electron at the time of recollision is given by

$$T = 2U_p [\sin(\phi) - \sin(\phi_0)]^2 \quad (4.4)$$

where

$$U_p = (e\xi_0)^2 / (4m_e\omega^2) \quad (4.5)$$

is the pondermotive energy. Solving for the kinetic energy for times when the electrons return to the nucleus results in a maximum energy for emission at $\phi_0 = 17^\circ, 196^\circ \dots$, giving a maximum value of the kinetic energy of $3.17U_p$ and are thus dubbed the cut-off electrons as this is the maximum classically allowed energy. In the quantum model, energies beyond the cut-off are allowed, but with a rapid decrease in yield. Electrons that are born within the same half-cycle prior to the cut-off electrons return to the core approximately one period later, whereas electrons that are born after the cut-off and return to the ion do so within half a period. Therefore these trajectories are dubbed the long and short trajectories respectively, relating to the time the electrons spend in the continuum.

If the returning electron energy exceeds the e-2e scattering energy, then it is possible to collisionally ionise the parent ion, resulting in correlated two-electron ejection. Ex-

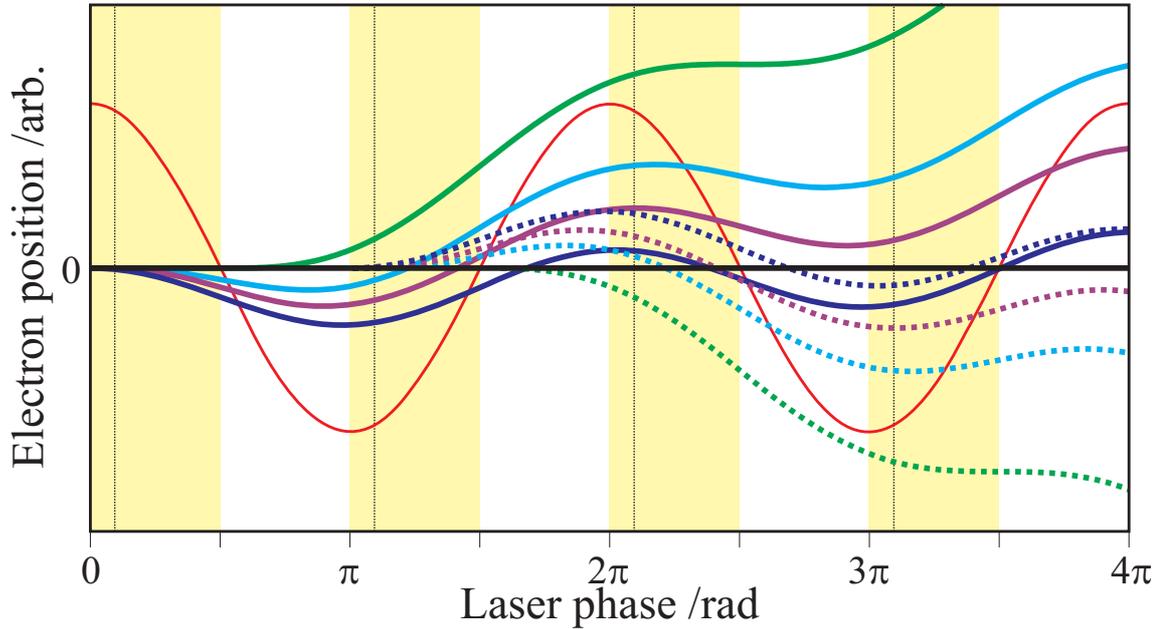


Figure 4.3: Classical electron trajectories from strong field ionisation, showing the electric field (red), cut-off trajectories (purple), long trajectories (blue), short trajectories (cyan) and trajectories that do not contribute to harmonic emission (green) — only electrons born within the shaded yellow regions return to the core. The vertical dotted lines show the birth time of the cut-off electrons. Solid and dotted trajectories correspond to different half-cycle birth times.

periments have shown that the yield of single electron and two-electron ionisation as a function of intensity match the model, showing its validity in prediction of the electron energy. Alternatively, the electron can undergo inelastic scattering, absorbing energy from the laser field. Both of these interactions yield emission of ATI photoelectrons and thus do not contribute to high harmonic generation, but are worth studying in order to check the validity of the simple model.^[184]

(iii) Recombination

The final process that can occur is that the returning electron recombines with the parent ion, emitting a photon with an energy equal to the sum of the electron kinetic energy and the atomic ionisation potential, I_p , i.e. a frequency of $\omega_e = (I_p + T) / \hbar$. If the ground state has negligible depletion, then the wave packet will pass the ion in the same way every laser period, thus the light emitted must be a harmonic of the driving laser frequency. In fact, due to the anti-symmetry of the process — the electron recombines from opposite sides every half laser cycle — only odd harmonics are present. Not only does this model predict harmonic emission, but it recovers the harmonic cut-off law: the maximum allowed photon energy is given by $\omega_{\text{cut-off}} = (I_p + 3.17U_p) / \hbar$, which has also been found experimentally.^[185]

Harmonic spectrum

The harmonic emission is calculated from the expectation value of the dipole operator, $\langle \psi | er | \psi \rangle$, where r is the dipole separation. It is assumed that the wavefunction can be written as the sum of the ground and continuum wavefunctions, $\psi = \psi_g + \psi_c$. The dominant contribution to the harmonics is derived from the $\langle \psi_g | er | \psi_c \rangle + c.c.$ terms. The assumptions: (1) the ground state is negligibly depleted, (2) the continuum wavefunction is constructed from the correspondence principle (i.e. the electron is treated as a classical particle in the continuum), (3) emission for phases $\phi_0 < 17^\circ$ and phases $\phi_0 > 17^\circ$ are added incoherently, and (4) the wavepackets spread linearly in time are used to calculate

the harmonic spectra. For hydrogen-like atoms, the continuum wavefunction can then be expressed as

$$\psi_c(x \approx 0, t) = \sum_h A_p(x \approx 0, t) \exp \left\{ i \frac{p_h(x \approx 0, t)}{\hbar} x - i \left[\frac{p_h(x \approx 0, t)^2}{2m_e \hbar} + I_p \right] t \right\}, \quad (4.6)$$

where the index of the sum is the harmonic order, h , and $p_h(x \approx 0, t)$ is the momentum of the electron that generates that particular harmonic. A_p is the amplitude normalisation factor that corrects for the wavepacket spreading in the continuum and must satisfy

$$\int d^3x A_p^2 = \left[\int_{E_h - \hbar\omega}^{E_h + \hbar\omega} dE P(E) \right] / V, \quad (4.7)$$

where $V = \pi r^2 p \delta t / m_e$ and $P(E)$ is the probability per unit energy per laser period of finding the electron traversing the ion with energy E , which is derived from the ADK ionisation probability and equations (4.2) and (4.3). Equation (4.6) states that the phase of harmonic h is given by the phase of an electron with canonical momentum $p_h(x \approx 0, t)$ resulting in emission of that frequency. The strength of the yield is dependent on the diffusion of electrons, and hence how long they spend in the continuum. Thus harmonic emission is dominated by the electrons that recollide on their first return.

Corkum found that such a simple model recovers many essential features of the harmonic emission: the dipole strength and thus the harmonic yield as a function of frequency matches that found in experiments, namely a rapid decay in yield for the first few harmonics, followed by an extended plateau until the cut-off, which the model recovers. The emission is also heavily dependent upon the ionisation, which acts as a temporal gate that results in an extremely short temporal resolution.

4.1.2 Quantum orbit model

Although the three step model provides a physical picture of HHG, and recovers some of the essential features, accurate modelling of the process requires a complete quantum

mechanical picture. In this section, I describe such a quantum model, as put forward by Lewenstein^[186]. Using some approximations, I then describe how this model recovers the assumptions made in the three step model and how it can be used to provide an intuitive quantum model.

The time dependent Schrödinger equation for the atom-field interaction is given in the length gauge by:

$$i \frac{\partial}{\partial t} |\Psi(\mathbf{r}, t)\rangle = \left[-\frac{\mathbf{p}^2}{2} + V(\mathbf{r}) + \boldsymbol{\xi}(t) \cdot \mathbf{r} \right] |\Psi(\mathbf{r}, t)\rangle, \quad (4.8)$$

where $V(\mathbf{r})$ is the atomic potential, $\boldsymbol{\xi}(t)$ is the laser electric field and $\mathbf{p} = -i\nabla$ is the kinetic energy operator. All quantities are in atomic units ($\hbar = 1$, $e = 1$, $m_e = 1$, $a_0 = 1$) and time is in units of radians (i.e. scaled by the central frequency). Thus the pondermotive energy is given by $U_p = \xi_0/4\omega^2$. Upon making the following assumptions: (1) only the ground state contributes to the harmonic emission, (2) the ground state experiences negligible depletion and (3) strong field approximation (SFA) — the electron is treated as a free particle in the electric field (i.e. the atomic potential is neglected in the treatment of the electron in the continuum), the expectation value of the dipole operator, $x_{\hat{\mathbf{n}}}(t)$ along the direction $\hat{\mathbf{n}}$ can be found exactly:

$$\begin{aligned} x_{\hat{\mathbf{n}}}(t) &= \langle \Psi(x, t) | x_{\hat{\mathbf{n}}} | \Psi(x, t) \rangle \\ &= i \int_0^t dt' \int_{-\infty}^{\infty} d^3\mathbf{p} \hat{\mathbf{n}} \cdot \mathbf{d}^*(\mathbf{p} - \mathbf{A}(t)) \\ &\quad \times \boldsymbol{\xi}(t') \cdot \mathbf{d}(\mathbf{p} - \mathbf{A}(t')) e^{iS(\mathbf{p}, t, t')} + \text{c.c.} \end{aligned} \quad (4.9)$$

$$S(\mathbf{p}, t, t') = \int_{t'}^t dt'' \left(\frac{[\mathbf{p} - \mathbf{A}(t'')]^2}{2} + I_p \right) \quad (4.10)$$

where I have introduced the canonical momentum in terms of the vector potential, $\mathbf{A}(t) = -\int_{-\infty}^t dt' \boldsymbol{\xi}(t')$ of the laser field and the electron velocity, \mathbf{v} :

$$\mathbf{p} = \mathbf{v} + \mathbf{A}(t). \quad (4.11)$$

The dipole matrix element depends on the model used for the atomic potential and for hydrogen-like atoms is given by

$$\mathbf{d}(\mathbf{p}) = i \frac{2^{7/2} \alpha^{5/4}}{\pi} \frac{\mathbf{p}}{(\mathbf{p}^2 + \alpha)^3}, \quad (4.12)$$

where α is a scaling factor to match the ionisation potential of the atom, I_p , such that $\alpha = 2I_p$. Equation (4.9) has an intuitive physical interpretation: (1) the second scalar product, $\boldsymbol{\xi}(t') \cdot \mathbf{d}(\mathbf{p} - \mathbf{A}(t'))$, is the probability amplitude for an electron to make the transition into the continuum at time t' with a momentum \mathbf{p} ; (2) the electron wavefunction is then propagated until time t , acquiring a phase equal to the quasi-classical action, $S(\mathbf{p}, t, t')$ according to equation (4.10); (3) finally the electron recombines at time t with a probability amplitude equal to $\mathbf{d}^*(\mathbf{p} - \mathbf{A}(t))$. According to the SFA, the effect of the Coulomb potential is negligible and thus the quasi-classical action describes the motion of an electron moving freely in the electric field with a constant canonical momentum.

It is possible to simplify equation (4.9) via the use of saddle point (a.k.a. stationary phase) approximations (SPA).^[187] For $(t - t')$ of the order of one laser period, the quasi-classical action varies much more rapidly than any other factor in equation (4.9). Therefore, the dominant contribution to the integral over \mathbf{p} comes from the stationary points of the classical action:

$$\nabla_{\mathbf{p}} S(\mathbf{p}, t, t') = 0, \quad (4.13)$$

where $\nabla_{\mathbf{p}} = \sum_i \partial / \partial p_i$. However, $\nabla_{\mathbf{p}} S(\mathbf{p}, t, t')$ is the difference between the position of the free electron at times t and t' , thus giving the first SPA:

$$\nabla_{\mathbf{p}} S(\mathbf{p}, t, t') = x(t) - x(t'). \quad (4.14)$$

The stationary points of the quasi-classical action correspond to momenta \mathbf{p} for which the electron born at time t' returns to the same position at time t . It is evident that $x(t)$ must be close to the origin as this is the only place where transitions to and from the ground

state can occur. This follows mathematically from the fact that the Fourier transform of the dipole matrix elements are localised around the nucleus on a scale comparable to the Bohr radius. The physical interpretation of equation (4.14) is that the dominant contribution to the harmonic emission comes from electrons that tunnel away from the nucleus, but then re-encounter it whilst oscillating in the electric field, thus justifying one of the basic assumptions of the 3-step model. Solving equation (4.13), one obtains the stationary value for the momentum,

$$\mathbf{p}^{\text{st}}_{\tau} = \int_{t-\tau}^t dt' \mathbf{A}(t'), \quad (4.15)$$

and the stationary value for the action,

$$S^{\text{st}}(t, \tau) = \int_{t-\tau}^t dt'' \left(\frac{[\mathbf{p}^{\text{st}} - \mathbf{A}(t'')]^2}{2} + I_p \right), \quad (4.16)$$

where the ‘orbit time’ $\tau = t - t'$ is the difference between the emission and recombinations times.

Restricting to the case of linearly polarised light only, i.e. $\boldsymbol{\xi}(t) = \xi_x(t)$ and $\mathbf{d}(\mathbf{p}) = d_x(\mathbf{p})$, substituting for the stationary action and momentum into equation (4.9) and extending the integration over infinity yields

$$x^{\text{st}}(t) = i \int_0^{\infty} d\tau \left(\frac{\pi}{\epsilon + i\tau/2} \right)^{3/2} d_x^* (p_x^{\text{st}}(t, \tau) - A_x(t)) \\ \times d_x (p_x^{\text{st}}(t, \tau) - A_x(t - \tau)) E(t - \tau) e^{-iS^{\text{st}}(t, \tau)} + \text{c.c.}, \quad (4.17)$$

The first term in equation (4.17), $[\pi/(\epsilon + i\tau/2)]^{3/2}$, results from the SPA and physically corresponds to electron wave packet diffusion in three dimensions. The effect of electron diffusion is effectively to limit the integral to return times of less than an optical cycle. Hence, the dominant contributions to the harmonic emission is from the two electron trajectories that result in the same momentum upon recombination within one optical

cycle, similar to the three-step model. Fourier analysis shows that the variation of S^{st} as a function of orbit time, τ , matches almost exactly the energy gained by the electron in the classical Newtonian treatment of the three-step model.^[186]

It is possible to simplify equation (4.17) via another SPA, this time over the orbit time, τ , provided $I_p \ll U_p$. The stationary values of the quasi-classical action given by equation (4.16) yields two more saddle point equations:

$$\frac{[\mathbf{p}^{\text{st}}(t, \tau) - \mathbf{A}(t)]^2}{2} + I_p = 0, \quad (4.18)$$

$$\frac{[\mathbf{p}^{\text{st}}(t, \tau) - \mathbf{A}(t)]^2}{2} + I_p = \omega_h. \quad (4.19)$$

and are a result of conservation of energy. Equation (4.18) corresponds to the value of the initial velocity:

$$\mathbf{v}(t - \tau) = \mathbf{p}^{\text{st}}(t, \tau) - \mathbf{A}(t - \tau) = 0. \quad (4.20)$$

Thus justifying the second assumption of the 3-step model: the electrons that contribute most to the harmonic emission are born with zero initial velocity. Equation (4.19) is obtained via a Fourier transform of the dipole and physically means that the harmonic with frequency ω_h is only emitted at times t at which the electrons returning to the nucleus have appropriate energy.

Further features of the harmonic emission can be revealed by Fourier analysis and saddle-point approximations. One feature is the recovery of the harmonic cut-off law, $\omega_{\text{cut-off}} = (I_p + 3.17U_p) / \hbar$. These features, depletion of the ground state and other features are beyond the scope of this thesis and are covered in more detail in^[186].

4.1.3 Time dependent Schrödinger equation

Solutions to the time dependent Schrödinger equation provide the most complete quantum mechanical description of the physics involved in HHG. However, due to constraints of the computational resources, the system is typically only solved in one spatial dimension. The

one dimensional system being solved is

$$i \frac{\partial}{\partial t} |\Psi(x, t)\rangle = \left[\frac{(p_x + A_x(t))^2}{2} + V(x) \right] |\Psi(x, t)\rangle, \quad (4.21)$$

where atomic units throughout unless otherwise specified, and the terms are the same as for section 4.1.2.

The Coulomb potential, $V(x)$ is normally modelled by a ‘soft-core’ potential, removing the singularity of the $1/r$ potential:

$$V(x) = \frac{-1}{\sqrt{x^2 + a^2}}, \quad (4.22)$$

where a is a softening parameter used to match the ionisation potential, I_p , of the atom in question. For example, a value of $a^2 = 0.673$ is used for Neon, which has an ionisation potential of $I_p = 21.5$ eV.

The value of $|\Psi(x, t = 0)\rangle$ used to seed equation (4.21) is calculated from the field-free ground state solution to the time-independent eigenvalue problem:

$$H |\Psi(x)\rangle = E |\Psi(x)\rangle, \quad (4.23)$$

$$H = -\frac{\nabla^2}{2} + V(x). \quad (4.24)$$

Although it is possible to solve this using standard eigenvalue problems, this is typically not desirable as the grid size used for the spatial dimension (for example, $N = 4096$)^[188] makes this unfeasible. However, the ground state solution only is required, which can be determined via the convergent iteration of the free-particle Green’s function multiplied by the potential.^[189] Equation (4.21) is then propagated using either a Crank-Nicholson method^[190] or the split-operator method^[191].

The harmonic spectrum, $\tilde{E}(\omega)$, is calculated from the wavefunction via the inverse Fourier transform of the expectation value of the dipole acceleration operator, $\ddot{D}(t)$:^[192]

$$\tilde{E}(\omega) = \omega^{-2} \mathfrak{F}^{-1} \left\{ \ddot{D}(t); t \rightarrow \omega \right\}, \quad (4.25)$$

$$\ddot{D}(t) = \langle \Psi(x, t) | -\nabla V(x) + \xi(t) | \Psi(x, t) \rangle. \quad (4.26)$$

The numerical solutions to the TDSE are thought to provide the most accurate description of the single atom response to the intense electric field, and indeed reproduce the same features found experimentally. However, there are significant difference between the measured and predicted spectra, which are usually explained when macroscopic propagation of the harmonic emission from every atom in the gas is also taken into account.

4.2 HHG: Macroscopic response

The three models discussed in the previous section only deal with the interaction of the intense low frequency field with a single atom. The harmonic field at the exit of the generating medium will, however, be described by the collective response of this interaction for every atom within the generating medium. As such, it is necessary to develop a model for the propagation of the electromagnetic field through the medium, incorporating the single-atom response for every atom, resulting in a complete macroscopic theory for high harmonic generation.

The equation of motion governing the propagation of the electromagnetic field is given by 2.34 in section 2.2.1. The polarisation of the medium can be calculated via solving Schrödinger's equation for each atom whilst simultaneously solving the wave equation. However, such an approach is extremely computationally demanding and is not realistic for personal computers. Fortunately it is possible to simplify the procedure due to the fact that the harmonic field is extremely weak compared to the low frequency field. As such,

it is possible to break the solution down into solving the propagation of the low frequency field, taking into account ionisation of the medium, then solving the propagation of the harmonic field, using the previously calculated intensity of the fundamental field.

The propagation of the fields can be calculated efficiently according to the recipe of Geissler *et al.*^[193] Using the slowly evolving wave (SEW) approximation, which is based on the polarisation and electric field varying slowly compared to the wavelength, λ_0 (i.e. $|\partial A/\partial z| \ll |A\omega/c|$ where $A = |E|$ is the envelope of the complex electric field),^[194] it is possible to reduce the second-order wave equation to a first-order propagation (FOP) equation that eliminates backward-propagating solutions:

$$\frac{\partial}{\partial z}\xi(r_{\perp}, z, \tau) - \frac{c}{2}\nabla_{\perp}^2 \int_{-\infty}^{\tau} d\tau' \xi(r_{\perp}, z, \tau') = -\frac{2\pi}{c}\frac{\partial}{\partial \tau}P[\xi(r_{\perp}, z, \tau)]. \quad (4.27)$$

Equation (4.27) uses Gaussian units in the moving co-ordinate frame: $\tau = t - z/c$, r_{\perp} refers to the radial distance from the optical axis z and ξ is the real electric field for all the propagation equations in this section. The FOP equation makes no assumption on the temporal profile of the pulse used, and as such can be used on arbitrarily short pulses.

4.2.1 Propagation of the intense low frequency field

The polarisation of the medium due to ionisation of electrons is given by Schrödinger's equation: $P = n_0 e \langle \Psi | x | \Psi \rangle$, where n_0 is the atomic particle density.^[193] Although it is possible to solve this numerically, it is much more computationally advantageous to consider the classical motion of the ionised electrons in the low-frequency field. Assuming no collisions, collective plasma processes or relativistic effects are present, the polarisation response reduces to:^[193]

$$\frac{\partial}{\partial t}P(\mathbf{r}, t) = I_p \frac{\frac{\partial}{\partial t}n_e(\mathbf{r}, t)}{\xi(\mathbf{r}, t)} + \frac{e^2}{m_e} \int_{-\infty}^t dt' n_e(\mathbf{r}, t') \xi(\mathbf{r}, t'), \quad (4.28)$$

where the electron density is given by the degree of ionisation:

$$n_e(\mathbf{r}, t) = n_0 \left\{ 1 - \exp \left[- \int_{-\infty}^t dt' W(\xi(\mathbf{r}, t')) \right] \right\}, \quad (4.29)$$

and $W(\xi)$ is the quasi-static ionisation rate given by the ADK formula — equation (4.1). The first term is the ionisation induced loss due to energy being transferred from the field to the ionised electrons. The second term takes into account any phase modulation effects such as blue-shifting and spectral broadening. Thus the full propagation equation becomes:

$$\begin{aligned} \frac{\partial}{\partial z} \xi(r_\perp, z, \tau) - \frac{c}{2} \nabla_\perp^2 \int_{-\infty}^\tau d\tau' \xi(r_\perp, z, \tau') = \\ - \frac{1}{2c} \int_{-\infty}^\tau d\tau' \omega_p^2(r_\perp, z, \tau') \xi(r_\perp, z, \tau') - \frac{2\pi I_p}{c} \frac{\partial}{\partial \tau} n_e(r_\perp, z, \tau) / \xi(r_\perp, z, \tau), \end{aligned} \quad (4.30)$$

where the plasma frequency is given by $\omega_p^2 = 4\pi n_e e^2 / m_e$. This equation is only valid in underdense plasmas, i.e. $\omega_p \ll \omega_0$ and Keldysh parameters $\gamma = \sqrt{I_p / 2U_p} < 1$.

4.2.2 Propagation of the weak harmonic field

The propagation of the weak harmonic field is slightly more complicated because it is generated by the low frequency field as it propagates along the z-axis. Thus the propagation is determined according to:^[188]

$$\begin{aligned} \left[\frac{\partial}{\partial z} + \alpha_{\text{xuv}} \right] \xi_{\text{xuv}}(r_\perp, z, \tau) - \frac{c}{2} \nabla_\perp^2 \int_{-\infty}^\tau d\tau' \xi_{\text{xuv}}(r_\perp, z, \tau') = \\ - \frac{2\pi}{c} n_0(r_\perp, z) \int_{-\infty}^\tau d\tau' \ddot{D}[\xi(r_\perp, z, \tau')] \end{aligned} \quad (4.31)$$

where α_{xuv} is the XUV absorption coefficient given by Henke *et al.*^[195], n_0 is the atomic gas density. The dipole acceleration, $\ddot{D}[\xi(r_\perp, z, \tau)] = \ddot{P}_{\text{xuv}}[\xi(r_\perp, z, \tau)] / n_0(r_\perp, z)$ is calculated from the quantum orbit model or the 1D TDSE solutions discussed in the previous section. Absorption of the XUV field would lead to ionisation of electrons. However, the effect this

has on the propagation of the driving field is negligible because the XUV intensity is extremely low.

4.2.3 Generalised phase-matching for HHG

The method of calculating the macroscopic response in the previous section is computationally extensive and obtaining insights into the physics of the process can become difficult. However, it is possible to understand the macroscopic response in some respect by considering the phase-matching properties of the generating medium. Such a method has been proposed by Balcou *et al.*^[196] and studied by Chipperfield *et al.*^[188].

In order to have a relatively strong harmonic signal at the exit of the generating medium, it is necessary to satisfy phase-matching conditions for the harmonics being produced. This requires the harmonics that are produced at one position in the gas to propagate with the correct phase velocity so that they are in phase with harmonics produced further downstream, resulting in constructive interference. Mathematically, one requires that the phase of a harmonic produced at position \mathbf{r}_1 and after propagation to position \mathbf{r}_2 to be in phase with harmonics generated at \mathbf{r}_2 . Provided the distance between the two positions is small, such that the phase velocity of the harmonics is constant over that distance, this can be written as:

$$\arg \left[E_q(\mathbf{r}_1) e^{i\mathbf{k} \cdot (\mathbf{r}_2 - \mathbf{r}_1)} \right] = \arg [E_q(\mathbf{r}_2)], \quad (4.32)$$

where $E_q(\mathbf{r})$ is the complex electric field of harmonic q at position \mathbf{r} and the wave vector is given by the spatial gradient of the harmonic spatial phase:

$$\mathbf{k} = \nabla \phi_q(\mathbf{r}). \quad (4.33)$$

In perturbative harmonic generation, phase-matching is satisfied when the harmonic wave vector matches the sum of the wave vectors of the driving fields, i.e. $\Delta \mathbf{k}_{pert} =$

$\mathbf{k}_q - \sum_i \mathbf{k}_i = 0$. In HHG however, the atomic dipole phase, ϕ_d is not just a function of the driving field, but also contains a contribution from the quasi-classical action due to the electron paths in the continuum. The phase of the dipole is given by

$$\phi_d = q\omega t_f - \frac{S(p, t_i, t_f)}{\hbar}. \quad (4.34)$$

The first term is due to the delay between recombination of the electron at time t_f and the reference time of the driving laser. The second term is due to the phase acquired by the electron on its sojourn from the nucleus. Thus the phase-matching condition is given by

$$\Delta\mathbf{k} = \mathbf{k}_q - q\mathbf{k}_1 - \nabla\phi_d. \quad (4.35)$$

Perfect phase-matching would require $\Delta k = 0$. For HHG, it is not possible to satisfy this condition over the whole generating region, due to the complexity of the fundamental, harmonic and dipole phase (as will be discussed below). Thus it is important to consider the coherence length, $L_{coh} = \pi/|\Delta\mathbf{k}|$, which is the length over which harmonics being generated are in phase with the harmonics previously generated upstream. For regions with a low coherence length, harmonics will rapidly dephase and thus not contribute significantly to the overall harmonic signal. Harmonics generated in regions with a large coherence length thus dominate the final signal.

Phase of the fundamental field

For a focussed beam, which is necessary to reach the intensity required for HHG, the phase of the fundamental beam is spatially dependent. Considering Gaussian beam propagation only, the local wave vector is given by

$$\mathbf{k}_G(r, z) = k_0\hat{\mathbf{z}} + \nabla \arg \left[\frac{1}{b + i2z} \exp \left(-\frac{k_0 r^2}{b + i2z} \right) \right], \quad (4.36)$$

where $b = 2z_R$ is the confocal parameter (twice the Rayleigh range), k_0 is the fundamental wave vector, r is the radial co-ordinate and z is the propagation direction. However, it is necessary to remember that the actual local wave vector can be much more complicated. For example, defocusing due to the free electrons generated via ionisation from the driving laser, spatial phase resulting from poor beam quality and STC can significantly change the wavefront. Ionisation induced defocusing will remain negligible if the driving intensity is smaller than the saturation intensity, $I \ll I_{\text{sat}}$. However, especially for few-cycle pulses, STC and a complex spatial phase may require numerical propagation of the driving field.

Phase of the harmonic field

There are two significant factors concerning the harmonic field that affect the phase-matching conditions for HHG. The first of these is from the dispersion due to the ionised medium. The induced plasma results in a modified refractive index given by

$$n = \sqrt{1 - \frac{\omega_p^2}{\omega^2}}, \quad (4.37)$$

where ω_p is the plasma frequency. This results in a phase mismatch between the fundamental and harmonic fields. Hence, the wave vector of the harmonic field is not necessarily q times the fundamental field, i.e. $\mathbf{k}_q \neq q\mathbf{k}_1$.

The second factor comes from the atomic phase resulting from the quasi-classical action of the electron, which is not present in perturbative harmonic generation. It has been shown by Lewenstein *et al.*^[197] that the atomic phase is approximately linearly proportional to the intensity of the fundamental field. This can be shown using the three-step model or the quantum orbit model. The phase acquired by the electron is proportional to the time the electron spends in the continuum, which for a given harmonic will, to first order, be linear with intensity. As a result of the spatial intensity distribution of the driving field, the phase due to the action will also be spatially dependent. This in turn results in a spatially dependent effective wave vector for the quasi-classical actions, given

by

$$\begin{aligned}
\mathbf{K}(r, z) &= \nabla \phi_d(r, z) \\
&= \frac{\nabla S(p, t_i, t_f, r, z)}{\hbar} \\
&\simeq \nabla \left[I_0 \frac{w_0^2}{w^2(z)} \exp\left(-\frac{2r^2}{w^2(z)}\right) \right], \tag{4.38}
\end{aligned}$$

where w_0 is the beam waist and $w(z)$ is the beam width at z , given by equation (2.15), and I_0 is the peak intensity.

Phase-matching of harmonics

The model proposed by Balcou^[196] has proven useful in explaining the behaviour of the harmonic emission. It has been used to describe how the harmonic yield varies as a function of the driving laser's focal position relative to the gas jet. For example, the position of maximum harmonic yield occurs in two positions, either just before or just after the focus, contrary to perturbative harmonic generation that has a maximum at the focal position. This is purely attributed to the atomic phase. In addition, the angular and spatial distributions of the harmonics vary significantly with focal position relative to the gas jet. This can also be explained, to some extent, through the phase-matching process.

Figure 4.4 shows four phase-matching geometries, corresponding to generating regions: (1) at the centre of the focus, (2) on-axis after the laser focus, (3) on-axis before the laser focus and (4) off-axis before the focus. In all cases, the laser intensity is assumed to be low such that ground state depletion and ionisation effects are negligible. Note however, that it is relatively straightforward to include these effects into the model. Figure 4.4 (a) and (b) show the local wave vectors of the fundamental and atomic dipole respectively. Before the focus, the local wave vectors of the fundamental field converge slightly towards the focus, but diverge slightly after the focus. The local atomic wave vectors, in contrast, always diverge from the centre of the focus. Thus it is clear that different spatial regions will result in different phase-matching conditions. For this analysis, one can consider the generating

medium having a thickness that is small compared to the confocal parameter, which is generally true in practice. For example, typical confocal parameters used in experiments

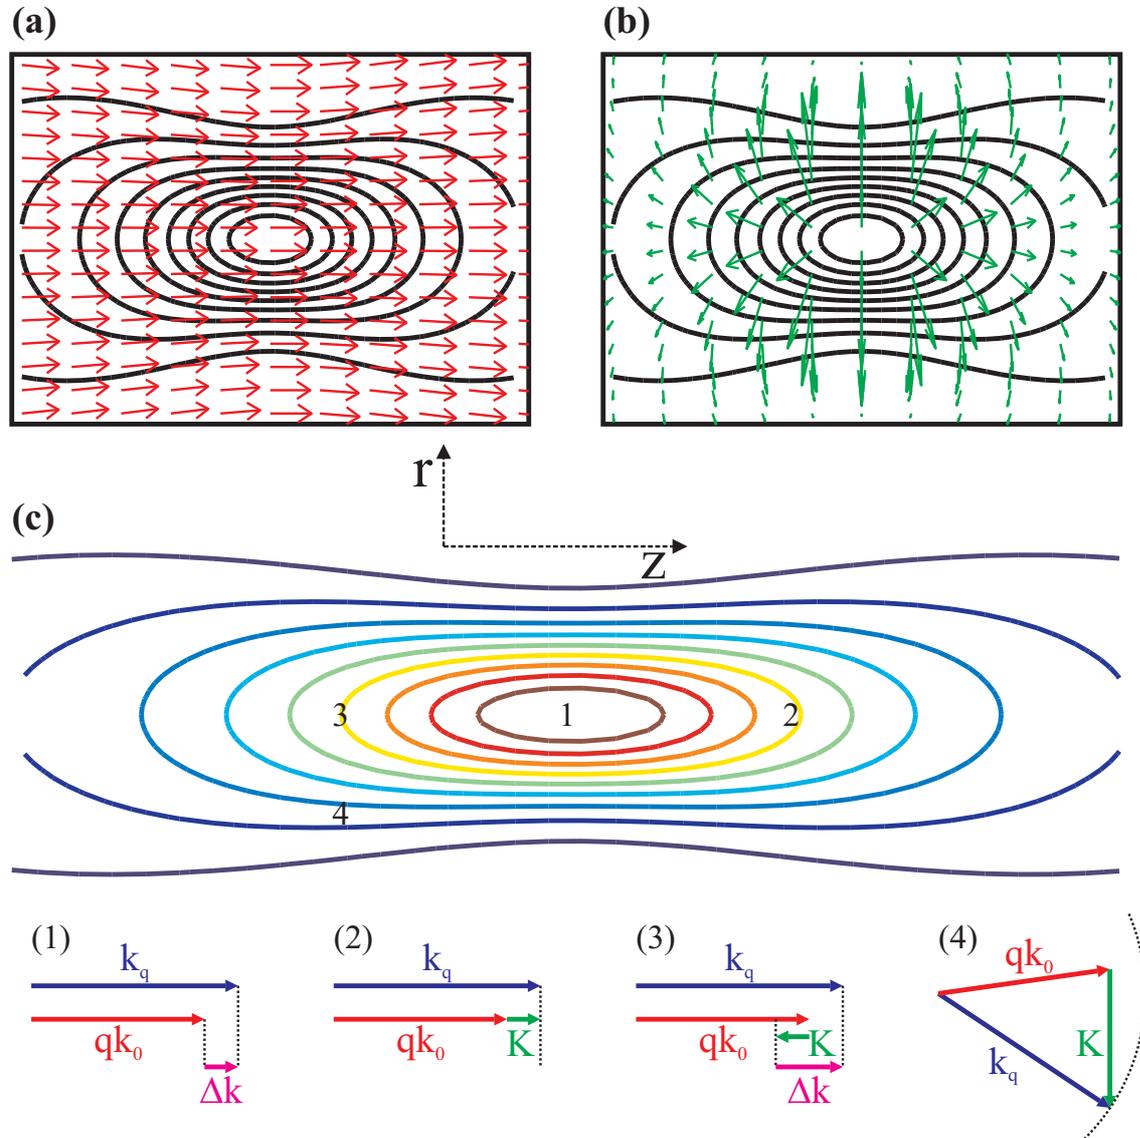


Figure 4.4: Schematic of generalised phase-matching for high harmonic generation. The propagation, z , is to the right and radial direction, r is vertical. (a) The local wave vectors due to focusing of a Gaussian beam (red arrows). The contours of the intensity is shown for reference (solid black). (b) Local wave vectors due to the atomic phase (green arrows). (c) Contours of beam intensity and locations (1–4) for different phase-matching conditions: (1) at the focus — no phase-matching occurs due to the different lengths of the harmonic wave vector, k_q and the fundamental wave vectors, qk_0 ; (2) on-axis after the focus — the atomic phase, K , allows for phase-matching; (3) on-axis before the focus — the atomic phase is aligned in the opposite direction to (2) and thus phase-matching cannot occur; and (4) off-axis before the focus — non-collinear phase-matching is fulfilled, allowing generation off-axis at an angle to the propagation direction.

are on the order of several centimetres; the interaction length for a gas jet is on the order of several hundred micrometres, i.e. up to a factor of one hundred. Figure 4.4 (c) shows the intensity contours for a focused Gaussian beam and the locations of the four generating regions. Below the contours are phase-matching diagrams for the four locations.

At the centre of the focus (1), the phase of the fundamental results in a phase mismatch equal to $2q/b$ due to focusing. However, there is no wave vector from the atomic phase to compensate the phase mismatch and thus phase-matching is not fulfilled. Just after the focus (2), the atomic phase results in an effective wave vector in the direction of propagation and thus compensates for the focusing phase mismatch. Hence collinear phase-matching is fulfilled. Before the focus (3), the atomic wave vector adds to the phase mismatch and thus collinear phase-matching cannot be fulfilled. However, off-axis before the focus (4), the combination of the atomic and fundamental wave vectors allows non-collinear phase-matching to be fulfilled. Thus it is clear that there exists two regions of high efficiency for HHG: (i) on-axis after the focus and (ii) off-axis before the focus. Not only do these two scenarios result in efficient harmonic generation, but they also result in different spatial distributions; the former generating low divergence harmonics on axis, the latter generating a higher divergent annular beam.

Another phenomenon that can be attributed to phase-matching conditions in HHG is the selection of different quantum trajectories. For example, the harmonics that are produced can either be generated in the cut-off regime or in the plateau regime. The former scenario occurs when $U_p \lesssim (\omega_q - I_p)/3.2$, where ω_q is the frequency of the observed harmonic in atomic units. In this regime, only one trajectory is associated with production of harmonics. If however, $U_p > (\omega_q - I_p)/3.2$, then the harmonic belongs to the plateau regime, and thus multiple quantum trajectories contributed to the emitted field. Each trajectory corresponds to a different orbit time (i.e. the time that the electron has spent in the continuum). The atomic phase will vary with intensity with a different constant of

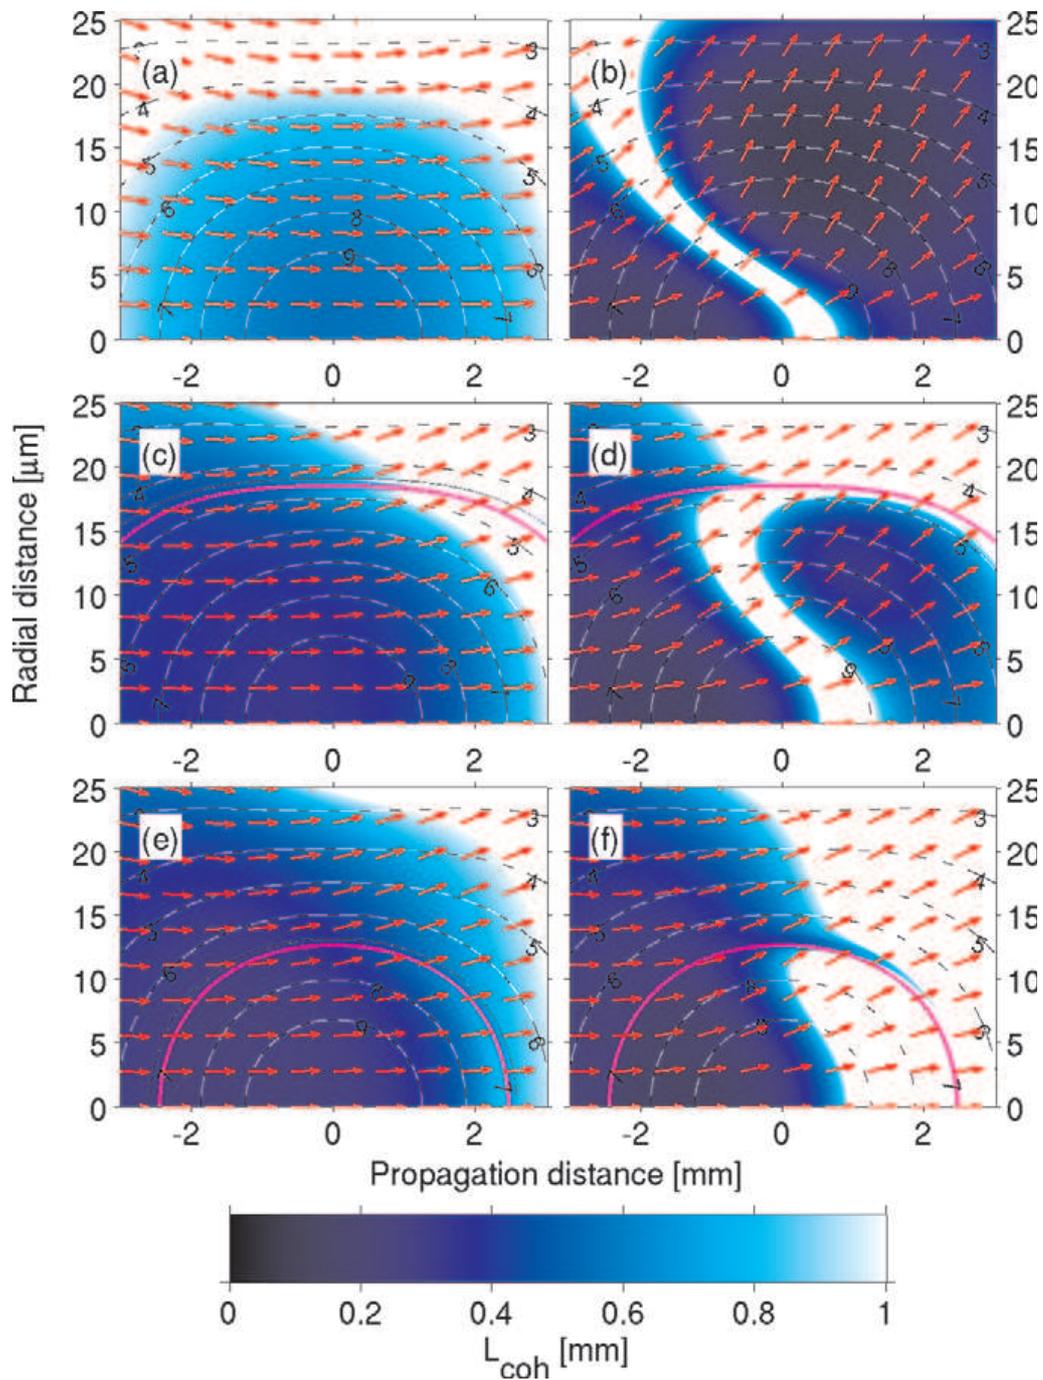


Figure 4.5: Phase-matching maps for HHG. White regions correspond to a long coherence length, dark blue to a short coherence length. The arrows represent the direction of best phase-matching. The thick magenta line represents the boundary between the cut-off and plateau regime for that particular harmonic. Figures (a), (c) and (e) correspond to the short trajectory, (b), (d) and (f) for the long for harmonics 21: (a) & (b), 35: (c) & (d) and 45: (e) & (f). Images taken from Chipperfield.^[188]

proportionality that is itself linearly dependent on the orbit time, i.e.

$$\phi_q^i = -\alpha_q^i I(r, t), \quad (4.39)$$

where the superscript denotes the i -th trajectory and $\alpha_q^i \propto \tau_q^i$. As a result of the dependence of the atomic phase on the quantum trajectory, each quantum trajectory will have different phase-matching conditions. The net result is that different quantum trajectories and generating regimes will be generated in different regions of the focus and with different divergence properties. This behaviour is clearly illustrated in figure 4.5, which shows the regions of long coherence length (i.e. small phase mismatch) for different harmonics and the two shortest quantum trajectories for the plateau regime. As the intensity in the wings of the pulse will be much lower than the centre, the outer spatial regions of the pulse will generate the given harmonic in the cut-off regime (and thus only one trajectory contributes). The spatial boundary between the plateau and cut-off regimes is also shown for each harmonic.

As one can see, the different trajectories in the plateau regime have very different phase-matching properties. The long trajectory is produced efficiently with a large divergence for $z \lesssim 0$, whereas the short trajectories are produced efficiently with a small divergence at $z \gg 0$. The cut-off regime is produced off-axis for low harmonics and on-axis for high harmonics with very different propagation properties to the plateau harmonics.

It is possible to utilise the dependence of the spatial and divergence properties of the harmonics on the trajectory as a result of the phase-matching conditions to measure the atomic dipole phase. The first experiment to demonstrate interference between two quantum trajectories and thus measure the atomic phase was performed by Corsi *et al.*^[198] In the experiment, two well separated sources of HHG are generated in a single atomic gas jet. The two sources are then imaged into a monochromator to select a single harmonic and then allowed to propagate into the far field. As a result, the two sources will diverge and overlap to produce interference fringes with two distinct regions. The central

region corresponded to the short trajectory and the outer to the long trajectory. Using a waveplate/polarising beamsplitter arrangement, it was possible to vary the intensity of one of the beams whilst keeping the other at a fixed intensity to act as a reference. Upon varying the intensity, it was found that the absolute phase of the fringes shifted. However, the shift was different between the two regions. By measuring the relative shift of the fringes between the two regions as a function of intensity, it was possible to deduce the atomic phase. Figure 4.6 shows two interference patterns, corresponding to two different intensities, clearly showing the differential shift in the absolute phase of the fringes.

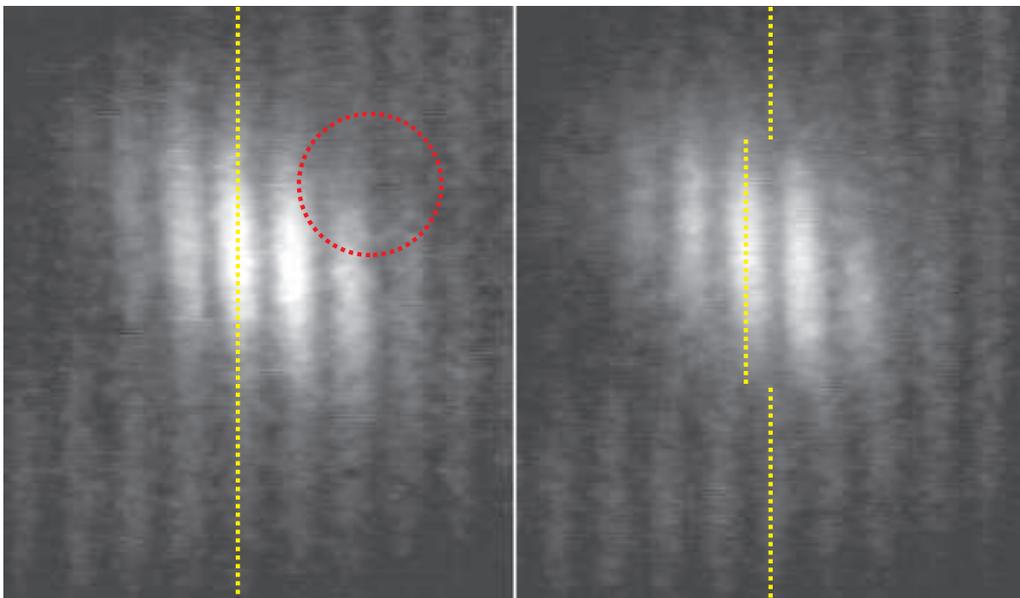


Figure 4.6: Interference fringes showing the intensity dependent dipole phase. Equal intensities were used in the left figure, resulting in no phase difference between the inner and outer regions (as shown by the dashed yellow line). An intensity difference between the two driving pulses of $\delta I = -3 \times 10^{13} \text{ Wcm}^{-2}$ generated the right interference pattern. Note the shift in the fringes between the inner and outer regions (shown by the dashed yellow lines), showing the intensity dependence of the short (inner region) and long (outer region) trajectories. Images taken from Corsi.^[198] The spatially filtered region selected by Zair^[199] is shown by the red circle.

The problem with the technique used by Corsi is that it averages the dipole phase over the whole duration of the generation process. This therefore limits the temporal resolution of the technique. Recently Zair *et al.*^[199] have extended the technique further, removing the temporal averaging and thus improving the temporal resolution. In their experiment, they spatially filtered a section of the harmonics from a single source in the far-field, such

that there was an approximately equal contribution from each trajectory. The spatially filtered harmonics were then imaged in a spectrometer, resulting in interference between different quantum trajectories from the same sample. As short pulses are used to generate the harmonics, the atomic dipole phase is time dependent, in analogy with the spatial dependence due to the finite beam size, resulting in a chirp of the harmonics. Clearly this effect is larger for few-cycle pulses. Therefore, by spectrally resolving the harmonics, it is possible to utilise the time-to-frequency mapping (i.e. chirp) of the harmonics due to the time-dependent dipole phase. As a result, it was possible to measure and control the electron trajectories with a resolution on the order of tens of attoseconds, the shortest timescale currently observed directly. This method promises to be beneficial for studying the dipole phase of more complicated media (e.g. molecules or clusters).

4.3 Attosecond metrology via photoelectron spectroscopy

4.3.1 Autocorrelation

Intensity autocorrelation is arguably the simplest temporal characterisation method and conceptually the same for XUV pulses as for optical pulses. However, due to the wavelength of the light being characterised (sub-100 nm), neither beamsplitters (in the traditional sense) or nonlinear crystal exist for such wavelengths. In order to generate two replicas of the XUV light, it is necessary to: (1) produce two identical sources, (2) spatially split the beam in two or (3) use a grating in a double pass configuration to act as a beamsplitter^[200]. The first method is hard to enforce due to the nonlinearity of the generation process and the high field intensities needed. The second method is sensitive to space-time coupling, which is known to exist in HHG, whilst the latter method is inefficient and geometrically complicated. For the nonlinear medium, it is necessary to use, for example, double ionisation of an atomic gas.^[201] However, the weak intensity and high photon energy of the XUV field results in a low cross-section for the double ionisation process. Hence the signal-to-background is extremely low and the technique is only likely to be viable for the lowest order harmonics.

The major advantages of the intensity autocorrelation are that it is relatively simple to implement and can be considered as a direct measurement on the temporal duration of the pulse, due to the intuitive nature of the method. However, an intensity autocorrelation cannot recover any phase information directly and requires some assumption on the temporal pulse shape. Thus this technique can only be used as an estimate of the temporal duration. Also, due to the need to use double ionisation, it can only be used for a limited range of frequencies, set by the atomic medium used.

4.3.2 Theory of atomic XUV photoionisation in a laser field

In this section, I describe how it is possible to use a low frequency (e.g. NIR) field to act as a phase modulator for XUV pulses in order to characterise the temporal field of the XUV pulse via photoelectron spectroscopy (PES). Use of the low frequency field as a phase modulator allows one to extend some of the temporal characterisation methods in the femtosecond and picosecond domains into the attosecond domain. After describing the ionisation process, I briefly outline a few techniques that can be used for temporal characterisation.

Initially, I consider the ionisation of an atom from the XUV field alone and use the single active electron approximation.^[202] The transition amplitude $a_{\mathbf{v}}$ from the ground state to the final continuum state $|\mathbf{v}\rangle$ with momentum \mathbf{v} can be calculated, via first order perturbation theory for times sufficiently large for the XUV amplitude to vanish, as^[203]

$$a_{\mathbf{v}} = -i \int_{-\infty}^{\infty} dt \mathbf{d}_{\mathbf{v}} \cdot \boldsymbol{\xi}(t) e^{i(W+I_p)t}, \quad (4.40)$$

where $W = \mathbf{v}^2/2$ is the energy of the final continuum state, $\boldsymbol{\xi}(t)$ is the XUV field, $\mathbf{d}_{\mathbf{v}}$ is the dipole matrix element for the transition and I_p the ionisation potential of the atom. Note that atomic units shall be used throughout.

From equation (4.40), it is clear that the photoelectron energy distribution is the convolution of the XUV spectrum with the response of the ionising medium. Using known

or calculated ionisation cross-sections, it is possible to deconvolve for the dependence of $|\mathbf{d}_v|$ on \mathbf{v} to obtain the XUV spectrum. Alternatively, one may use an XUV spectrometer to calibrate this. It is also necessary to correct for the phase of the dipole element for temporal characterisation. However, the relative bandwidth has been small up to now and so it has been possible in general to neglect this. As the pulse duration approaches the response time of the medium (i.e. the pulse spectrum becomes comparable to the medium spectral response function), this must be taken into account.

Using the SFA, it is possible to calculate the spectral distribution in the presence of both an XUV and low frequency field. As the XUV photon energy is much larger than the ionisation potential (i.e. $\omega_{\text{xuv}} \gg I_p$), the SFA is satisfied for even moderate IR intensities. The probability amplitude for a transition to final state $|\mathbf{v}\rangle$ is given by^[21,203]

$$a_v = -i \int_{-\infty}^{\infty} dt \mathbf{d}_{\mathbf{p}(t)} \cdot \boldsymbol{\xi}_{\text{xuv}}(t - \tau) \exp \left\{ i \left[I_p t - \int_t^{+\infty} dt' \frac{\mathbf{p}^2(t')}{2} \right] \right\}, \quad (4.41)$$

where τ is the delay between the IR and XUV fields; and $\mathbf{p}(t) = \mathbf{v} + \mathbf{A}(t)$ is the instantaneous momentum of the free electron in the vector potential $\mathbf{A}(t)$. Note that in the absence of an IR field, equation (4.41) reduces to that of (4.40).

Equation (4.41) has an intuitive picture: the ionisation occurs with a probability related to the XUV field, followed by classical electron motion in the presence of the IR field. This can be seen more easily by rearranging equation (4.41) to obtain

$$a_v(\tau) = -i \int_{-\infty}^{\infty} dt \mathbf{d}_{\mathbf{p}} \cdot \boldsymbol{\xi}_{\text{xuv}}(t - \tau) e^{i\Phi(t)} e^{i\Omega t} \quad (4.42)$$

$$\Phi(t) = - \int_t^{+\infty} dt' \left[\mathbf{v} \cdot \mathbf{A}(t') + \frac{\mathbf{A}^2(t')}{2} \right], \quad (4.43)$$

where $\Omega = W + I_p$. Thus the IR laser acts as a temporal phase modulator for the XUV field. However, due to the scalar product in equation (4.43), the phase modulation is dependent on the angle of detection of the emitted electrons relative to the polarisation of the IR field. It is also clear from equation (4.42) that it is not the envelope of the IR

but the IR electric field itself that acts as a time-nonstationary filter, and thus provides sub-femtosecond temporal resolution.

The effect the phase modulation has on the PE spectrum is dependent on the duration of the XUV field relative to the period of the low frequency laser, giving rise to three regimes of operation. The three regimes correspond to when the XUV duration is: (1) much shorter than, (2) comparable to or (3) much longer than the optical period of the low frequency field. These three scenarios are depicted in figure 4.7. These three regimes result in (1) a shift of the PE velocity distribution, (2) streaking of the PE velocity distribution or (3) the appearance of sidebands in the velocity distribution respectively.

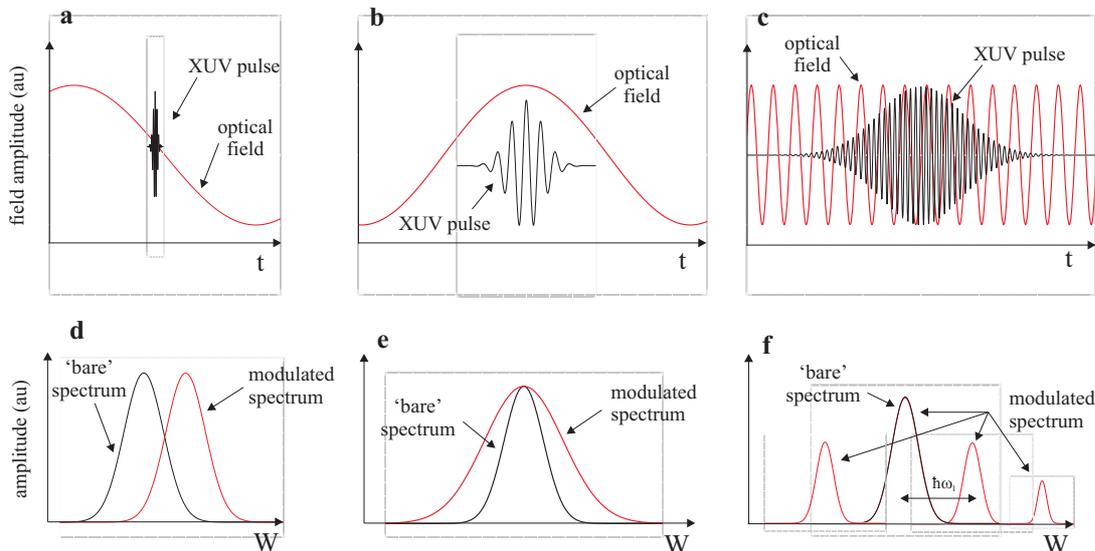


Figure 4.7: Three different regimes of performing characterisation of XUV attosecond pulses via PES. (a)–(c) Three regimes of operation. (d)–(f) Resultant PE spectrum. (a) The XUV pulse is much shorter in duration than the dressing field period and aligned with the zero crossing resulting in a linear temporal phase. (d) The resulting PE spectrum is shifted in frequency by an amount proportional to the gradient of the dressing field. (b) The XUV pulse is slightly shorter than the dressing field period and aligned with the peak/trough of the dressing field, resulting in a primarily quadratic temporal phase. (e) The PE spectrum is broadened by an amount that is dependent on the strength of the dressing field and the XUV pulse duration relative to the dressing field period. (c) The XUV pulse spans many optical cycles of the dressing field. (f) The resulting PE spectrum consists of sidebands separated by the dressing field frequency.

Attosecond spectral shearing interferometry

If the XUV pulse is much shorter than the optical period of the dressing laser field, then it is possible to perform attosecond spectral shearing interferometry — one method of implementing SPIDER via PES.^[204] The method is conceptually extremely simple. Two time-delayed replicas of the XUV pulse are temporally aligned with adjacent zeros of the low-frequency electric field. The delay between the two XUV pulses will generate fringes in the PE spectrum, allowing the phase between them to be extracted via a Fourier filtering routine. The dressing laser field will cause the two PE velocity distributions to spectrally shift in opposite directions, thus generating the spectral shear required for SPIDER.

This technique has all the advantages associated with spectral shearing interferometry: excellent robustness to noise, simple and direct phase retrieval algorithm and acquires a 1D data set. However, this technique suffers from the disadvantage of requiring a high resolution PE spectrometer, which is an extremely demanding task. In addition, it is necessary to synchronise the timing of the IR and XUV fields with interferometric stability. However, it is possible to use the same IR field that generated the harmonics for the spectral shearing routine, thus ensuring that they are phase locked. Finally, the range of pulse durations is limited by the requirement that the dressing laser electric field must vary linearly with time. However, it is possible to use a longer wavelength dressing field if necessary (or utilise a different approach below). The biggest drawback is the need to generate two identical copies of the XUV field.

Attosecond streak camera

If the duration of the XUV pulse is comparable to the dressing laser period, then the spectral shift applied to the PE velocity distribution will become time dependent. If the XUV pulse is temporally aligned with the peak of the electric field, then the phase modulation will be approximately quadratic in time. This will result in a linear stretch (or streaking) of the PE spectrum. Comparison of the spectral width with and without

the streaking allows one to estimate the XUV pulse duration. This method, known as the attosecond streak camera^[205], has the advantage of being experimentally simple to implement. However, it is necessary to ensure synchronisation of the timing of the XUV and dressing laser fields. Thus the fundamental field used to generate the harmonics is often used as the dressing field.

XUV simplified chronocyclic tomography

XUV simplified chronocyclic tomography (SCT)^[21] can be considered as an extension of the atomic streak camera that allows one to unambiguously characterise the XUV field using a direct routine. The method is exactly as described in section 2.5.4. Two independent PE spectra are obtained for when the XUV field is aligned with the peak and the trough of the electric field, which corresponds to positive and negative quadratic temporal phase. When combined with the PE spectrum with no modulation, it is possible to recover the XUV spectral phase. It is necessary to ensure that the phase applied is equal in magnitude, of opposite sign and quadratic in nature for the modulated spectra. This requires the dressing field to contain many cycles to ensure that the envelope of the pulse does not affect the modulation and sets a limit on the maximum pulse duration that can be measured.

RABBITT

If the XUV pulse is much longer than the dressing field period, then sidebands will occur in the PE spectrum. If the dressing field is the same that is used to generate the harmonics, then the sideband will consist of a contribution from each adjacent harmonic, resulting in interference. If the sideband intensity is measured as a function of the delay between the XUV and dressing laser fields, sidebands will appear modulated with a period equal to twice the phase of the dressing laser field and an absolute offset equal to the phase

difference between the adjacent harmonics, i.e.

$$I_q \sim A \cos(2\phi_1 + \phi_{q-1} - \phi_{q+1}), \quad (4.44)$$

where I_q is the intensity at the even harmonic (i.e. sideband) q , ϕ_1 is the phase of the dressing laser, $\phi_{q\pm 1}$ is the phase of the XUV harmonics. If one obtains the phase offset for many sidebands, then it is possible to obtain the phase difference between multiple harmonics, and thus determine whether they are phase locked. This technique was the first experimental verification of attosecond bursts and is called resolution of attosecond beating by interference of two-photon transitions (RABBITT).^[206] Note that this technique is not valid for isolated attosecond pulses as it only gives the phase difference between adjacent harmonics. This is enough information to obtain the duration of individual pulses in a pulse train, but not the intensity envelope of the pulse train.

X-SPIDER

Another form of XUV-SPIDER that is similar to the RABBITT technique above can be used for XUV pulses that are long compared to the period of the dressing laser field.^[21] In this case, two time delayed XUV replicas are used to ionise a gas in the presence of a chirped dressing field. The dressing field will cause sidebands to form in the PE spectra of each individual XUV pulse. The frequency of the sideband is determined by the frequency of the dressing field. This method can be considered as a direct implementation of conventional SPIDER for PES, except that the SPIDER signal is encoded in sidebands of a PE spectrum. This method has the same advantages and disadvantages as the attosecond spectral shearing technique above, except that it works in a slightly different parameter range.

FROG-CRAB

The above techniques all rely on a specific relationship between the duration of the XUV field relative to the period of the dressing laser field. However, it is possible to completely characterise the XUV field without knowledge of the relative durations via frequency-resolved optical gating for complete reconstruction of attosecond bursts (FROG-CRAB).^[207] The method is completely analogous to conventional FROG in the optical domain. In conventional FROG, one uses an amplitude filter to gate the probe pulse. However, it is perfectly acceptable to use a phase only gate. Thus equation (4.42) is equivalent to a spectrogram where $P(t) = \mathbf{d}_p \cdot \boldsymbol{\xi}_{\text{xuv}}(t - \tau)$ and $G(t) = e^{i\Phi(t)}$. Using this method, it is not only possible to obtain the XUV field, but also the dressing laser field via the use of the PCGPA deconvolution algorithm. The limits of validity of this technique are set by the size of the phase modulation that can be applied, the resolution of the PE spectrometer and the range of travel for the stage used to scan the delay between the gate (i.e. dressing laser) and probe (i.e. XUV) fields. At present, this is the only technique that has been able to completely characterise isolated attosecond pulses.

4.4 Summary

High harmonic generation (HHG) can occur when an intense low-frequency field interacts with a gas, such that $\omega \ll I_p$, where ω is the laser frequency and I_p is the ionisation potential. The process is fairly well described by a simple three step model. If the intensity of the pulse is comparable to the ionisation potential (e.g. $I_0 \sim 10^{13}\text{--}10^{15} \text{ Wcm}^{-2}$) then the laser potential will suppress the atomic potential so that tunnel or over-the-barrier ionisation can occur. The freed electron is then driven by the laser field, crossing the atom every half laser cycle. As the electron passes the atom, it may recombine, emitting a high-energy photon. The frequency spectrum of the emitted radiation consists of peaks located at the odd harmonics of the driving field. The maximum photon energy that can be released is governed by the well-known cut-off law, $\omega_{\text{cut-off}} = I_p + 3.2U_p$, where U_p

is the pondermotive energy for the electron oscillating in the laser field and is given by $U_p = E_0^2 / (4\omega^2)$. This cut-off law can be found by solving for the maximum energy of the electrons that may return to the ion using simple classical mechanics. The fine details of the harmonic emission can be found by considering the quantum mechanical picture and solving Schrödinger's equation. However, this only gives the emission for a single atom. The total harmonic emission is governed by the coherent sum of all the atomic responses in the medium, which can be calculated by the propagation of the low-frequency and harmonic fields, using the single-atom response for the polarisation of the medium.

The ionisation process acts as a temporal gate for the laser field, due to the high nonlinearity of the ionisation process. Hence, ionisation is localised around the peaks of the electric field (e.g. the electron bursts from an 800 nm laser will be confined to several hundred attoseconds). It is this temporal gating that may be utilised to reach attosecond timescales. If a few-cycle pulse is used, then it is possible to confine the emission of electrons to a single burst, and thus produce a single isolated attosecond burst. However, this requires driving pulses with a duration of several femtoseconds and with the carrier envelope offset (CEO) stabilised so that there is a single electric field maximum.

Current schemes to temporally characterise attosecond pulses generated via HHG have been based around photoelectron spectroscopy (PES). In general, the XUV field is used to ionise electrons from an atomic gas in the presence of a low frequency dressing field (usually the same field used for the HHG process). The dressing laser field then acts as a temporal phase modulator. Depending on the relative durations of the XUV and dressing laser fields, it is possible to perform several different measurement techniques. However, one may consider the dressing laser field as a temporal gate function to the XUV field. As a result, the emitted photoelectron (PE) spectrum, measured as a function of the delay between the XUV and dressing laser fields can be considered as a spectrogram of the XUV field. Using the PCGPA, it is possible to deconvolve the XUV and dressing fields from the spectrogram and thus completely characterise the XUV fields. This method has been used to measure pulses as short as 130 as.^[176]

Although PES has proven a viable method of characterising attosecond pulses, it is not without its disadvantages. Typically the SNR is very low, requiring integration for long periods of time. This puts stringent demands on the CEO stability of the laser for characterising isolated attosecond pulses. In addition, the spectral resolution is limited, the devices are expensive and require expertise to use. It is therefore unlikely that PES methods will be accessible to the wider community. Spectral interferometry (SI), on the other hand, is generally much less expensive and requires less expertise. In addition, it is possible to obtain a SNR that is high enough to obtain single-shot data and thus reducing the stability constraints for the driving laser. However, complete characterisation of attosecond pulses has not been demonstrated. This is because complete characterisation requires a time-nonstationary (TNS) filter. At present, the only TNS filters available for the XUV involve photoionisation. However, it may be possible to utilise the time-nonstationary character of the HHG process itself, by controlling the generation in a controlled manner. In the following chapters, I describe the steps that I have taken to enable complete characterisation of attosecond pulses from HHG via spectral interferometry.

Spectral interferometry for characterising space-time coupling of XUV attosecond pulses from high harmonic generation

Currently, most XUV attosecond pulse characterisation methods for high harmonic generation (HHG) have neglected any space-time coupling (STC) of the pulses, either by integrating over the spatial profile or by spatial filtering. The generalised phase-matching (GPM) model (section 4.2.3) and simulations of the HHG process including full propagation effects^[188,208] have shown that the XUV pulses produced via HHG exhibit significant STC. The physical origin of this is that each frequency will have a different dipole phase due to the different times the electron has spent in the continuum, and thus will experience different phase-matching conditions. In addition, different generation regions will experience different laser intensities, and thus have a different dipole phase, again affecting the phase-matching.

It has been shown that by changing the focusing geometry and/or spatial filtering in the far-field, it is possible to select specific quantum trajectories, thus utilising the STC of the harmonic emission.^[199] It has also been shown theoretically that the shortest

pulse duration is obtained by spectral and spatial filtering. However, such a process will inevitably reduce the intensity of the XUV further, possibly to the extent of rendering such pulses useless for further experiments.

Having a means to measure the STC of the XUV field easily may allow one to optimise the pulse to minimise the amount of spatial and/or spectral filtering required. In fact, optimisation of high harmonic yield is still an active area of research; having such measurement capabilities would prove extremely useful in such research. Another active area of research lies in studying the generation process itself and extending the analysis to more complex generating systems than single atoms (e.g. molecules or clusters). Having complete characterisation of the emitted field (i.e. spectral *phase* and intensity) will not only improve the understanding of the generation process, but may yield further insight into the dynamics of the generating media beyond what can be gained from spectral intensity or photoelectron spectrometry alone.

Further information into the dynamics during the HHG process can only be gained by measurement of the full spatio-spectral intensity and phase, enabling a complete characterisation of the full electric field, $E(x, y, z, t)$. Obviously such a task will prove extremely difficult — very few measurement methods exist which provide such information in the optical regime. However, it should be possible to extend current measurement techniques into the XUV regime, specifically for measurement of emission from HHG such that it is possible to obtain the spatio-temporal field in one spatial dimension, i.e. $E(x, z_m, t)$ where z_m marks the position of the x - y plane where the measurement took place. If it is possible to make some assumptions on the field distribution in the x - y plane, such as cylindrical symmetry, then this measurement is sufficient to numerically propagate the pulse to an arbitrary z -position, as described in section 2.2.

Frequency-domain interferometry^[209,210] (FDI) and spectral interferometry (SI) are extremely robust and proven methods for measuring spectral phase and are easily extended to measuring the spatio-spectral phase in one spatial dimension by spatially resolving the

spectral interferogram. However, SI requires the availability of a known reference pulse with a frequency spectrum spanning that of the pulse to be measured for each location in the beam. Typically this is derived by replicating the pulse before the system to be studied using a beamsplitter. Unfortunately, no conventional beamsplitter exists for XUV light (although one may use a grating in a double pass configuration to replicate a beamsplitter, but this results in some angular dispersion and thus pulse front tilt). Alternatively, one can spatially split the beam in two, provided it does not exhibit any STC, which is known not to be the case for high harmonics. Thus the reference beam must be derived from a second source of XUV generated with the same conditions.

The general approach to performing XUV SI is to generate two sources of XUV, one of which remains constant to act as a local oscillator (i.e. reference). The second source then acts as the control, whereby the driving field is changed and the interference pattern is measured as a function of this change. Using the local oscillator, it is relatively straightforward to obtain the change in the XUV phase as a function of the control parameter: for any given measurement, the phase difference between the two sources is given by equation (5.1). If one source is kept constant, then the difference of the extracted phase between two measurements, $\tilde{\theta}_1$ and $\tilde{\theta}_2$ will yield the variation in the XUV phase as a function of the control parameter:

$$\begin{aligned}\tilde{\theta}_n(\omega, x) &= \arg \left\{ \mathfrak{F}_2^{-1} \left[f(\tilde{t}, \tilde{k}) \times \mathfrak{F}_2 \left[\tilde{I}_n(\omega, x) \right] \right] \right\} \\ &= \tilde{\phi}_n^A(\omega, x) - \tilde{\phi}_n^B(\omega, x),\end{aligned}\tag{5.1}$$

$$\tilde{\phi}_1^A(\omega, x) - \tilde{\phi}_2^A(\omega, x) = \tilde{\theta}_1(\omega, x) - \tilde{\theta}_2(\omega, x),\tag{5.2}$$

where $f(\tilde{t}, \tilde{k})$ is the filter used in the Fourier domain to extract the AC sideband; $\tilde{\theta}$ is the extracted phase difference between the two interfering beams; $\tilde{\phi}$ is the phase of one individual beam; the subscript numerals denote different measurements and the superscript letters denote the source arm and $\tilde{\phi}_2^A(\omega, x) \neq \tilde{\phi}_1^A(\omega, x)$ and $\tilde{\phi}_2^B(\omega, x) = \tilde{\phi}_1^B(\omega, x)$. Thus by varying the properties of one of the beams, either by changing the drive pulse or by

placing a sample in the path of one beam only, it is possible to measure the variation in the phase as a function of the variable control parameter. So far, reference based spectral interferometry like this has been used to measure the effect of different generating conditions^[211,212], the dispersion of materials^[213], the spatial and temporal coherence of the XUV^[214–217], to perform Fourier transform spectral interferometry^[218] and to show the different divergence of the long and short trajectories^[198].

The problem with performing such a method is that it does not provide the absolute phase of the XUV field — only the phase as a function of the control parameter. This is because there is no relationship between the two sources of XUV. In order to obtain the absolute phase via SI, not only is it necessary to be able to control the generation in a known manner, but it must be possible to relate the fields generated from the two sources to each other. The first goal in such a task is to be able to generate two *identical* sources, which until now has proven challenging and has not been demonstrated to my knowledge. This ability will be crucial, for example, in performing XUV SPIDER via direct detection of the XUV photons, and the steps which I have taken towards this goal will be outlined below.

The generation of two identical sources of XUV radiation via HHG relies on the ability to control and generate two identical driving pulses. The parameters which need to be controlled are: (1) temporal shape (i.e. pulse spectrum and phase, including the CEP), (2) beam wavefront (i.e. the beam profile and spatial phase), (3) pulse energy and (4) generating medium. Control of these parameters is easier for longer pulses with smaller bandwidths than for few-cycle pulses with nearly an octave bandwidth. This is because dispersion becomes less of an issue, allowing one to use thicker optics to prevent wavefront aberrations and polariser/waveplate combinations to control the power and it is possible to precisely control the spectral phase using pulse shapers. Additionally, many cycle pulses are generated directly out of a CPA system and thus do not suffer power loss due to pulse compression, providing more scope for controlling the pulse parameters with this excess

energy (e.g. controlling the power via a polariser/waveplate combination, pulse shaping and wavefront correction with adaptive mirrors).

The story for controlling few-cycle pulses, which are necessary for generation of isolated attosecond pulses, is much more problematic and poses many technical challenges. Firstly, it is necessary to use as few and thin transmissive optics as possible, in order to minimise the dispersion. However, thin optics can lead to wavefront aberrations. Secondly, it is difficult to obtain optics with sufficiently large bandwidth for these pulses. As an example, it may not be possible to use the polariser/waveplate arrangement often used for longer pulses. Over these bandwidths, pulse shapers have limited controllability and tend to be extremely inefficient. Thirdly, the generation of these pulses tends to come from compression of the output of a CPA, which results in a reduction in pulse energy to the level where generation of isolated attosecond pulses is not possible if one wants to split these pulses into two (e.g. 0.5-1 mJ per pulse before replication). Alternative generation methods (e.g. OPCPA) have recently become an alternative route which may solve some of these problems, but are currently not as commonplace in attosecond laboratories. Lastly, these pulses are more likely to experience space-time coupling beyond the user's control (e.g. nonlinear spatial chirp and spatially dependent chirp due to nonlinear effects), leading to further complications in the HHG process.

In this chapter, I outline the procedures and measurements which we have taken to ensure that the two sources of XUV radiation are identical. I then detail an experiment of spectrally resolved lateral shearing interferometry in the XUV which can be used to extract the complete spectrally resolved spatial phase of the pulse. Using this phase, it is possible to measure STC of the pulses generated and thus confirm the generalised phase-matching model proposed by Balcou.

5.1 Experimental setup

To generate the two driving pulses, we employed a compact folded Mach–Zehnder interferometer (figure 5.1). We used two identical fused silica beamsplitters (from Layertec) with 50:50 transmission:reflection (T:R) ratio with a thickness of 1 ± 0.1 mm. The beamsplitters are thin enough to minimise the dispersion, but thick enough to prevent wavefront aberrations during mounting. The beamsplitters are arranged such that each arm of the interferometer has one transmission only. This then ensures equal and minimal dispersion in each arm. The small amount of dispersion introduced by the interferometer can be compensated via chirped mirrors.

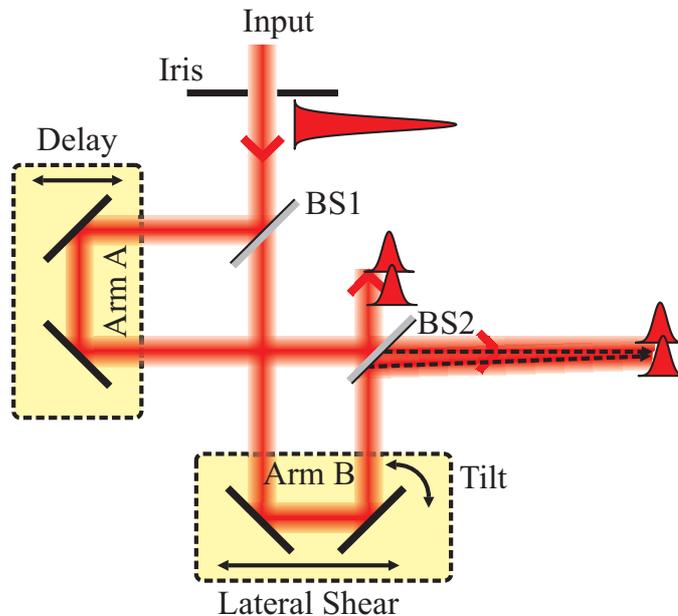


Figure 5.1: Mach–Zehnder interferometer used to generate two identical few-cycle pulses. BS1 & BS2: Beamsplitter (50:50). The two arms of the interferometer are labelled arm A and arm B. A delay is set using arm A, whilst the tilt and lateral shear is introduced in arm B.

As HHG is a nonlinear process, the two driving pulses need to be completely separated in either space or time. For performing spectrally resolved lateral shearing interferometry, the ideal scenario would be to have the two pulses co-propagating (with a small lateral displacement) and a delay between them. This enables sampling the spatial domain at the sampling limit and simultaneously allowing filtering of the AC sideband in the Fourier

domain due to the delay, regardless of the complexity of the spatial profile — compare this to SEA-SPIDER which introduces fringes in the spatial domain to enable extraction of the spectral phase gradient produced with a spectral shear. However, there are two reasons why this cannot be employed in general for isolated attosecond pulses produced via HHG: (1) the spectral resolution of the spectrometer being used requires a delay which is of the same order as the pulse width (FWHM) resulting in a distortion of the peak field due to the leading (trailing) edge of the succeeding (preceding) pulse; and (2) the intensity required to generate the desired wavelength causes significant ionisation after the preceding pulse, resulting in different generating conditions for the second pulse, and hence the harmonics from the two sources will not be identical.

An alternative arrangement is to introduce both a lateral shear (i.e. spatial offset) and tilt between the two beams at zero delay. As the beams are focused into the gas target/jet after the interferometer and the harmonics are generated close to the focus, the lateral shear and tilt introduced by the interferometer result in a tilt and lateral shear respectively in the focal region. Since the fulcrum of the tilt is not one focal length away from the parabolic mirror, and the harmonics are not necessarily generated directly at the focus (enabling different phase-matching conditions), the tilt and lateral shear at the focus are coupled. It is therefore necessary to employ additional measurements and a specific alignment procedure, which will be explained later.

Introducing a tilt between the two pulses after the interferometer ensures that the two foci are completely spatially separated. Thus there is no interference between the two beams inside the gas jet/target. This has the advantage of generating two independent, yet coherent sources of XUV radiation. After propagation, the two sources will diffract and overlap, resulting in fringes in the spatial dimension, where the resolution is sufficient for our application. This alleviates the spectral resolution requirements, which is typically the limiting factor for such large bandwidths.

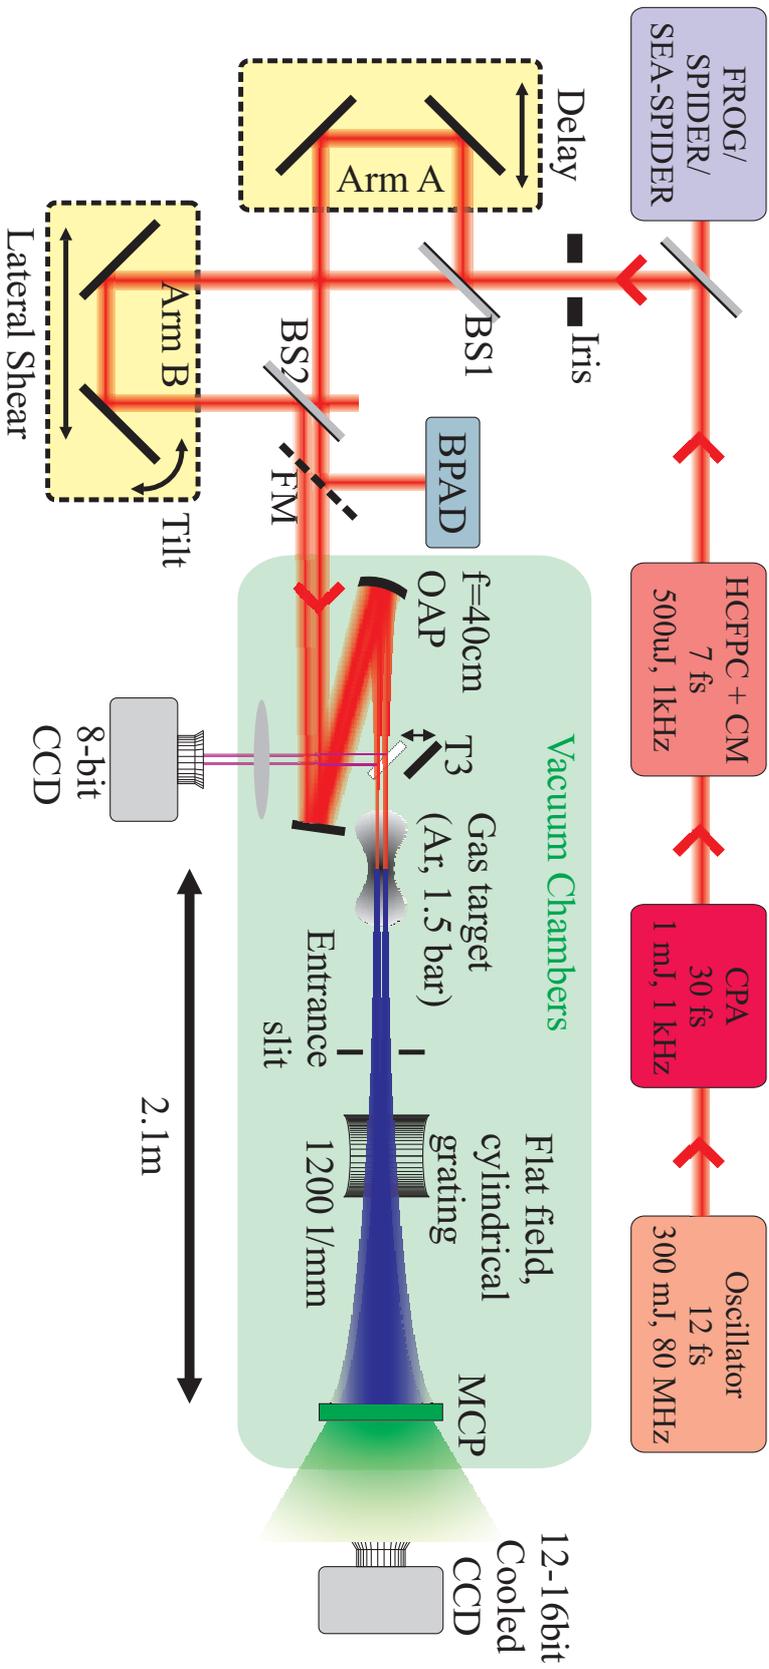


Figure 5.2: Experimental setup for performing XUV LSI via HHG. Few-cycle pulses are generated from a HCFPC system and characterised by a SEA-SPIDER, SPIDER or FROG device. After the interferometer, the pulses can be sent into the BPAD for alignment, measured with a spectrometer or the SEA-SPIDER, SPIDER or FROG using the flip mirror (FM). Alternatively they are sent into the vacuum chambers where they are focused via a 40 cm off-axis parabola (OAP) into a gas target or gas jet. The harmonics diffract to the flat field spectrometer where they are detected on a micro-channel plate (MCP) connected to a phosphor screen which is captured by a 12-bit CCD (which is also capable of 16-bit ‘fusion’ acquisition by integrating over multiple shots). Using T3, it is possible to translate a mirror into the beams to pick them off before the focus and image the foci onto an 8-bit CCD.

After the interferometer, the pulses are focused into either a solenoid gas jet (with an interaction length of $\sim 300 \mu\text{m}$) or a differentially pumped tube target (interaction length of $\sim 4 - 6 \text{ mm}$) by an 8° 40 cm off-axis parabolic mirror. The gas jet is synchronised with the laser at either 1 kHz or 2 Hz. The gas target allows a continuous flow at a constant pressure. The target was found to be more consistent than the gas jet, but has an extremely long interaction length and thus produces a more complicated and very different harmonic signal compared to the gas jet. After the gas, the XUV and NIR pulses are allowed to diffract into a spectrometer, with an image plane located 2.1 m from the source. The spectrometer employs a Hitachi non-uniformly ruled grazing incidence cylindrical grating, with a radius of curvature of 5.649 m giving an image plane at 237 mm from the grating,^[219] which images an entrance slit onto a micro-channel plate (MCP) attached to a fluorescent screen. The spectral resolution is in the range of $\Delta\lambda/\lambda \sim 500-1000$. Optimisation and a narrower entrance slit could increase the resolution up to a maximum of $\Delta\lambda/\lambda \sim 3000$, at the cost of photon flux. The reflectivity is approximately 20% for wavelengths in the range 10-20 nm. The curvature of the grating is perpendicular to the grooves, thus imaging the spectral components. The non-uniform rulings ensure a flat field image. The grating is flat in the spatial dimension of the image, allowing the beams to continue to diffract. The size of the grating is $30 \times 50 \text{ mm}$ (30 mm parallel to the rulings — i.e. across the ‘spatial’ dimension). A CCD is used to image the fluorescence from the screen (figure 5.2).

5.2 Alignment procedure

Due to the strong STC of the harmonics, alignment between the two sources is critical to ensure that the phase which is extracted can be related to the physical field. In addition, the fragility of few-cycle pulses results in conditional constraints on controlling their parameters. Thus it is crucial to use the specific alignment procedure coupled with the diagnostics which we have employed. It is this procedure and the diagnostics which allows self-referencing SI to be performed, which has not been possible in previous experiments.

This procedure is summarised as follows:

1. Interferometer setup:
 - (a) Ensure delay is adjusted without any lateral displacement.
 - (b) Ensure lateral shear results in a horizontal motion only.
 - (c) Ensure the tilt results in a horizontal motion only.

2. Interferometer alignment:
 - (a) Align two beams to be collinear using the beam pointing alignment device (BPAD) such that they overlap in both the near and far fields.
 - (b) Perform SI between the two arms to measure the dispersion imbalance of the interferometer.
 - (c) Characterise the individual pulses using e.g. SEA-SPIDER, SPIDER or FROG.
 - (d) Measure the wavefront aberration between each arm by introducing a tilt between the two beams and overlapping the spatial profiles on a CCD.
 - (e) Adjust the wavefront mismatch using the adaptable mirror.
 - (f) Once the wavefront mismatch is eliminated, introduce a small lateral shear between the two arms to measure the absolute wavefront via LSI.
 - (g) Measure the power of each arm. Adjust using reflective ND filters if necessary.
 - (h) Realign arms to be collinear.
 - (i) Optimise the HHG from one arm of the interferometer only.
 - (j) Ensure that the spatial/spectral intensity is identical from both arms. If not, repeat the alignment procedure.

3. Generated spatially separated sources:
 - (a) Adjust the tilt of arm B, observing the direction the XUV image moves.

- (b) Adjust the lateral shear such that the XUV image returns back to its original position. Continue to adjust so that the image from arm B passes to the other side of arm A.
- (c) Measure the spatial separation of the two foci.
- (d) Repeat until the two foci are completely separated and the separation of the XUV beams is as desired.

The first stage of the alignment is concerned with the initial setup of the interferometer. It is paramount that any adjustment of the delay, tilt or lateral displacement of the two beams with respect to each other does not result in any change in the beam pointing (other than that which is desired). This is checked by measuring the positions of the beams on a CCD for the two extremes of the travel of the translation stages and ensuring that the positions do not change in the case of the delay, or that the beams only move horizontally in the case of the tilt or lateral displacement. This is also checked regularly during the experiment and before any critical measurements to ensure that any drift in the laser pointing has not affected the interferometer alignment. The pinhole before the interferometer is closed to produce a small spot on the CCD and thus improve the accuracy in the position of the beams.

The second stage of the alignment procedure is to ensure that the two output beams from the interferometer are identical and thus produce identical harmonics. In order to distinguish between a difference in the pulse parameters from each arm of the interferometer and the alignment of the laser with respect to the gas jet/target and the XUV spectrometer, the two output arms are initially aligned to be collinear. This is performed with the aid of the beam pointing alignment device (BPAD) — figure 5.3 — immediately after the interferometer, which measures the near-field and far-field distribution of the pulses on the same CCD. The lens used in the device has the same focal length and thus mimics the parabolic mirror used to generate the harmonics. The narrow band-pass filter

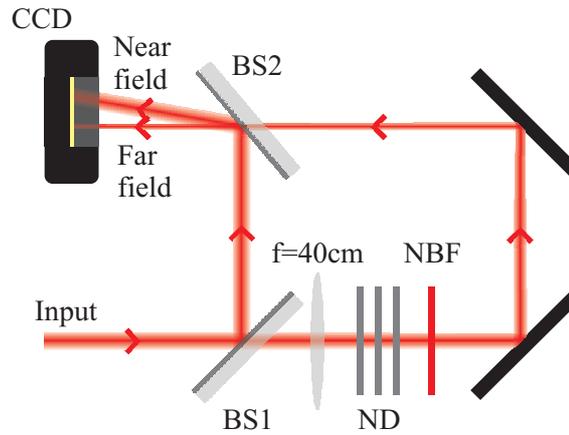


Figure 5.3: Beam pointing alignment device to help align the two arms of the interferometer with respect to each other. Measures both the near field and far field, and the approximate beam profile in the focus of the harmonics. BS1 & BS2: beamsplitter; ND: neutral density filters; NBF: narrowband filter.

is to prevent chromatic aberration due to the lens. The beams are aligned collinearly when the spots from each arm overlap with each other in both the far and near fields. Again, the pinhole before the interferometer is closed to reduce the spot size on the CCD for this measurement and thus improve the accuracy in their positions.

Once the two arms of the interferometer are aligned collinearly, they are sent into a spectrometer (Ocean Optics HR4000) with a delay between them in order to perform SI. The phase which is extracted from the interferogram using the Fourier filtering routine corresponds to the dispersion imbalance between the two arms of the interferometer. This dispersion mismatch is shown in figure 5.4 with the linear term, corresponding to the delay (~ 139 fs), and the CEP difference (measured to be 1 rad) removed. In these particular experiments, we were not sensitive to the CEP phase as we are measuring the plateau harmonics. However, this needs to be measured and set to zero via adjustment of the beamsplitters / using a piece of glass in the arms of the interferometer if CEP is critical. Note that it is not possible to measure the CEP (or spectral phase) of the laser, only the difference between the two arms via this measurement. Due to slight differences in the beamsplitter thickness and relative tilt between them, there will inevitably be some difference in dispersion between them. For example, a difference in thickness of 0.1 mm

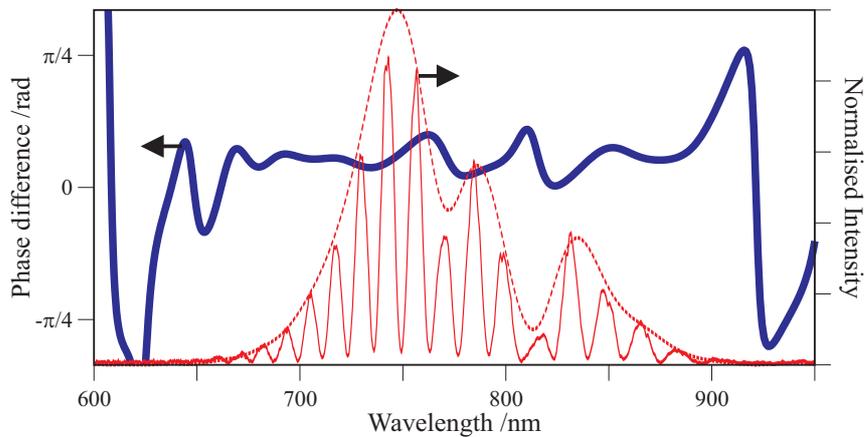


Figure 5.4: Phase difference between the two arms of the interferometer measured via SI. The measured interferogram (solid red) and the extracted amplitude (dashed red) is shown for reference as well as the extracted phase difference (solid blue) which is the dispersion imbalance between each arm.

will cause a broadening of ~ 2 fs over the FWHM, comparable to the pulse duration of 7-10 fs. If this is the case, it is possible to adjust the dispersion using a thin piece of glass in each arm and to change the angle relative to the beam in one arm to add additional dispersion. However, figure 5.4 clearly shows little dispersion between the two arms and so this was not necessary. The two arms are also sent, individually, into a SEA-SPIDER.

A significant finding from our experiments was that the harmonic emission was extremely sensitive to any wavefront aberration on the driving pulses, thus it was necessary to keep this minimal and, more importantly, the same for both arms. The difference in the wavefront aberrations can easily be verified in a similar manner to SI by introducing a tilt between the two arms and displace the two beams at the interferometer such that the two arms perfectly spatially overlap on a CCD (figure 5.5 (a)). The fringes introduced by the tilt allows one to extract the phase difference between the two arms via a Fourier filtering algorithm. If this is found to be negligible, then LSI can be performed by introducing a small displacement between the two beams at the CCD, or RDI can be performed by diverging / converging the beams with a lens before the interferometer (figure 5.5 (b)). As ultrabroadband pulses are being used, a narrowband filter can be used to increase the fringe contrast.

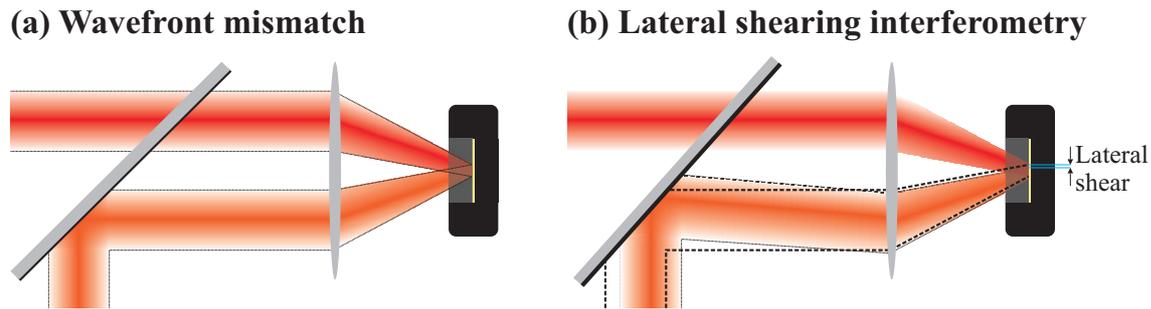


Figure 5.5: Measuring the wavefront aberrations of an interferometer. (a) Measuring the difference in the wavefront of each arm — the two beams perfectly spatially overlap on the CCD with a tilt between them. (b) Provided there is no wavefront aberration between the two arms, the complete spatial phase is measured via LSI by bringing the beams together with a tilt and small lateral shear.

It was found, due to using relatively thin optics, that careful mounting of the beamsplitters is crucial. Using beamsplitters manufactured in the same batch, or choosing a matching pair from a selection, offered the best chance of matching the wavefront aberrations. However, it was also found that there may be a slight difference in the radius of curvature, as little as $\lambda/10$ over the whole beam diameter, resulting in a change in the focal positions of the two sources on the order of 0.1-0.5 mm (figure 5.7), which was enough to change the harmonics produced. The source of this curvature could either be from using thin beamsplitters which are slightly bowed, or from using non-collimated beams such that the beamsplitters act as very weak lenses in transmission. In order to compensate for this, we used a specially built mount for one of the mirrors (figure 5.6) such that by pushing on the centre of the back of the mirror, a slightly curved surface is generated. This was found to be sufficient to adjust the focal position by up to a millimeter. We found that the gas target was more sensitive to the focal position compared to the gas jet; this is likely to be due to the extended interaction length of the gas target.

Finally, the pulse energy is checked using a power meter. As the shot-to-shot fluctuations of the pulse energy was on the order of several percent, a power meter is precise enough for this purpose. However, a photodiode can be used to measure this more precisely if required. To adjust the relative power of the two arms, such that they are equal, a variable reflective neutral density (ND) filter can be used in one of the arms. Alternatively,

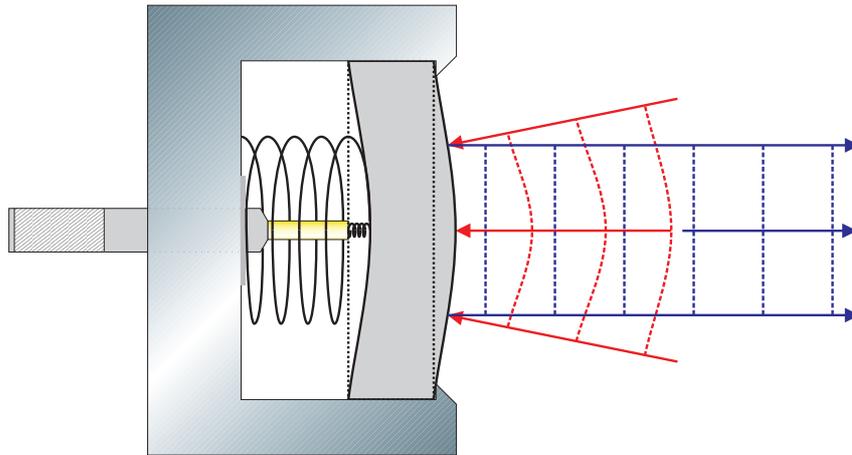


Figure 5.6: A simple adaptive mirror to compensate a spherical wavefront. An actuator is used to push against the back of the mirror via a stiff spring to generate a small spherical front. A soft spring is used to keep the mirror in place. The incoming, slightly spherical wavefront (red) is compensated so that the reflected beam is collimated (blue).

a thin transmissive ND filter can be inserted into each arm and the phase compensated with chirped mirrors. However, it was found that neither of these were necessary as the symmetry of the interferometer ensured equal power in both arms.

Once all of the parameters (dispersion, power and wavefront aberrations) have been checked and both arms found to be identical, within the limits of the measurements described above, the output of one arm of the interferometer — chosen to be arm A because the alignment of this arm is not adjusted — is aligned into the XUV chamber and the harmonics are recorded on the spatially resolved XUV spectrometer. After the alignment for HHG is optimised, the spatial-spectral intensity of the two individual arms are measured and compared to each other. Due to the high nonlinearity and significant STC of the HHG process, it is assumed that any change in the driving pulses will change the spatial-spectral intensity (i.e. it is unlikely that a change in the driving pulses will cause a phase only change in the emitted emission). This is confirmed by modifying the three properties of the driving pulses (dispersion, power and wavefront aberrations). As we find that the two spatial-spectral intensity distributions are remarkably similar (figure 5.8), we conclude that the two sources of XUV pulses are identical.

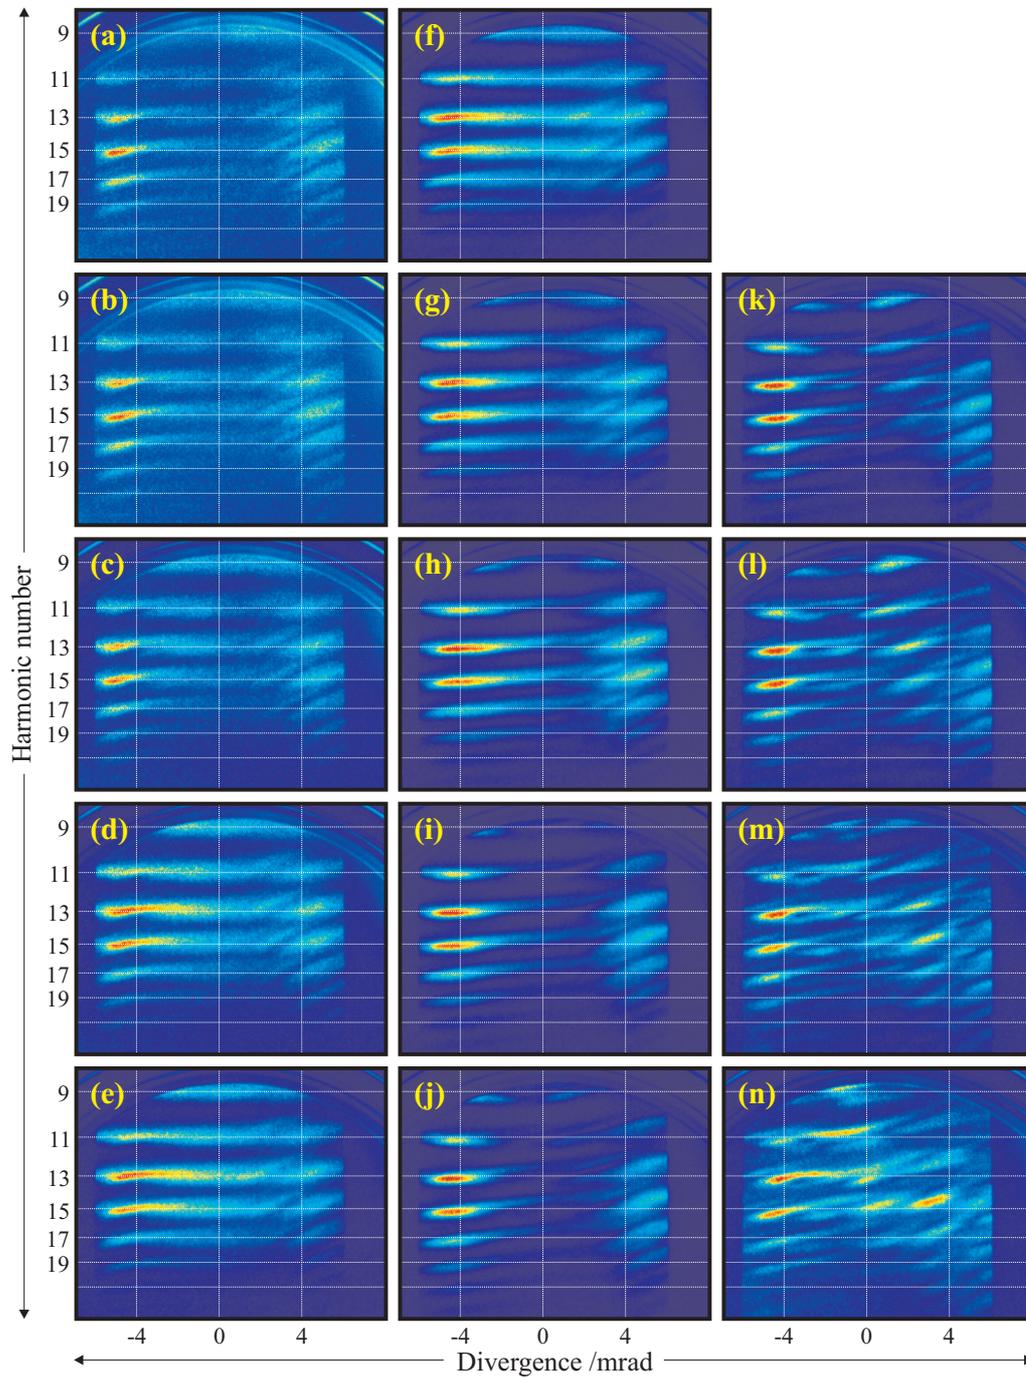


Figure 5.7: The dependence of the spatial-spectral intensity of high harmonics as a function of the focal position relative to the gas jet along the propagation direction. The relative positions are $z = z_0 + \dots$ (a) -2.5 mm, (b) -2.0 mm, (c) -1.5 mm, (d) -1.0 mm, (e) -0.5 mm, (f) -0.0 mm, (g) 0.5 mm, (h) 1.5 mm, (i) 2.0 mm, (j) 2.5 mm, (k) 3.0 mm, (l) 3.5 mm, (m) 4.0 mm (n) 4.5 mm, where z_0 is some unknown quantity near 0 mm. Fig. 5.8 corresponds to (g) in this figure, and is used as the reference for the position of the harmonics.

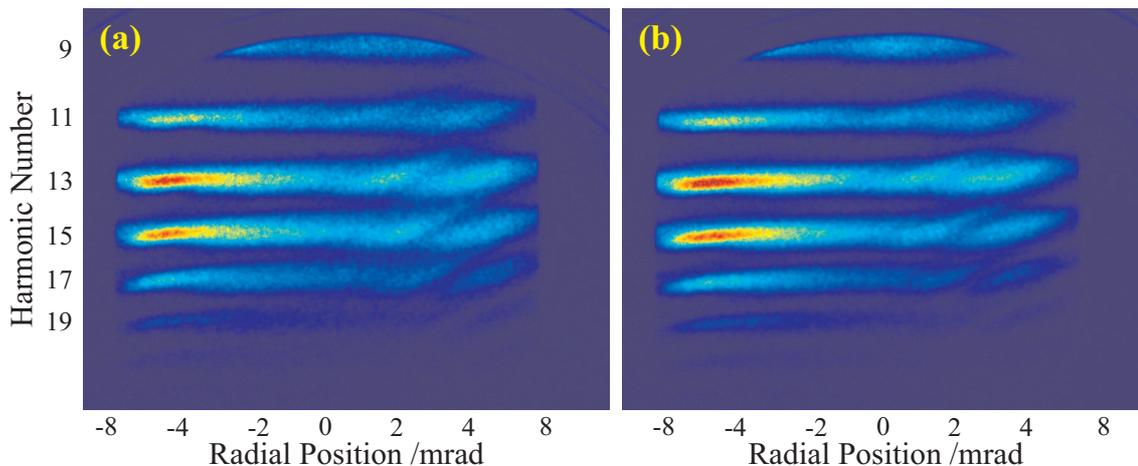


Figure 5.8: Spatial-spectral distribution of harmonics from each arm of the interferometer: (a) arm A and (b) arm B.

If the XUV spectrometer has sufficient spectral resolution, SI on the two collinear beams could be performed by introducing a delay between the two driving pulses and ensuring that the reconstructed phase is indeed flat over the whole distribution. Unfortunately our spectrometer only has a resolution which supported a delay of about 10 fs, which is too small to ensure that the two pulses are temporally separated.

The third and final stage of the alignment enables one to produce two spatially separated sources of the XUV radiation. The alignment procedure described below enables one to perform a spatially resolved lateral shearing interferometry (SRLSI) measurement. For this experiment, it is necessary to ensure that the two sources are spatially separated in the gas region, but overlap in the far-field as shown in figure 5.9. In addition, the tilt between the two beams may need to be known for the analysis.

Firstly, the tilt of arm B is adjusted whilst observing the direction in which the image of the XUV pulse moves. Then, the lateral shear is changed such that the XUV image moves back in the opposite direction and passes to the opposite side of arm A. This process is repeated until the two foci are completely spatially separated, and the lateral displacement of the XUV spatial distributions are as desired. This procedure is illustrated in figure 5.10.

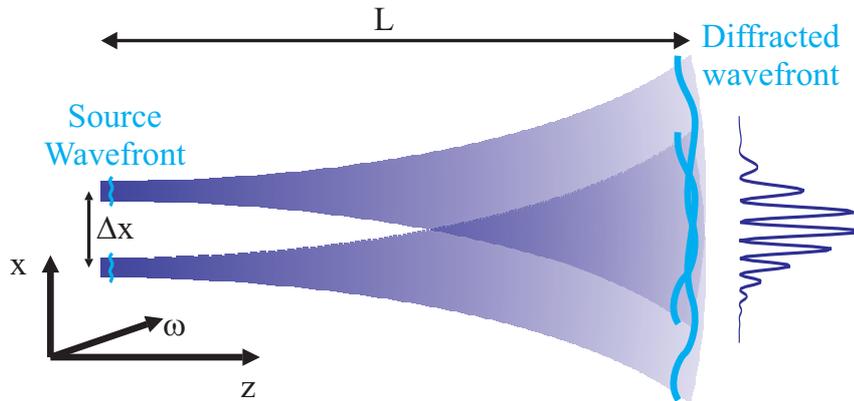


Figure 5.9: Propagation of XUV fields for performing lateral shearing interferometry.

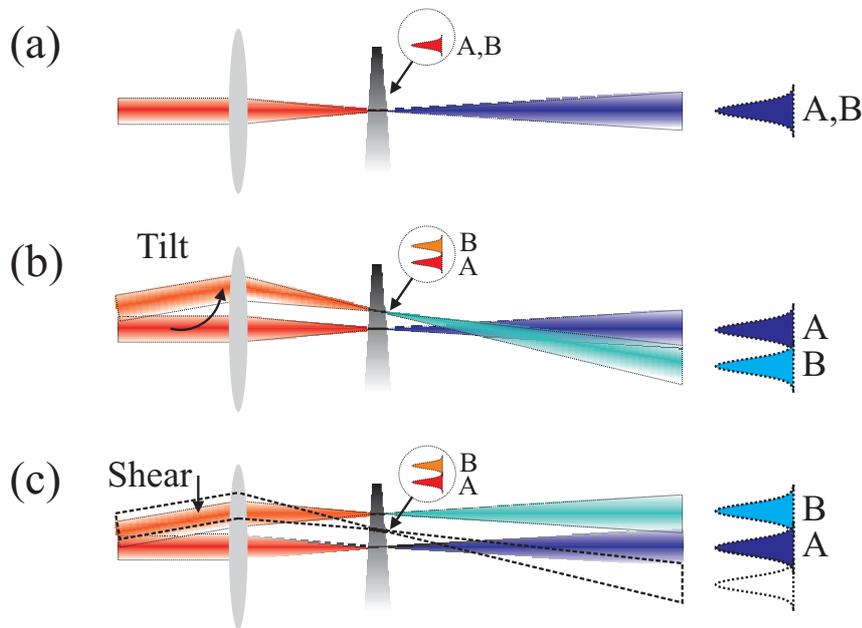


Figure 5.10: Walk-off procedure for aligning XUV beams. (a) Initially the beams are collinear. (b) The tilt is adjusted so that the foci are separated. (c) The lateral shear is adjusted so that the XUV beams propagate parallel with respect to each other. Iteration between (b) and (c) is required due to the coupling between the tilt and shear after the focusing parabola.

The separation of the foci is measured by inserting a pick-off mirror after the parabola, but before the foci, which are then imaged onto a CCD with a lens (using a narrow band-pass filter to reduce chromatic aberrations, and ND filters to prevent damage) with approximately $\times 2$ magnification. To ensure that the CCD is placed in the image plane, and to measure the exact magnification of the imaging system, we insert an iris with a

diameter slightly less than the focal spot size into the focal region. By optimising the power transmitted, we can ensure that the iris is at the centre of the focus. The CCD is then translated until a sharp image of the iris is formed. By measuring the diameter of the iris and comparing this to the image size, the magnification factor can be calculated.

Once the two output arms of the interferometer have been aligned such that the focal positions are spatially separated and the XUV beams are travelling parallel with each other, the relative positions of the two outputs in the near- and far-field can be measured using the BPAD, thus simplifying the alignment procedure in future. However, the positions of the foci and the XUV beams should always be measured independently for any experiments, as this information will be required for the analysis, as well as ensuring correct alignment. Care also needs to be taken when performing the walk-off procedure to ensure that the spatio-spectral intensity does not change. This is normally a sign that the pulse is either clipping in the case of the gas target or moved outside the central region of the gas jet. As the foci are typically separated by just over a beam waist (i.e. about $150\ \mu\text{m}$), the gas density should be equal for both beams.

5.3 XUV lateral shearing interferometry results

In general, the method for recovering the complex spatial field of the harmonics from the LSI method would be to perform the Fourier transform filtering routine, remove the phase resulting from a difference in the tilt between the two beams and then to concatenate the phase difference to obtain the spatial phase in the image plane. However, this phase would primarily consist of the phase due to diffraction (i.e. quadratic spatial phase) which does not itself contain much information. However, the pulse can then be numerically propagated to the source to enable examination of the finer details of the STC.

Although the above procedure is valid, it is possible to make some approximations to the spatial phase by using the fact that the image plane corresponds to the far-field of the XUV source provided the propagation distance, L is large. In the far field limit, the

spatial phase from one source at the image plane is given by

$$\Phi(\omega, x) \simeq -\frac{k}{2L}(x - x_0)^2 + \phi[\omega, x - (x_0 + L\theta)] + \tilde{\phi}_\omega(\omega). \quad (5.3)$$

The first term is purely due to diffraction in the far-field limit. x_0 is the spatial position of the source. The second term is due to any wavefront aberration from the source, propagated to the far field. Therefore the functional form of this term will be identical to the functional form of the momentum phase. L is the propagation distance between the source and image plane and θ is the angle of propagation. The final term is due to a spectrally dependent phase only, which is related to the temporal pulse profile. As will be shown below, this final term cannot be measured via LSI.

The extracted phase difference from two sources with a separation δx and relative tilt between them, $\delta\theta = \Delta x/L$ is given by

$$\tilde{\theta}_{\text{LSI}}(\omega, x) \simeq \tilde{\phi}[\omega, x - L\theta] - \tilde{\phi}[\omega, x - (L\theta + \Delta x)] + \frac{k\delta x}{L}(x - x_0) - \frac{k\delta x^2}{2L} \quad (5.4)$$

$$\simeq \Delta x \frac{\partial \tilde{\phi}(\omega, x - L\theta)}{\partial x} + K[(x - x_0) - \Phi_0], \quad (5.5)$$

where $K = k\delta x/L$ and $\Phi_0 = \delta x/2$. The first term corresponds to the gradient of the phase (scaled by the lateral shear in the image plane) in momentum space, where $k_x = k/L(x - L\theta)$ is the momentum co-ordinate. If the source exhibits angular dispersion, such that $\theta \rightarrow \theta(\omega)$, then the spatial phase in the far-field will appear to be spatially chirped. The latter term in (5.5) corresponds to an absolute phase which is dependent on frequency the lateral shear at the source and the position of the source. If the source exhibits spatial chirp, then this will appear as a frequency dependent absolute phase in the fringes. Simulations have shown that a resolution of $2 \mu\text{m}/(\text{rad/fs})$ is possible for the parameters used in this experiment.

If one wants to measure the divergence of the beam, i.e. $\phi(\omega, x) = \beta(\omega)x^2$, then it is necessary to perform an accurate calibration of the spectrometer, lateral shear (δx)

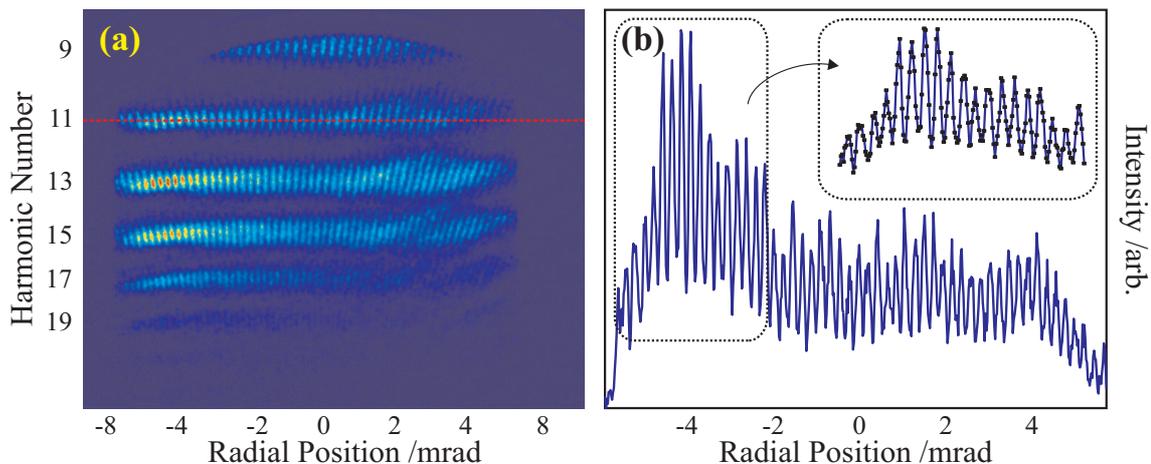


Figure 5.11: (a) Spatially and spectrally resolved interference pattern of XUV radiation generated from two sources of HHG. The individual spectra are shown in figure 5.8. (b) Lineout of the interferogram along the dashed red line in (a). (b: inset) Zoom of the interferogram showing the sampling frequency of about 10 points per fringe.

at the source and the propagation distance (L) so that the frequency dependent linear phase, $K(x - x_0)$ can be removed. This is because the divergence appears as a linear phase difference. This is analogous to removing the delay term in conventional SPIDER. In order to measure the spatial chirp in the image plane (i.e. angular dispersion at the source), it is necessary to remove the frequency dependent absolute phase, $K\Phi_0$. As it has not been possible to perform these calibrations accurately, I have not been able to perform a complete reconstruction of the spatial distributions of the harmonic field. However, it is still possible to perform a qualitative analysis and highlight the features which can be observed via XUV LSI.

The interference pattern corresponding to the simultaneous detection of both arms in 5.8 is shown in figure 5.11. It is clear that the fringe periodicity shown in figure 5.11 is a function of frequency, as described by equation (5.5). Figure 5.11 (b) shows a lineout of the fringes along the dotted line in (a). The inset shows the spatial sampling frequency, which corresponds to about 12 points per fringe at this harmonic. The Nyquist sampling limit requires a minimum of two points per fringe, which is sufficient to sample up to the 60th harmonic. Note that the current limitation in our setup is the imaging of the MCP

onto the CCD. Using a magnification system, this can be increased to allow sampling of higher frequencies. As the lateral shear is set by focal spot size, the fringe periodicity can also be reduced by moving the MCP further away from the source. However, this will also result in a smaller signal due to a larger beam size.

For the interference pattern shown in figure 5.11, I integrated the signal over 500 laser shots (0.5 s). The high fringe contrast signifies the similarity between the two sources and the stability of the system. Reducing the integration time can improve the contrast ratio, at the expense of signal intensity. However, we have also performed the experiment with single-shot acquisition using a gas jet.

As the fringe periodicity is a function of frequency, the AC sidebands become stretched in the Fourier domain, resulting in lower noise rejection. Thus, the phase on each harmonic is extracted independently by filtering before performing the Fourier filtering algorithm. Since the reconstruction algorithm is local, this process can be performed on a continuous spectrum as well.

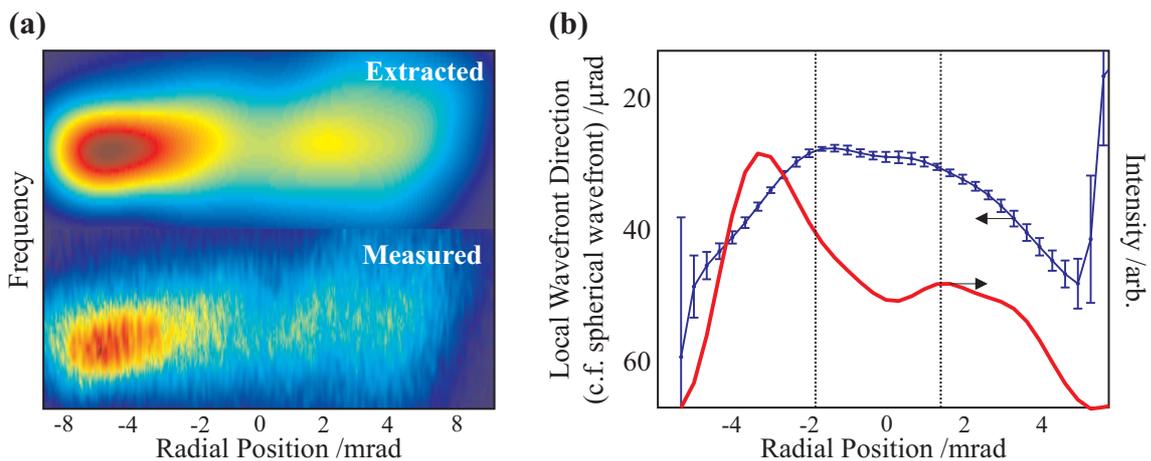


Figure 5.12: Extracted amplitude and phase from an XUV LSI measurement on harmonic 11 from argon in a gas target. (a) Intensity distribution extracted via the Fourier filtering routine and compared with the independently measured intensity space-frequency distribution. (b) Lineout of the extracted phase difference, corresponding to the local wave vector along the spatial co-ordinate (blue). A lineout of the extracted intensity is also shown (red). The dotted vertical lines highlight the boundaries of the piece-wise nature of the local wave vector.

For the following discussion, I consider the reconstruction of harmonic 11 only. Figure 5.12 (a) shows the extracted amplitude of the filtered sideband and the spatio-spectral intensity from one source only, showing excellent agreement. This comparison should always be performed to check that there are no artifacts from the filtering routine, and that the two sources are indeed the same.

The amplitude and phase difference (with the linear term corresponding to diffraction removed) is shown in figure 5.12 (b). As this extracted phase difference is approximately equal to the scaled gradient of the spatial phase, this represents the x-component of the local wave vector of the pulse. It is worth noting the apparent piece-wise nature of the wave vector, the boundaries are illustrated by the two vertical dotted lines. From the model described by Balcou^[196] and Chipperfield^[188], the different quantum orbits have different phase-matching conditions. This results in efficient generation in different regions of the gas with different wave vectors and divergence. In the far-field, the different quantum orbits should be spatially separated, and the phase of the orbits should also differ. This would suggest that the different regions shown correspond to different quantum orbits. However, this can only be confirmed by simulating the generation and propagation of the XUV for the exact pulse parameters. Unfortunately, not all of the pulse parameters have been measured, due to limited beamtime on the attosecond system, thus this is currently not possible. However, it is expected that these data can be taken soon and the simulations performed to confirm this theory.

An important feature which can be seen in figure 5.12 is that the intensity profile is asymmetric and the wave vector is symmetric (which requires an asymmetric phase) about the centre of the beam. However, theory would suggest that the beam should be cylindrically symmetric. The most probable cause for this asymmetry is either due to beam alignment or from an asymmetric driving pulse. One can indeed see that the image of the driving pulses are themselves asymmetric. This is another area which needs to be optimised before one can fully understand the system.

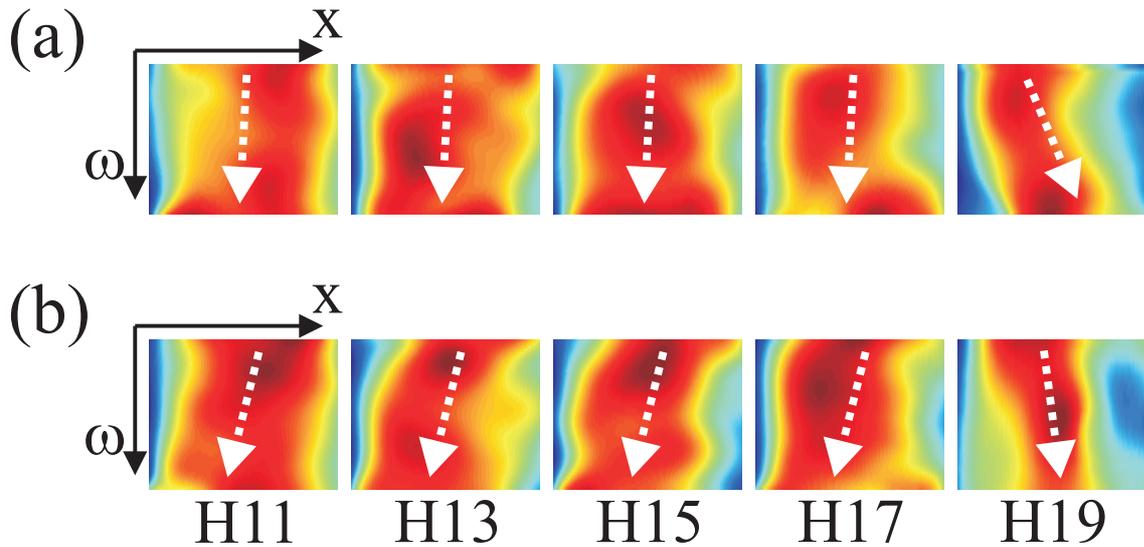


Figure 5.13: Extracted space-frequency phase from the central region of the beam. (a) Gas is positioned slightly after the focus. (b) Gas is positioned slightly before the focus. The STC among the plateau harmonics (H11–H17) are very similar. The cut-off harmonic (H19) shows quite different structure from the plateau harmonics. If the harmonics are produced before the focus, then spatial chirp is seen on the extracted phase (highlighted by the dotted white arrows), corresponding to angular dispersion at the source.

Figure 5.13 shows the space-frequency distribution of the extracted phase for the central portion of each of the harmonics, corresponding to a divergence value of less than 2 mrad and two different gas positions relative to the focus. The key feature of concern here is the contour which represent the boundaries between the different trajectories, and how the position at which the boundary occurs changes as a function of frequency. The boundary between the piece-wise sections of the local wave vector is clearly visible towards the left and right edges of each image, and is marked yellow on this colour map. The darker regions are likely to be artifacts of the reconstruction caused from noise in the data and can generally be ignored.

For the plateau harmonics, corresponding to H11–H17 in this experiment, the dependence of the position of the boundary, as a function of frequency, is approximately independent of the harmonic number. The cut-off harmonic (H19) however, has quite a different boundary structure compared to the plateau harmonics generated from the same pulse. This is expected due to the different generating regimes and thus different dependence of the atomic phase on intensity.

If the gas is placed before the focus (b), then the harmonics are angularly dispersed, resulting in PFT. This is clearly shown by the spatial chirp in the far field (i.e. the position at which the boundary occurs is approximately linearly dependent on frequency). The spatial chirp is indicated by the dotted white arrows. In comparison, the harmonics corresponding to generation after the focus, figure 5.13 (a), appear to have very little spatial chirp. This is as one would expect from the Balcou model, which suggests harmonics produced before the focus are generated off-axis with an angular divergence, and thus the non-collinear phase-matching can result in angular dispersion. Harmonics generated after the focus are produced on-axis as only collinear phase-matching is satisfied, thus there is no angular divergence. This angular dispersion, which gives rise to pulse-front tilt could result in a lower temporal resolution than one would expect from PES methods.

5.4 Summary

In order to utilise the HHG process to enable complete characterisation of attosecond pulses via spectral interferometry, it is necessary to generate two identical sources of the XUV radiation. This requires careful control of the driving pulse parameters, such as the pulse energy, dispersion and wavefront, as well as an identical generating medium. It is necessary to ensure that these parameters are identical, other than the control parameter used to perform the characterisation. In the generation of isolated attosecond pulses, it is necessary to use a CEO stabilised few-cycle laser pulse. Such pulses exhibit an octave spanning bandwidth. As a result, the availability of devices for controlling their parameters is much more limited compared to many-cycle pulses. In addition, the sensitivity to these parameters also increases. For example, only a small amount of dispersion is required to significantly stretch a few-cycle pulse. I have therefore designed an experimental configuration, system of alignment procedures and diagnostics that are required to ensure that it is possible to control the driving pulses with the accuracy required to perform spectral interferometry.

As a test of the ability to control the driving pulses, I have performed a set of lateral shear interferometry (LSI) experiments. LSI enables the spectrally-resolved spatial phase (i.e. wavefront) of an ultrashort to be measured via the interference of two spatially separated (i.e. laterally-sheared) pulse replicas in a spectrometer. By spectrally resolving the interferogram, it is possible to fully characterise the space-time coupling (STC) of the pulses. However, it is not possible to characterise the spectral phase, and thus perform complete temporal characterisation, of the pulse.

The advantage of performing LSI is that it is an entirely linear process. As a result, it is possible to vary the interferogram by simply changing the beam pointing. For example, it should be possible to change the lateral shear and tilt between the two beams, and ensure consistent reconstructed phase between the different measurements.

In addition, by measuring the STC of the harmonics, it should be possible to gain further insight into the dynamics of the HHG process. For example, through the coupling of the atomic phase to the phase-matching properties of the medium, it should be possible to use LSI as a simple means to accessing the atomic phase of the medium. Currently, I have only investigated simple atoms (e.g. argon) to test the method. However, it should be possible to extend the method to characterise more complicated species (e.g. molecules or clusters) which may yield new insights into their behaviour.

Finally, it should be possible to use LSI to optimise the generation process. For example, pulse-front tilt, which arises due to angular dispersion, results in a lower temporal resolution for time-resolved spectroscopy if the method integrates over the whole spatial profile of the beam. Due to the low SNR of PES methods, this is typically required. As it was not possible to fully characterise the parameters used in our LSI method, I have shown qualitatively that it is possible to extract particular STC features. For example, I showed how the angular dispersion varies as a function of the relative positions of the gas target and the centre of the focus. It is clear therefore, that SI may prove to be extremely useful in attosecond time-resolved spectroscopy and compliment PES techniques.

Chapter 6

Spectral phase interferometry for direct electric-field reconstruction of XUV attosecond pulses from high harmonic generation

To date, the only demonstrated method for complete characterisation of attosecond pulses has been via photoelectron spectroscopy (PES), whereby XUV photoionised electrons are streaked by an intense IR field; the dynamics of which is determined by the phase of the IR at the time of ionisation. This method, FROG-CRAB^[207], allows reconstruction of both the IR and XUV fields via an iterative deconvolution based on a principle component generalised projection algorithm (PCGPA). Using this method, it has been possible to characterise isolated pulses as short as 130 as.^[176] Lower order harmonics have also been measured via two-photon absorption^[220], but these do not give sub-fs pulses and it is unlikely that this technique can be extended to the higher harmonics necessary to obtain attosecond pulses.

Although FROG-CRAB has been proven to extract the temporal and spectral characteristics of attosecond pulses, it is not the most ideal method for general use. As it uses a photoelectron spectrometer, specific expertise is required in addition to a relatively large expense, thus not lending itself to the general scientific community. Secondly, the detection efficiency for XUV attosecond PES is extremely small, due to the low cross-section of XUV ionisation and the relatively weak intensity of the XUV field. Thirdly, the traces require scanning the delay between the two fields, with each delay step requiring many hundreds of laser shots to obtain sufficient signal for the reconstruction. Thus, they require extremely stable lasers, especially CEP stability over several hours, which has proven to be an extremely difficult task, limiting this capability to only a handful of groups worldwide. Fourthly, the non-instantaneous and non-uniform response of the photoionisation process requires the deconvolution of the spectrally dependent dipole matrix element from the acquired trace. Until now, this has only been performed for the intensity profile, and no phase dependence has been taken into account. This has been because the spectral bandwidth for these measurements has been deemed small enough for this effect to be negligible. However, as the pulse duration decreases and spectral bandwidth increases, this may need to be considered. Finally, the field produced via HHG is known to exhibit significant STC. In fact, it is possible to utilise this to select particular features of the generation process through phase-matching conditions. However, such STC may lead to inaccuracies in PE based methods as they require averaging over a significant portion of the field to obtain sufficient intensity for ionisation. This averaging can also reduce the temporal resolution of the experiment.

In order for attoscience to progress to the wider scientific community, it is necessary to simplify the measurement procedure, and to significantly reduce the acquisition and processing times involved. Not only that, but PES only provides half of the information of laser-matter interactions, in the sense that it records the reactive response (i.e. how the medium reacts to the incident field). Measuring the scattered field should provide the

active response (i.e. the dynamics of how the system relaxes back to equilibrium) and thus compliment PES methods to obtain the complete picture.

Although spectral interferometry (SI) has already proven itself as a possibility in the XUV and spectral phase interferometry for direct electric field reconstruction (SPIDER) is well established in the IR, visible and UV spectral regions, extending SPIDER to the XUV has not been so forthcoming.

Methods for performing SPIDER via photoionisation have been proposed, but it has proven experimentally challenging. The main difficulty for photoelectron (PE) SPIDERS is the high spectral resolution required to resolve the spectral fringes, which is extremely difficult to achieve in photoelectron spectrometers. As the PE sensitivity is very low, it is also necessary to achieve interferometric stability on a single attosecond timescale over the entire duration of the acquisition, probably a more stringent condition than for FROG-CRAB. Although it may be possible to perform PE-SPIDER in a spatially encoded arrangement by ‘imaging’ the electron momentum distribution via velocity map imaging^[221], and thus possibly reducing the spectral resolution constraints, this avenue has yet to be fully explored.

For spectral interferometry to be established in the XUV attosecond regime, it is necessary to extend its capability for direct detection of the photons, whereby spectral interferometry should excel over PES techniques. Unfortunately the difficulty in doing so is that currently there does not exist a time-nonstationary filter for XUV pulses, which is required to perform temporal characterisation. Current time-nonstationary filters operate on the photoelectrons, not the XUV field. Thus it is necessary to utilise the nonlinearity of the high harmonic generation process in a known way to produce two XUV pulses which can be interfered in an XUV spectrometer with sufficient information relating the spectral phases of each pulse to allow reconstruction of the spectral phase of an individual pulse.

In this chapter, I outline a method of performing SPIDER in the XUV region which is based on the direct detection of the photons via the control of the generation process. Via

numerical simulations, I demonstrate theoretically that the process is based on a sound principle. As this technique is based on spectrally/temporally shaping the driving field, I outline a few methods in which this can be achieved experimentally.

6.1 XUV SPIDER via HHG

One of the underlying features of HHG is that the spectrum consists of peaks centred around the odd harmonics of the driving field, corresponding to the emission and recollision of electrons every half cycle of the laser field in an antisymmetric fashion. Thus it is reasonable to assume that if one were to shift the central frequency of the driving laser by an amount $\delta\omega \ll \Delta\omega$, where $\Delta\omega$ is the bandwidth of the drive pulse, then the harmonics should shift in frequency by an amount $\Omega = n\delta\omega$ where n is the harmonic number. Thus it seems possible to shift the harmonic frequency by simply driving with a slightly frequency shifted replica of the IR field. This process is a direct consequence of conservation of energy.

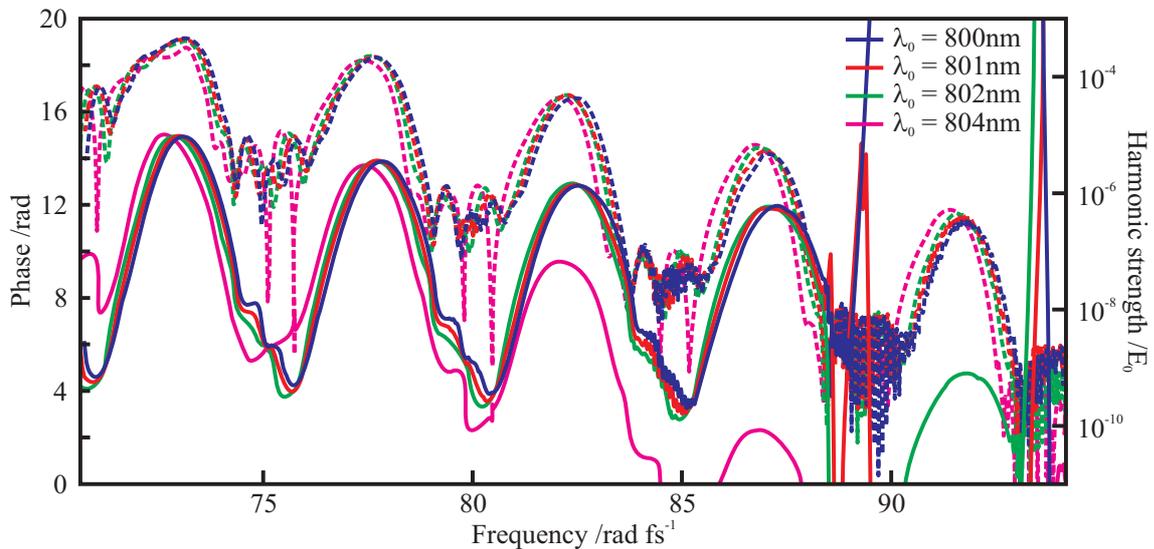


Figure 6.1: Simulation of the single atom response for different driving frequencies. The dashed lines correspond to the amplitudes and the solid lines to the phases for driving wavelengths of 800 nm (blue), 801 nm (red), 802 nm (green) and 804 nm (magenta). Data obtained by E. Cormier's TDSE simulations.

In order to demonstrate quantitatively that a small shift in the frequency of the driving pulse does indeed lead to sheared harmonics, we have performed numerical simulations to solve the 3D TDSE for the single atom response in cylindrical co-ordinates in argon ($I_p = 0.581$), further details can be found in these papers and references therein.^[22,23] These simulations were performed by E. Cormier and the data consisted of harmonic spectral phase and amplitude for different pulse parameters. I then used this data to simulate SPIDER interferograms, the SPIDER reconstruction and analyse the suitability of certain pulse parameters as discussed below.

For these simulations, the pulses were assumed to be FTL (i.e. no spectral phase) with a Gaussian envelope:

$$\boldsymbol{\xi}(t) = \sqrt{I_0} e^{-\frac{\alpha}{2} \left(\frac{t}{\Delta t}\right)^2} \cos[(\omega_0 + \delta\omega)t] \hat{\boldsymbol{x}}, \quad (6.1)$$

where $\alpha = 4 \ln 2$ if Δt is the FWHM intensity duration, and atomic units are used for the simulation. The FWHM duration was 19.94 fs with a peak intensity of $I_0 = 1.7 \times 10^{14} \text{ Wcm}^{-2}$ (giving the cut-off harmonic as $n_{\text{cut-off}} = 31$, i.e. $\lambda_{\text{min}} = 25.5 \text{ nm}$) and carrier wavelength of 800 nm. The phase of the harmonics for various shears, $\delta\omega$ is shown in figure 6.1. It is clear that the different harmonics are indeed sheared replicas of each other in terms of both the frequency and phase. However, the amplitude of the harmonics increases slightly as the driving frequency decreases. This is likely to result from the integration of the electric field to obtain the vector potential, resulting in a scaling of the vector potential by the carrier frequency. In addition, there are also other discrepancies, near 80 rad/fs for example. This is likely to be due to numerical errors in the simulation, as it is clear that the spectral intensity is falling to zero and thus possibly causing an unphysical phase shift. Whether this effect is truly physical or not, it is clear that the smaller shears still result in frequency shifted harmonics.

To observe that the harmonics are indeed sheared replicas, figure 6.2 shows the same data except plotted against harmonic number (i.e. $\omega / (\omega_0 + \delta\omega)$). Clearly the harmonic

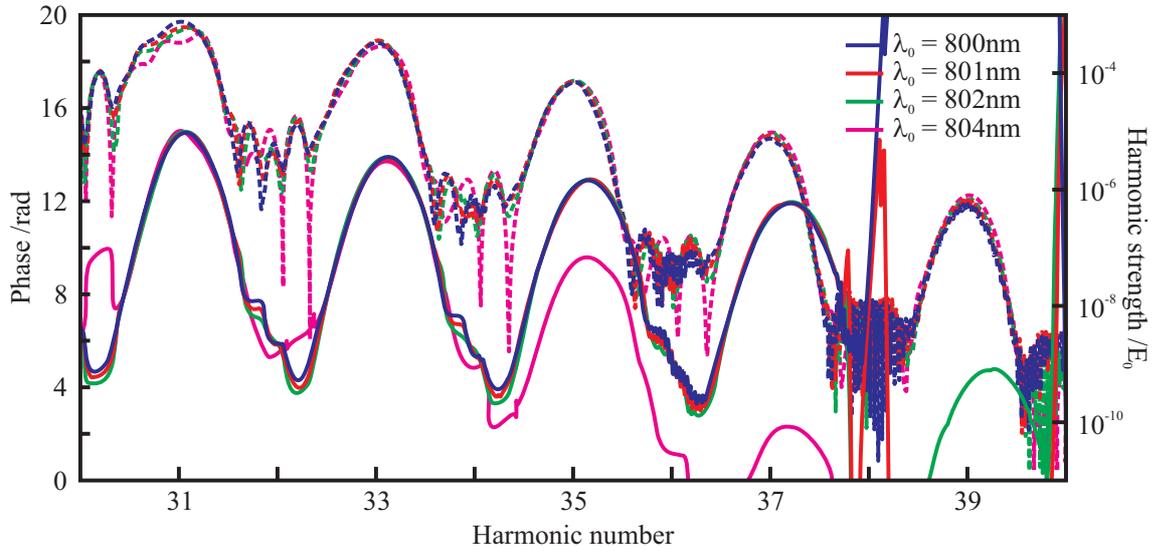


Figure 6.2: Overlap of harmonics driven with different central frequencies. This is the same plot as for figure 6.1 except the harmonics are plotted against harmonics number (i.e. the frequency axis is scaled by the inverse of the central frequency of the driving pulse). The harmonics clearly overlap in terms of harmonics number. Data obtained from E. Corimer’s simulations.

intensity and phase show excellent agreement, with slight discrepancies at lower field intensities which are probably due to numerical inaccuracies.

From figure 6.2, it is clear that the shear is dependent on the harmonic number, n , i.e. $\Omega = n\delta\omega$. Using this information, the shear used for the SPIDER reconstruction is no longer constant but can be calculated using

$$\begin{aligned}\Omega &= \omega - \omega' \\ &= \frac{\delta\omega}{\omega_0}\omega,\end{aligned}\tag{6.2}$$

where ω and ω' are the normal and sheared frequency axes respectively; $\delta\omega = \omega_0 - \omega'_0$ is the shear at the fundamental wavelength and ω_0 is the fundamental carrier frequency.

As the shear is small, it is only possible to measure the phase across any individual harmonic, and thus the time window which needs to be extracted is slightly less than that of the driving pulse duration, assuming no dispersion of the harmonics after generation. Thus for 30 fs pulses (FWHM), the spectral shear needs to be less than 2.3 nm in order to

measure the 31st harmonic. Therefore it is possible to use the desired shear to measure these harmonics.

Just as is in the optical domain, there are two different geometries which can be used to perform XUV-SPIDER: (1) collinearly with a time delay (as in conventional SPIDER) and (2) propagation from separate sources with a tilt between them (as in SEA-SPIDER). Each geometry has its set of advantages and disadvantages, and will be discussed below.

6.1.1 Spectrally encoded XUV-SPIDER

The geometry for performing spectrally encoded (i.e. conventional) XUV-SPIDER is shown in figure 6.3. In this case, two spectrally sheared replicas of the driving pulse are sent collinearly into the gas jet to produce two spectrally sheared replicas of the attosecond pulse/pulse train. The delay between the two IR pulses needs to be sufficient to avoid temporal overlap of the two pulses and to ensure that the spectral fringes have sufficient periodicity to enable extraction of the sidebands in the Fourier domain. However, the separation also needs to be small enough that it is possible to resolve the fringes with the spectrometer.

The advantage of this technique is that the two pulses will experience the same gas density and thus the same dispersion after their generation. Producing the two sheared driving pulses can also be technically simpler for this geometry, for example using a pulse shaper. However, it is necessary to note that the first pulse may cause ionisation which will affect the succeeding pulse, thus limiting the peak intensity and therefore the maximum

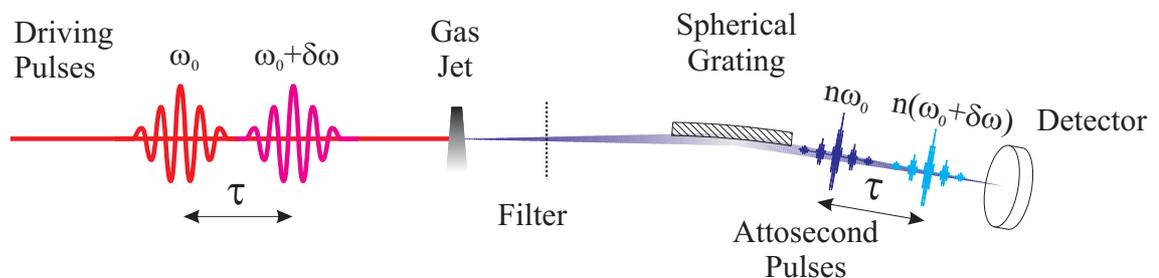


Figure 6.3: Experimental setup for performing spectrally encoded XUV-SPIDER. Note that it is also possible to spatially resolve the interferogram, at a cost of detection efficiency.

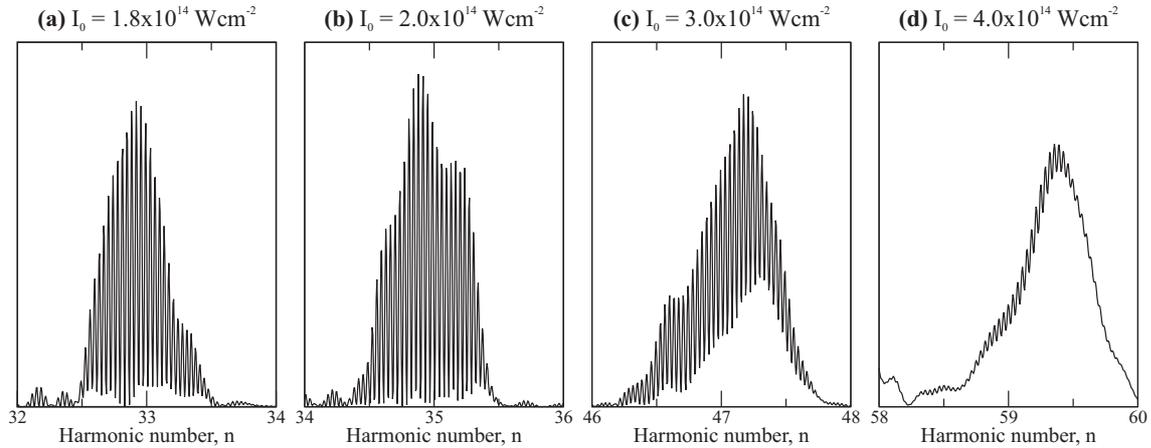


Figure 6.4: Simulation of spectrally encoded XUV-SPIDER for different driving intensities. As the intensity increases, the effects of ionisation result in a decreasing fringe contrast. The parameters used for the simulation are the same as described above. Provided $I_0 \lesssim 2.0 \times 10^{14} \text{ Wcm}^{-2}$ in Argon, the affect of the preceding pulse on the medium is remains negligible. The harmonics depicted correspond to the cut-off region, hence the harmonic number shown increases with intensity. These figures were obtain from E. Cormier, who performed full 3D TDSE simulations of a double-pulse train in Argon.

frequency. This is shown by the simulation results in figure 6.4. Thus this route is unlikely to be suitable for characterisation of isolated attosecond pulses, which are generated in the cut-off region by relatively intense few-cycle pulses. Another disadvantage is the requirement of a high spectral resolution, which is very difficult to achieve. For example, the Imperial College setup provides sufficient resolution to sample fringes due to a delay of 13 fs, which is insufficient to ensure two 7 fs pulses are temporally separated. This will also cause problems for resolving the shorter pulses which have a larger bandwidth at a lower wavelength.

6.1.2 Spatially encoded XUV-SPIDER

Figure 6.5 shows the setup required to perform spatially encoded XUV-SPIDER. In this case, the two sheared driving pulses are focused to two spatially separated regions in the gas jet. The two sheared XUV pulses then propagate such that they perfectly spatially overlap at the detector. As they started from two separate sources, there will be spatial fringes due to their tilt, which will allow the phase to be extracted just as for SEA-SPIDER

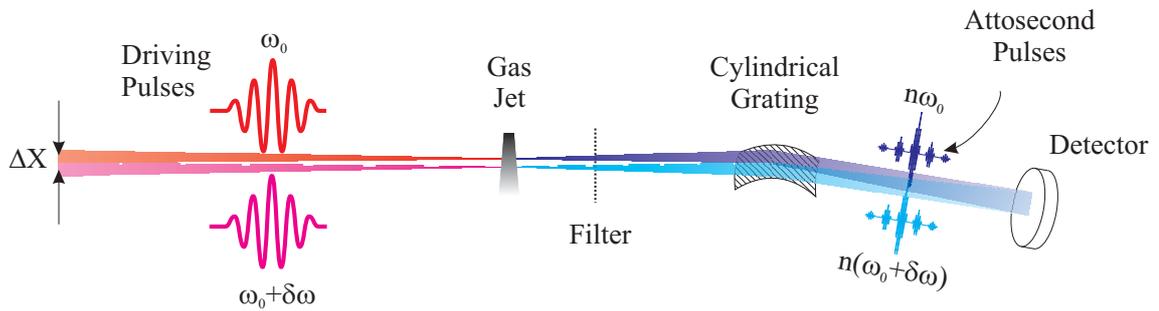


Figure 6.5: Experimental setup for performing spatially encoded XUV-SPIDER. Note that the two sources are spatially separated, but the XUV fields perfectly spatially overlap at the detector.

in the optical domain. It is necessary to note, as was found in the previous chapter, that it is important to ensure spatial overlap of the two XUV fields at the exit of the spectrometer, otherwise any STC will distort the results.

The two main advantages of this geometry are: (1) the spectrum can be sampled at the sampling limit and (2) there is no restriction on the intensity of the pulses used for driving the system. The downside is that as the two foci are in different locations in the gas jet, they may experience different densities and thus produce different harmonics. However, the previous chapter has shown that it is possible to generate identical harmonics this way, thus it should also be possible to generate two sheared replicas. A simulation of the SEA-SPIDER interferogram for the data in the previous section is shown in figure 6.6.

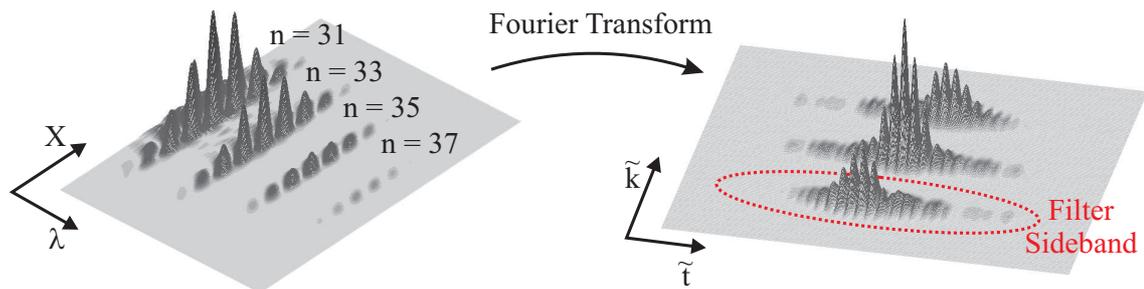


Figure 6.6: Simulation of spatially encoded XUV-SPIDER and the Fourier filtering process. A 2D filter in the Fourier domain allows the phase to be extracted in the same way as SEA-SPIDER in the optical domain. These interferograms were produced using data from E. Cormier via a TDSE simulations of the single atom response from two independent pulses with differing frequencies.

6.2 Shearing the driving fields

One of the biggest advantages of XUV-SPIDER is that the detection is performed using a standard XUV spectrometer, yielding the highest sensitivity and simplicity for spectral measurements. In addition, it is very straightforward to spatially resolve the spectrum, and thus obtain complete information of the STC of the pulse. However, the disadvantage is that it is necessary to control the driving fields with extremely high precision and accuracy. There are many different routes which can be used in order to shear the driving field, and I detail them below.

6.2.1 Nonlinear frequency mixing

One method which may be possible to shear the driving fields is through nonlinear frequency mixing in a similar manner for generating the sheared pulses in SPIDER. Such a method will be dependent on obtaining sufficient intensity at the desired wavelength, thus is likely to be suitable for much longer pulses where damage in the nonlinear medium is less likely to be a problem. Using this method, it may be possible to generate replicas with almost any desired shear. One issue may be that the resultant pulses occur at a different, possibly undesirable, frequency to the fundamental. An alternative method is to utilise a cascaded $\chi^{(2)}$ process which results in the same frequency as the fundamental, and by controlling the phasematching it may also be possible to control the central wavelength of the upconverted pulse.

6.2.2 Pulse shaper

Probably the most obvious and controllable method to shear the driving fields is to use a pulse shaper. The merits and capabilities of pulse shaping is dependent on the method used to generate the driving pulses and on the regime of their parameters such as duration/bandwidth.

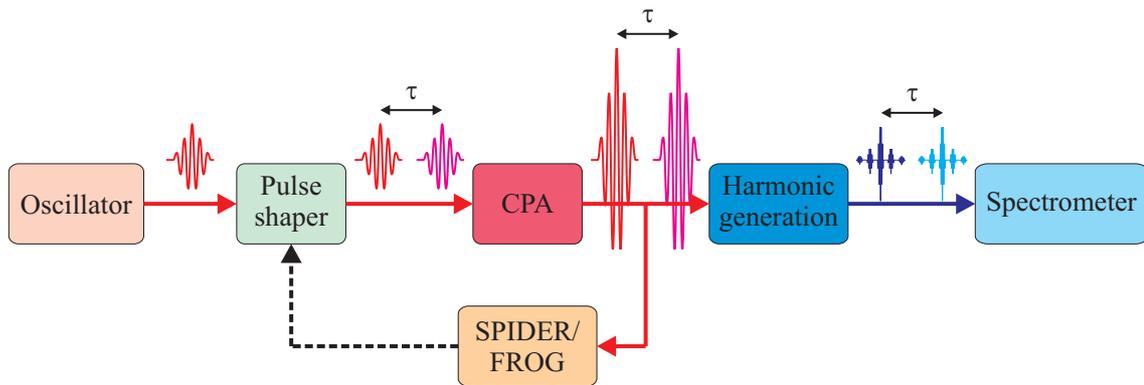


Figure 6.7: Implementation of spectrally encoded SPIDER using a pulse shaper between the oscillator and CPA lasers. A SPIDER or FROG can be used to measure the individual shaped pulses at the output of the CPA system to ensure that they are sheared replicas, using a feedback loop if necessary.

For many cycle pulses which are generated via a CPA system, one can place the pulse shaper between the oscillator and amplifier (figure 6.7). This method has already been demonstrated using an APODF with 50 fs pulses.^[222] It was shown that generating spectrally sheared XUV pulses was possible and that reconstruction of the chirp matched that as expected from the theory. Currently this is the only implementation of XUV SPIDER, but it has experimentally demonstrated that it is feasible. Placing the shaper before the amplifier allows one to produce two time delayed pulses with no significant reduction in their power. However, as the amplification is nonlinear, it is necessary to have a complete characterisation device after the CPA to ensure that the two pulses are indeed spectrally sheared copies.

If few-cycle pulses are generated via pulse compression (e.g. HCFPC or filamentation) after the CPA, this method may still be possible if one can use a closed-loop feedback system with a global search algorithm to ensure that the two pulses are indeed sheared replicas of each other. However, it is probably best to place the shaper after the compression stage, although this will inherently suffer from loss in power, and many pulse shapers operate with limited power and bandwidth.

In the case of few-cycle pulses generated via OPCPA systems, it may be possible to use a pulse shaper before the amplification stage and generate the shear in the same way as for many cycle pulses from a CPA.

It is necessary to note that shearing a pulse requires replicating the pulse spectral intensity and phase at a different frequency. This is not possible to do directly using a linear shaper. However, it is possible to simulate the process by applying amplitude masks to generate the desired shifted spectra, each of which are encompassed by the original spectrum before filtering. It is also necessary to measure the spectral phase and then apply the correct compensation phase to create the sheared replicas. Of course this also means that it may be possible to remove any spectral phase completely, and thus only amplitude shaping would be required.

6.2.3 Amplitude shaping

For near FTL pulses, such that the spectral phase is approximately flat, amplitude shaping only may be sufficient to shear the driving pulse. One simple method may be to clip the spectral phase in a 4-f line. This may work as the temporal envelope will not be affected by small changes in the spectral amplitude (provided it is near FTL). It may be necessary to perform the spectral clipping in an antisymmetric manner (i.e. clip the high frequencies of one pulse and the low frequencies of the other) such that the two pulses have the same bandwidth and thus same pulse duration. However, using a zero dispersion line may prove difficult for extremely large bandwidths.

For our experiments, we propose to use a novel, simple and universal approach to shifting the frequency of the driving laser of near FTL pulses. In this case, we propose to perform the amplitude shaping via a beamsplitter with a transmission/reflectivity which varies linearly as a function of frequency (appendix B) in place of BS2 in the interferometer in figure 5.1 from the preceding chapter. The alignment procedure is exactly the same as in the preceding chapter. However, it is necessary to ensure the two XUV pulses perfectly spatially overlap at the exit of the spectrometer (rather than maintaining the lateral shear). In addition, it is necessary to calibrate zero delay between the two XUV pulses. This cannot be done directly in the IR because the delay between the XUV pulses may not match that of the IR due to the geometry of the beam paths. Thus, spectral interferometry

is first performed for the IR pulses with the 50:50 beamsplitters for a range of positions of the translation stage to obtain the delay between the two arms of the interferometer as a function of the position of the translation stage. Next, spectral interferometry is performed in the XUV. From these two measurements, it is possible to obtain an accurate value for the offset between the delay of the IR pulses and the delay of the XUV pulses. After changing BS2 to the amplitude shaping beamsplitter, the delay stage is re-calibrated via SI of the IR. Then using the offset calculated for the 50:50 beamsplitters, the delay stage is set to the position which gives zero delay between the XUV pulses. This calibration is critical because linear chirp appears as a linear phase in the SPIDER interferogram, thus it is important to ensure zero delay for accurate reconstruction.

Using these beamsplitters, we have seen that the spectrum of the harmonics does indeed shift by a small amount. However, we have also observed a change in the spectrum as a function of other pulse parameters, e.g. intensity, chirp, wavefront aberrations and focal position. It is necessary to ensure that these parameters are identical for the two pulses, so that the shear can be attributed to the shift in central frequency. It has previously been shown that the chirp of the IR pulse is transferred to the XUV field,^[223] for example, and thus this effect would be detrimental to a SPIDER measurement. It is essential therefore that the alignment procedure outlined in the previous chapter is followed and all the necessary diagnostics performed to ensure that the interferogram corresponds to a SPIDER measurement.

Using this interferometer allows one to perform both spatially and spectrally encoded SPIDER. It is also universal in the sense that it does not depend on the laser system used to generate the pulses, and is suitable for any pulse duration, bandwidth or carrier frequency, depending on the beamsplitter used for amplitude shaping. In addition, the device is very inexpensive, only requiring a high precision piezo translation stage (e.g. Physik Instrumente GmbH P-753.3CD with 38 μm piezo travel and closed feedback loop via a capacitance sensor) and specially produced beamsplitters (Layertec GmbH custom de-

sign). It is possible to change the reflectivity curve to match the bandwidth and shear required for the particular experiment. Using the amplitude shaping via beamsplitters in this ‘antisymmetric’ fashion also ensures no spectral intensity is lost. However, using an interferometer does mean that half of the energy is wasted in the other output port. Note that a pulse shaper will be even more inefficient: half of the energy is lost from the amplitude shaping in addition to the intrinsic inefficiency of the device. If the shaper is placed before an amplifier, this loss may be recovered. However, the amplification may change the pulse shape.

In our experiments, we have been limited by the pulse energy available (approximately $400\ \mu\text{J}$ per pulse after pulse compression and thus $100\ \mu\text{J}$ per pulse after the interferometer). This limits the peak power for a 7 fs pulse to $\sim 10^{14}\ \text{Wcm}^{-2}$ per pulse in the gas jet (when using the interferometer) and thus it is not possible to reach the regime of isolated attosecond pulses. Such a capability is now possible on the latest generation laser systems which can output several mJ per pulse and thus, after compression, are suitable for isolated attosecond production in conjunction with our interferometer.

6.2.4 Space-Time Coupling

An alternative geometry for producing spectrally sheared XUV pulses may be realised via utilisation of STC of the driving pulse. If the harmonics are generated with a pulse which exhibits spatial chirp in the focus, then the harmonic field should also be spatially chirped. Interference to produce a SPIDER trace can be achieved via: (1) generating two sources and overlap them with a lateral shear; (2) generating one source then spatially splitting the beam and overlapping the two halves; or (3) performing a Young’s double slit experiment on a single spatially chirped beam such that each slit will correspond to an XUV pulse of a different carrier frequency.

The first two methods have the ability to adjust the shear simply, whereas the double slit method is restricted because the spatial chirp is limited by the slit size and the shear set by the slit separation and spatial chirp. Method two, however, may be experimentally

difficult to implement. The latter two methods rely on generating one source of XUV radiation, and therefore do not suffer from any loss of pulse energy as a result of generating two pulses. It should therefore be possible to characterise isolated attosecond pulses using one of these techniques on the Imperial College laser system. The Young's double slit method also relies on spatial filtering of the XUV and may be problematic for low intensities. However, such an arrangement has already been demonstrated for measuring the spatial coherence of high harmonics.^[214]

The problem with all these methods is that they rely on mapping frequency to the spatial position of the XUV field. However, we have already found that the spatial chirp of the harmonics match that of the driving pulse, although further experiments are required. In addition, it has already been demonstrated theoretically and experimentally that the harmonics exhibit significant STC. Thus, the phase difference between any two points will not only result from the spectral shear but also from the spatial phase of the pulse. However, the amount of STC can be reduced to an insignificant level over a sufficient region of the beam for this method to work. It should also be possible to measure the STC of the pulse via LSI or by performing a calibration using a pulse without spatial chirp.

Generating the spatial chirp

It should prove relatively straightforward to apply spatial chirp to the driving pulse. However, it is necessary to remember that the spatial chirp needs to exist in the focus of the beam, which corresponds to applying angular dispersion prior to the focusing optic. As only a small amount of chirp is required, this can be achieved using a thin wedge. An alternative approach may be to chirp the pulse before any compression techniques for few-cycle pulses, or prior to the amplifier in CPA systems as these methods tend to preserve the spatial chirp. As typical intense ultrashort laser systems (for example CPAs and OPCPAs) rely on space-time coupling for their operation, introducing spatial chirp can also be achieved via a slight misalignment of the compressors. As these are normally

arranged to provide precise control over the STC of the pulses, this may prove an easy and controllable method for introducing the spatial chirp.

6.3 Summary

It is possible to perform XUV SPIDER via direct detection of the photons by utilising the nonlinearity of the high harmonic generation process to create two spectrally sheared replicas of the XUV pulse. Via numerical solutions to the time dependent Schrödinger equation, we have demonstrated that driving the harmonics with a slightly frequency shifted pulse does indeed spectrally shear the harmonics which are produced, provided the frequency shift of the fundamental pulse is small. The technique has also been demonstrated experimentally using a pulse shaper to generate the sheared fundamental pulses.^[222] This technique thus promises to provide more simplicity in the temporal characterisation of attosecond pulses and attosecond pulse trains and is inexpensive compared to PES methods. In addition, the sensitivity of XUV spectrometry provides a higher signal to noise ratio, sufficient to allow single shot acquisition.

Two geometries for performing XUV-SPIDER have been discussed: (1) collinear with a delay between the two pulses to generate fringes in the spectral domain and (2) spatially encoded arrangement via the generation of two spatially separated XUV pulses which spatially overlap after propagation to generate fringes in the spatial domain. The former method can be achieved using a 1D spectrometer, although spatially resolving the interferogram will provide information on the space-time coupling of the pulse.

By using the spectrally encoded arrangement, it is possible to generate the shear at the fundamental wavelength using a traditional pulse shaper. However, it is necessary to measure the spectral phase and amplitude and apply the correct filters to generate pulses which appear to be spectrally sheared replicas of each other. Provided ionisation is limited, using this arrangement also guarantees that both pulses experience the same optical paths and thus ensures equal dispersion and generating conditions. However, this

limits the maximum pulse energy which can be used, otherwise significant ionisation from the first pulse will not only cause the second pulse to experience a different generating region, but can also cause blue shifting and defocusing. Additionally, the need to sample high periodicity fringes means that a high resolution spectrometer is required, thus limiting the wavelength and bandwidth which can be measured.

The spatially encoded geometry has the advantage that both sources are spatially separated, and thus they will not suffer from ionisation effects. Thus it is possible to characterise isolated attosecond pulses using this method. By encoding the phase in the spatial fringes, it is possible to sample the spectrum at the sampling limit, and so a low resolution XUV spectrometer can be used. However, it is necessary that both pulses which are produced are identical, but spectrally sheared XUV pulses. This can be achieved by following the alignment procedure and diagnostics presented in the previous chapter.

There are a few methods for generating the spectrally sheared replicas of the fundamental pulse. For the spectrally encoded arrangement, it is possible to utilise a pulse shaper, which allows full control of the shaping process. However, this route is not possible for the spatially encoded method unless two shapers are used. For pulses which are close to their Fourier transform limit, such that they do not exhibit any spectral phase, then it is possible to perform amplitude shaping only, thus simplifying the process further. I propose to do this by using a beamsplitter, which has a linear dependence of its reflectivity/transmission on frequency, in an interferometer to generate two frequency shifted pulses. The interferometer has been designed to allow full control of the relative beam pointing between the two beams to ensure the two foci are spatially separated, but the XUV fields overlap at the detector to remove any effects from the spatial phase.

It may also be possible to utilise space-time coupling by introducing a spatial chirp to the fundamental field prior to the high harmonic generation. As a result, each spatial position will correspond to a spectrally sheared version of the unmodulated attosecond pulse. The spectral phase can then be extracted from an interferogram generated by: (1)

overlapping two spatially chirped sources with a later shear; (2) spatially overlapping two halves of one spatially chirped source or (3) performing a Young's double slit experiment on the spatially chirped source. The latter two methods have the advantage of only requiring a single source of XUV radiation, possibly overcoming the issues associated with generating two identical sources, and reducing threshold value for the pulse power of the laser source.

Conclusion

Time-resolved spectroscopy (TRS) experiments measure dynamical systems via the interaction of light with matter. In general, the system that is to be studied is excited by a pump and the dynamics are probed via a series of small time intervals afterwards, allowing a movie of the dynamics to be built up. As the size of the system being studied decreases, the timescale of the dynamics also tends to decrease. Molecular dynamics can be studied on a femtosecond timescale. However, electronic and nuclear dynamics occur on a sub-femtosecond timescale, therefore requiring attosecond pulses to enable the dynamics to be captured. At present, it is possible to generate pulses at optical wavelengths directly from lasers with durations as short as a few optical cycles, i.e. several femtoseconds. As the duration of the pulses is limited by the carrier frequency, it is necessary to go to smaller wavelengths to generate shorter pulses, e.g. a 100 as pulse requires a wavelength of less than 30 nm. At present, it is not possible to generate these pulses directly from a laser. High harmonic generation (HHG), which involves the interaction of an intense optical pulse with a gas, can be used to generate the small wavelengths and short durations required to perform attosecond spectroscopy. However, the detection and characterisation of the pulses produced is not trivial.

Complete characterisation requires either a reference pulse with known spectral phase, or a time-nonstationary (TNS) filter that has a response time shorter than the pulse

duration to enable self-referencing methods to be performed. As these are the shortest pulses produced, there does not exist a known reference, hence it is necessary to use a self-referencing measurement. Currently, the only available TNS filter for XUV attosecond pulses is photoionisation of electrons from atoms in the presence of a low frequency dressing field. In photoelectron spectroscopy (PES), electrons are ionised by the XUV field so that the pulse spectrum is determined by the electron velocity distribution. If the ionisation occurs in the presence of a dressing field then, under the strong field approximation (SFA), the electrons will be driven classically by the dressing field. The electron spectrum will then pick up a temporal phase which is given by the quasi-classical action of the electron motion. Therefore, the dressing field acts as a temporal phase modulator which can be used to characterise attosecond pulses.

PES is the only method which has characterised isolated attosecond pulses. However, the SNR ratio is extremely small, resulting in very long integration times. This puts stringent stability demands on the driving lasers, including carrier envelope phase (CEP) stabilisation. In addition, the expense and complexity of photoelectron spectrometers limits these methods to few groups worldwide.

In PES, the electron spectrum is a convolution of the XUV spectrum and the response function of the ionising medium. It is possible to deconvolve the magnitude of the response function either from calculations or independently measured XUV spectrum from an XUV spectrometer. However, it is not possible to deconvolve the phase function, which is required for temporal measurements. At present, the response function is shorter than the pulse durations currently being measured, therefore it has a negligible effect. However, for a pulse with a duration reaches as short as tens of attoseconds, the phase of the response function must be taken into account. The only means to do so will be via the detection of XUV photons. However, there does not exist a TNS filter that can be applied to the XUV photons. The only method of complete characterisation of the XUV field is to control the generation process. This then requires excellent control and characterisation of the

driving pulses used for HHG. The parameters which need to be measured and controlled are: (1) pulse power, (2) dispersion, (3) wavefront and (4) the generating medium.

Controlling many-cycle pulses is relatively trivial. For example, very efficient shapers can be used to control the dispersion, and thus temporally shape the pulses. It is possible to control the pulse power with minimal effect on the pulse shape, e.g. through a polarising beamsplitter/waveplate combination. It is possible to implement this control with minimal space-time coupling (STC), dispersion and wavefront aberration (for example, it is possible to use thick beamsplitters, which have better surface quality and less distortion than thin beamsplitters).

Controlling few-cycle pulses is an extremely challenging task. For example, even 1 mm of glass can cause significant temporal broadening. In addition, there does not exist a simple means to control the dispersion of these pulses. The most common method of dispersion control is via chirped mirrors. However, these result in significant energy loss and have limited shaping capabilities (e.g. it is only possible to add small, discrete quantities of negative GDD). In addition, the bandwidth of the mirrors is limited so that only pulses as short as 5–7 fs can be supported. Finally, it is very difficult to control the pulse power without adding dispersion or wavefront aberration to the pulse. Therefore, a significant amount of work has been needed in developing methods to characterise and control few-cycle pulses with the required accuracy and precision.

I have developed an experimental configuration, system of alignment procedures and diagnostics that are required to ensure the possibility to control the driving pulses with the accuracy to perform spectral interferometry. I have performed a set of lateral shearing interferometry (LSI) experiments to test the ability to control the driving pulses. However, LSI is not only a method to test the ability to perform self-referencing spectral interferometry, but it can be used to measure the spectrally dependent spatial phase (i.e. wavefront) of attosecond pulses generated via HHG. Due to the dependence of the atomic phase on the medium and intensity of the driving pulse, the phase-matching conditions vary signif-

icantly with the focusing geometry. For example, the harmonics are produced efficiently on-axis when the gas is positioned slightly after the focus. However, if the gas is positioned slightly before the focus, then efficient harmonic generation occurs off-axis with a large divergence, thus producing an annular beam. In addition, the different quantum trajectories result in a different atomic phase for each trajectory, resulting in different phase-matching conditions. From the LSI data, I have shown that it is possible to measure STC of the emitted field (e.g. I showed how the angular dispersion varies as a function of the gas position relative to the focus) and that it is possible to identify different quantum trajectories or generating regimes (i.e. whether the harmonics are produced in the cut-off or plateau regimes). However, it was not possible to obtain accurate values for the parameters used in the LSI experiment. Therefore it was not possible to perform a quantitative analysis of the harmonic wavefront or STC. However, I have shown that it is possible to control the few-cycle driving pulses with required accuracy and a qualitative analysis of the harmonic field showed that it is possible to measure STC of XUV attosecond pulses with high accuracy.

In chapter 6, I showed how it is possible to extend the method of spectral phase interferometry for direct electric field reconstruction (SPIDER) to enable the complete characterisation of attosecond pulses via the direct detection of the XUV photons. The method is based on generating two sources of harmonics from two spectrally-sheared replicas of the driving pulses. As a result of the spectral shear of the driving pulse, the harmonics are also spectrally sheared copies of each other, provided the spectral shear remains small. I demonstrated that XUV-SPIDER is viable through numerical simulations of the 3D time-dependent Schrödinger equation.

I proposed two possible geometries for performing XUV-SPIDER: (1) spectrally encoded and (2) spatially encoded. In the former method, the harmonics are generated by two collinear, spectrally-sheared and temporally separated driving pulses, resulting in fringes in the spectral domain. This geometry may allow the use of pulse shapers for

generating the sheared driving pulses. However, due to ionisation from the first pulse in the gas, it is necessary to limit the power used so that the second pulse experiences the same medium as the first pulse. For argon, it was shown that the peak intensity must be below $2 \times 10^{14} \text{ Wcm}^{-2}$ to prevent this. However, it is unlikely that spectrally encoded XUV-SPIDER is suitable for characterising isolated attosecond pulses. This is because the spectral resolution required to resolve the fringes is too demanding for typical XUV spectrometers and the driving pulse energy is limited to a value that is insufficient to generate an isolated attosecond pulse.

As spatially encoded XUV-SPIDER encodes the spectral phase in the spatial dimension, it is possible to sample the spectrum at the sampling limit, enabling the spectrum of attosecond pulses to be measured. In addition, the two harmonic sources are generated from two spatially separated regions in the gas jet, therefore there is no limit on the driving pulse energy. However, it is possible that the two generating regions may be slightly different and thus the two harmonic pulses will not be sheared replicas. However, through the LSI experiments, I have shown that this is not the case provided the correct alignment procedures are followed.

There exist many methods for complete characterisation of near-infrared (NIR) femtosecond pulses. Spectral phase interferometry for direct electric field reconstruction (SPIDER) and frequency-resolved optical gating (FROG) are the two most commonly implemented techniques. However they are not ideally suited to the measurement of few-cycle NIR pulses in their conventional form. This is because few-cycle pulses tend to exhibit: (1) an octave spanning bandwidth, (2) modulated spectrum and (3) space-time coupling. In addition, it is often required to have online reconstruction, to enable optimisation of the laser system, and single-shot acquisition so that it is possible to characterise and minimise the shot-to-shot fluctuations.

Conventional SPIDER encodes the spectral phase in the spacing of spectral fringes and therefore significantly oversamples the pulse spectrum. This is a problem for few-cycle

pulses due to the excessively large bandwidth. In addition, the highly modulated spectrum can cause problems for the reconstruction algorithm. However, as SPIDER measures a 1D data set to reconstruct a 1D field, it is possible to measure STC by spatially resolving the interferogram.

FROG has the advantage over SPIDER because it is possible to record the spectrum at the sampling limit. However, it requires a 2D data set to reconstruct the 1D field. If the pulse exhibits STC, then this will distort the FROG trace (e.g. spatial chirp introduces an asymmetry in single-shot SHG-FROG and pulse-front tilt spatially shifts the spectrogram). It may be possible to determine the STC provided it has a simple functional form. However, FROG cannot measure arbitrary STC, which is typically exhibited by few-cycle pulses.

I have developed a variant of spectral interferometry which uses a spatially encoded arrangement (SEA-) for SPIDER. This method encodes the phase in the spacing of spatial fringes. As a result, it is possible to record the interferogram at the sampling limit. SEA-SPIDER has many other features which make it ideally suited to measuring few-cycle pulses. For example, it is not necessary to make a replica of the test pulse. If this were necessary, then this will always result in additional dispersion which needs to be calibrated, or worst still an imbalance in the dispersion of the two replicas. The two signal pulses are guaranteed to interfere at zero delay. Therefore there is no additional chirp added to the reconstructed pulse as a result of an error in the calibration of the delay used in conventional SPIDER. The value of the shear can be adjusted to enable a range of pulse shapes to be measured. For complicated pulses that have regions of low spectral intensity, this feature can prove extremely useful, or even necessary, to ensure accurate reconstruction.

The SEA-SPIDER interferograms are intuitive. The contours of the fringes directly map the gradient of the spectral phase (i.e. the chirp) of the pulse. Therefore it is possible to perform online optimisation of the laser without the need to reconstruct the pulse.

In addition, the intuitive nature and redundancy of the interferograms also allow certain features of the pulses to be identified. For example, a low intensity sub-pulse resulting from insufficient compensation of the GDD in the wings of the spectrum from a few-cycle oscillator, which employs chirped mirrors for the intracavity dispersion, could be identified by the slope of the fringe contours in the interferogram. This feature was not identified in the conventional SPIDER because the SNR was too low to enable reconstruction of the phase in these spectral regions. Finally, SEA-SPIDER enables complete and arbitrary STC characterisation of the pulse, since it spatially resolves the spectrum. This was shown by measuring the pulse-front tilt of the few-cycle oscillator.

I have shown that SEA-SPIDER is an extremely robust and accurate method for characterising few-cycle NIR pulses, even if they exhibit STC. The method was compared to other complete characterisation methods, and it was shown that it was possible to identify features that were not reconstructed with the conventional methods. SEA-SPIDER should prove extremely useful for both femtosecond and attosecond spectroscopy. As attosecond pulses are generated from the interaction of intense few-cycle pulses with a gas, SEA-SPIDER can be used to optimise the driving pulse and for seeding into simulations of the HHG process.

7.1 Outlook

Lateral shearing interferometry

Up to now, I have only been able to perform a qualitative analysis of the LSI results. This is because I have not been able to perform all the measurements necessary to accurately calculate the parameters used for the experiment. However, I have managed to develop a systematic process for obtaining the required parameters and for controlling the driving laser. Therefore a quantitative analysis of the results is only limited by the availability of the laser system to perform the calibration measurements to ensure that reconstructions of the spatial phase are consistent with the theory. It is expected that these measurements

will be performed in the immediate future and then compared with simulations of the HHG process. In addition, it should be possible to measure the STC of the harmonics as a function of the driving parameters, such as the distance between the gas jet and the focus, or as a function of spatial filtering, intensity, duration and amount of STC on the pulse. Another route to explore is to measure how the STC on the harmonics varies with different gas species. Using other noble gases should provide a simple relationship as it is only the ionisation potential which varies. This can be used to confirm our results. However, more complicated species such as molecules or clusters may provide more complicated and possibly more interesting results.

XUV SPIDER

It is not possible to perform XUV SPIDER until it has been demonstrated quantitatively that it is possible to generate identical harmonics from two independent sources. This therefore involves the completion of the LSI experiments. Once this has been achieved, it should be possible to apply a spectral shear between the driving pulses. In addition, there are many different routes of generated the sheared harmonics. Not all of these methods have been examined experimentally. For example, applying a spatial chirp to the beam. However, mapping the spectral shear to the spatial profile of the harmonics requires knowledge of the STC of the harmonics, which will also require the LSI experiments to be completed.

Attosecond TRS via spectral interferometry

The overall goal of XUV spectral interferometry is to provide a simple and robust means of performing attosecond spectroscopy. Currently, the only methods available are based on photoelectron spectroscopy. Therefore, it is necessary to make some assumptions about the response function of the ionising medium. Spectral interferometry will complement PES methods and provide a means to calibrate the response function.

In addition, the geometry for the experiments for performing LSI and XUV-SPIDER also allows reference-based spectral interferometry to be performed. This is achieved by keeping the driving parameters of one arm of the interferometer fixed whilst varying the parameters in the other arm. For example, the pulse intensity, dispersion or wavefront could be adjusted. Another possibility is to measure the phase of the harmonics as a function of the alignment angle of molecules. This can be achieved by applying an alignment laser to one source of the harmonics only. Alternatively, one may perform homodyne spectral interferometry. This would involve scattering one of the XUV sources from a sample and interfering it with the other source. As it is likely that the scattered field is very weak, the reference arm will act as an intense local oscillator. Due to the interference between the two beams, this effectively increases the signal from the weakly scattered pulse.

It is clear that there are many experimental configurations available, and their implementation is currently only limited by the amount of time that is invested into investigating their feasibility. In this thesis, I have demonstrated an accurate means of measuring the few-cycle driving pulses used for HHG in the form of SEA-SPIDER. This should allow one to seed simulations of HHG with a more accurate representation of the driving pulses, including full space-time coupling. In addition, I have demonstrated that it should be possible to control these pulses with the required accuracy and precision for performing self-referencing XUV spectral interferometry from HHG. The developments that I have presented only mark the start of what might prove to be a very useful and interesting area of research, not only for physicists but for a wide range of traditional disciplines such as biology, chemistry and materials scientists.

Appendix **A**

Maxwell's equations and the electromagnetic wave equation

A.1 Maxwell's equations

The description of electric (\mathbf{E}) and magnetic (\mathbf{B}) fields are given by the four Maxwell equations (in SI units):

$$\nabla \cdot \mathbf{D} = \rho \tag{A.1a}$$

$$\nabla \cdot \mathbf{B} = 0 \tag{A.1b}$$

$$\nabla \times \mathbf{E} = -\frac{\partial \mathbf{B}}{\partial t} \tag{A.1c}$$

$$\nabla \times \mathbf{H} = \mathbf{J} + \frac{\partial \mathbf{D}}{\partial t}, \tag{A.1d}$$

where ρ is the electric charge density; \mathbf{J} the current density; $\mathbf{D} = \varepsilon_0 \mathbf{E} + \mathbf{P}$ is the electric displacement in a medium of polarisability \mathbf{P} ; $\mathbf{H} = \mathbf{B}/\mu_0 - \mathbf{M}$ is the magnetic field strength in a material of magnetisation \mathbf{M} ; ε_0 and μ_0 are the electric permittivity and magnetic permeability of free space respectively. Equation (A.1a) is Gauss' law for the electric field and physically corresponds to the fact that the total flux through any closed surface is equal to the charge contained within that surface. Equation (A.1b) is Gauss' law

for magnetism and states an absence of magnetic monopoles — i.e. there is no equivalent of charge for magnetism. The latter two equations are Faraday's law (A.1c) and Ampère's law (A.1d) which relate the electric and magnetic fields to the material properties.

A.2 The wave equation

Taking the form of Maxwell's equations (A.1) for an isotropic medium using the material equations for the electric and magnetic field strengths ($\mathbf{D} = \varepsilon \mathbf{E}$ and $\mathbf{B} = \mu \mathbf{H}$, where ε and μ are the relative permittivity and permeability respectively) in the absence of any currents and free charges ($\mathbf{J} = \rho = 0$) so that dividing (A.1c) by μ and taking the curl, followed by substitution of (A.1d) differentiated with respect to time yields

$$\nabla \times \left(\frac{1}{\mu} \nabla \times \mathbf{E} \right) = -\varepsilon \ddot{\mathbf{E}}. \quad (\text{A.2})$$

Using the vector identities $\nabla \times (u\mathbf{v}) = u(\nabla \times \mathbf{v}) + (\nabla u) \times \mathbf{v}$ and $\nabla \times \nabla \times = \nabla \nabla \cdot - \nabla^2$, (A.2) becomes

$$\nabla^2 \mathbf{E} - \varepsilon \mu \ddot{\mathbf{E}} + (\nabla \ln \mu) \times \nabla \times \mathbf{E} - \nabla (\nabla \cdot \mathbf{E}) = 0. \quad (\text{A.3})$$

Using the material equations for \mathbf{D} , and the vector identity $\nabla \cdot u\mathbf{v} = u \nabla \cdot \mathbf{v} + \mathbf{v} \cdot \nabla u$ with (A.1a) so that $\nabla \cdot \mathbf{E} = -1/\varepsilon \mathbf{E} \cdot \nabla \varepsilon = 0$ can be substituted into (A.3) to yield

$$\nabla^2 \mathbf{E} - \varepsilon \mu \ddot{\mathbf{E}} + (\nabla \ln \mu) \times \nabla \times \mathbf{E} + \nabla (\mathbf{E} \cdot \nabla \ln \varepsilon) = 0. \quad (\text{A.4})$$

In a similar manner, it is possible to obtain an equation for \mathbf{H} alone:

$$\nabla^2 \mathbf{H} - \varepsilon \mu \ddot{\mathbf{H}} + (\nabla \ln \varepsilon) \times \nabla \times \mathbf{H} + \nabla (\mathbf{H} \cdot \nabla \ln \mu) = 0. \quad (\text{A.5})$$

Equations (A.4) and (A.5) describe the time and spatial dependence of electric and magnetic fields for isotropic media. However, if the medium is also homogeneous then $\nabla (\ln \varepsilon) = \nabla (\ln \mu) = 0$, and thus these equations reduce to the electric and magnetic wave

equations:

$$\nabla^2 \mathbf{E} - \varepsilon \mu \ddot{\mathbf{E}} = 0 \quad (\text{A.6})$$

$$\nabla^2 \mathbf{H} - \varepsilon \mu \ddot{\mathbf{H}} = 0, \quad (\text{A.7})$$

where the velocity of propagation is given by $v^2 = 1/(\varepsilon\mu)$. The refractive index of a material is given by $n = c/v = \sqrt{\varepsilon\mu/(\varepsilon_0\mu_0)} = \sqrt{\varepsilon_r\mu_r}$ where $\varepsilon_r = \varepsilon/\varepsilon_0$ and $\mu_r = \mu/\mu_0$ are the relative permittivity and permeability. In a vacuum, where $\mu = \mu_0$ and $\varepsilon = \varepsilon_0$, the speed of light is given by

$$c = \frac{1}{\sqrt{\varepsilon_0\mu_0}} \quad (\text{A.8})$$

Equation (A.6) represents three independent scalar wave equations, one for each directional component of the electric field vector in Cartesian co-ordinates for an isotropic, homogenous medium. Gauss' law, (A.1a), requires a spatially bounded beam to have an electric field component in the propagation direction, although each non-zero electric field component will satisfy (A.6). The electric field component in the propagation direction is only large in the vicinity where the spatial confinement is on the order of a wavelength (i.e. for numerical apertures $\gtrsim 0.6$). In addition, extending the treatment to non-isotropic media results in a coupling between the individual electric-field components (for further details, see e.g. Born and Wolf^[33]).

A.3 Analytic representation of the electromagnetic field

Any real signal can be represented by a Fourier transform of a Hermitian function:

$$f(t) = \int_{-\infty}^{\infty} d\omega F(\omega) e^{-i\omega t}, \quad (\text{A.9})$$

where

$$F(\omega) = F^*(-\omega) \quad (\text{A.10})$$

guarantees that the signal is real. Clearly, the plane wave solutions (2.2) to the electric-field wave equation (A.6) results in the function, $F(\omega)$, being written in the form, for the electric field:

$$F(\omega) = \int_{-\infty}^{\infty} d\mathbf{k} \tilde{\xi}_i(\omega, \mathbf{k}) e^{i\tilde{\phi}(\omega, \mathbf{k})} e^{i\mathbf{k} \cdot \mathbf{r}}, \quad (\text{A.11})$$

where $\tilde{\xi}_i(\omega, \mathbf{k}) \geq 0$ are the real amplitudes coefficients of the plane wave solutions with frequency ω and wave vector \mathbf{k} for the i^{th} component of the electric field vector. For convenience, it is possible to introduce the analytic electric field, $E_i(t, \mathbf{r}) = |E_i(t, \mathbf{r})| e^{i\phi(t, \mathbf{r})}$, which is calculated via the single-sided Fourier transform over the positive frequency amplitudes:

$$E_i(t, \mathbf{r}) = \int_0^{\infty} d\omega \int_{-\infty}^{\infty} d\mathbf{k} \tilde{\xi}_i(\omega, \mathbf{k}) e^{i\tilde{\phi}(\omega, \mathbf{k})} e^{i\mathbf{k} \cdot \mathbf{r}} e^{-i\omega t}. \quad (\text{A.12})$$

It is clear from equations (A.9)–(A.12) and equation (2.2) that the real electric field is given by twice the real part of the analytic electric field:

$$\begin{aligned} \xi_i(t, \mathbf{r}) &= 2\Re\{E_i(t, \mathbf{r})\} \\ &= E_i(t, \mathbf{r}) + E_i^*(t, \mathbf{r}) \\ &= \int_0^{\infty} d\omega F(\omega) e^{-i\omega t} + \int_0^{\infty} d\omega F^*(\omega) e^{i\omega t} \\ &= \int_0^{\infty} d\omega F(\omega) e^{-i\omega t} + \int_{-\infty}^0 d\omega F(\omega) e^{-i\omega t} \\ &= f(t). \end{aligned} \quad (\text{A.13})$$

It is also possible to introduce the analytic electric field spectrum, $\tilde{E}_i(\omega, \mathbf{k}) = |\tilde{E}_i(\omega, \mathbf{k})| e^{i\tilde{\phi}(\omega, \mathbf{k})}$ corresponding to the positive frequency components only:

$$\tilde{E}_i(\omega, \mathbf{k}) = \begin{cases} \tilde{\xi}_i(\omega, \mathbf{k}) e^{i\tilde{\phi}(\omega, \mathbf{k})} & \text{if } \omega \geq 0 \\ 0 & \text{otherwise,} \end{cases} \quad (\text{A.14})$$

such that the analytic electric field and analytic electric field spectrum are related via a

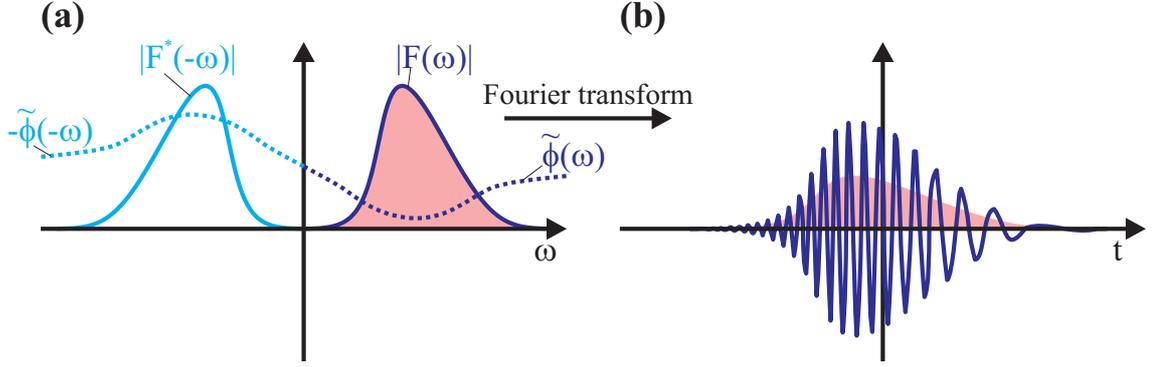


Figure A.1: Analytic representation of a real signal. (a) The real signal can be decomposed into positive (blue) and negative (cyan) Hermitian conjugate Fourier frequency components (the absolute value is shown as a solid line, the phase as a dashed line). (b) The real signal is the Fourier transform over all frequency components. The analytic signal (shaded region) is given by the Fourier transform over the positive frequency components only.

Fourier transform:

$$E_i(t, \mathbf{r}) = \int_{-\infty}^{\infty} d\omega \int_{-\infty}^{\infty} d\mathbf{k} \tilde{E}_i(\omega, \mathbf{k}) e^{i\mathbf{k}\cdot\mathbf{r}} e^{-i\omega t} \quad (\text{A.15})$$

The analytic representation is shown schematically in figure A.1, where the real electric field is given by the Fourier transform of both the positive and negative frequency components which are Hermitian conjugates of each other. The analytic signal is the Fourier transform of the positive components only (setting the negative components to zero) and is thus half the amplitude of the real field.

The analytic representations follow the usual mathematics of complex numbers and allows a more convenient means of performing operations on the electric field. Note that such a representation is only meaningful in the case of linear optics and when the electric field spectrum tends to zero as the frequency approaches zero

$$\lim_{\omega \rightarrow 0^+} \tilde{E}_i(\omega, \mathbf{k}) = 0. \quad (\text{A.16})$$

In the case of nonlinear optics, it is necessary to include the interactions between the positive and negative frequency components. For example, four wave mixing can result in the

interaction of three positive frequencies, yielding the third harmonic; or the interaction of two positive and one negative frequency, yielding the fundamental frequency propagating through an intensity dependent refractive index (known as the Kerr effect).

A.4 Fourier transforms

Using the definition of the Fourier transform pairs:

$$\begin{aligned} F(\omega) &= \int_{-\infty}^{\infty} dt f(t) e^{-i\omega t} \\ &= \mathfrak{F}\{f(t); t \rightarrow \omega\} \end{aligned} \quad (\text{A.17})$$

$$\begin{aligned} f(t) &= \frac{1}{2\pi} \int_{-\infty}^{\infty} d\omega F(\omega) e^{i\omega t} \\ &= \mathfrak{F}^{-1}\{F(\omega); \omega \rightarrow t\}, \end{aligned} \quad (\text{A.18})$$

it is clear from equation (A.15) that the analytic signal for the electric field, $\mathbf{E}(\mathbf{r}, t)$ is given by the Fourier transform over frequency of the analytic spectrum, followed by the three dimensional inverse Fourier transform over each wave vector component, scaled by factors of 2π . However, due to the constraints on the wave vector, this three dimensional Fourier transform actually reduces to a two dimensional Fourier transform over two components of the wave vector. The two dimensional Fourier transform pairs are defined as

$$\begin{aligned} F(k_x, k_y) &= \int_{-\infty}^{\infty} dx \int_{-\infty}^{\infty} dy f(x, y) e^{-i(k_x x + k_y y)} \\ &= \mathfrak{F}_2\{f(x, y); x \rightarrow k_x, y \rightarrow k_y\} \end{aligned} \quad (\text{A.19})$$

$$\begin{aligned} f(x, y) &= \frac{1}{4\pi^2} \int_{-\infty}^{\infty} dk_x \int_{-\infty}^{\infty} dk_y F(k_x, k_y) e^{i(k_x x + k_y y)} \\ &= \mathfrak{F}_2^{-1}\{F(k_x, k_y); k_x \rightarrow x, k_y \rightarrow y\}. \end{aligned} \quad (\text{A.20})$$

Note that these definitions correspond to the conventional mathematical forms of the Fourier transform pairs, which correspond to inverted versions of the definitions of the

Fourier transform pairs used in optics. However, I have used the same sign convention for the phase, frequency and wave vector components so that it is merely the definition of the forward/backward transform which differs. Thus smaller times correspond to earlier arrival times, and positive phase corresponds to normal dispersion.

A.4.1 Cylindrically symmetric beams

If the beam has cylindrical symmetry, then the Fourier transform over the two spatial co-ordinates, (x, y) , can be converted to a Hankel transform over the radial co-ordinate, r :

$$\begin{aligned}
 E(r) &= \int_{-\infty}^{\infty} dk_x \int_{-\infty}^{\infty} dk_y \tilde{E}(k_x, k_y) e^{i(k_x x + k_y y)} \\
 &= 2\pi \int_0^{\infty} dk_r \tilde{E}(k_r) k_r J_0(k_r r) \\
 &= \mathfrak{H} \left\{ \tilde{E}(k_r); k_r \rightarrow r \right\},
 \end{aligned} \tag{A.21}$$

where $k_r^2 = k_x^2 + k_y^2$ and J_0 is the zeroth order Bessel function of the first kind. The inverse Hankel transform is then given by:

$$\begin{aligned}
 \tilde{E}(k_r) &= \frac{1}{2\pi} \int_0^{\infty} dr E(r) r J_0(k_r r) \\
 &= \mathfrak{H}^{-1} \{ E(r); r \rightarrow k_r \}.
 \end{aligned} \tag{A.22}$$

Note the scaling of the transforms, which do not follow the symmetrical form normally seen, which is due to using the asymmetric Fourier transform pairs in the definition of the temporal and spatial fields (equations (A.19) and (A.20) respectively).

The Hankel transform can prove useful in both numerical and analytical treatments of the electric field. Numerically, it allows one to represent the spatial field in terms of a one dimensional, rather than two dimensional, complex quantity and thus can reduce memory constraints. Although there does not exist a fast Hankel transform, in analogy

to the fast Fourier transform (FFT), it is possible to compute the integral via a matrix operation provided the data is sampled on a specific grid, resulting in the quasi discrete Hankel transform (QDHT).^[224]

Amplitude shaping beamsplitters for XUV-SPIDER

Below are the data for the reflectivity of the beamsplitters used for amplitude shaping of the driving pulses to be used for XUV-SPIDER via HHG (chapter 6). In each figure, the reflectivity for the p (i.e. vertical) polarisation at 45° incidence angle is shown as a function of wavelength; note that it is possible to select an arbitrary transmission curve and for either polarisation and arbitrary incidence angle within the limits of the fabrication process. Figure B.1 shows the transmission curve for the broadband 50:50 beamsplitter. The curve is approximately flat over the whole bandwidth of our laser (~ 600 - 950 nm). Figure B.2 shows the reflection coefficient for a beamsplitter which causes a shift in the mean wavelength of ~ 1 - 2 nm depending on the bandwidth. This shift should be sufficient to characterise the highest order harmonics with a wavelength of ~ 12 nm (e.g. 61st harmonic of 750 nm pulses). The size of the shift in the mean wavelength for the reflectivity curve shown in figure B.3 is on the order of 10 nm. This should be suitable for measuring mid-order harmonics with a wavelength on the order of tens of nanometers (e.g. 31st harmonic of 750 nm pulses).

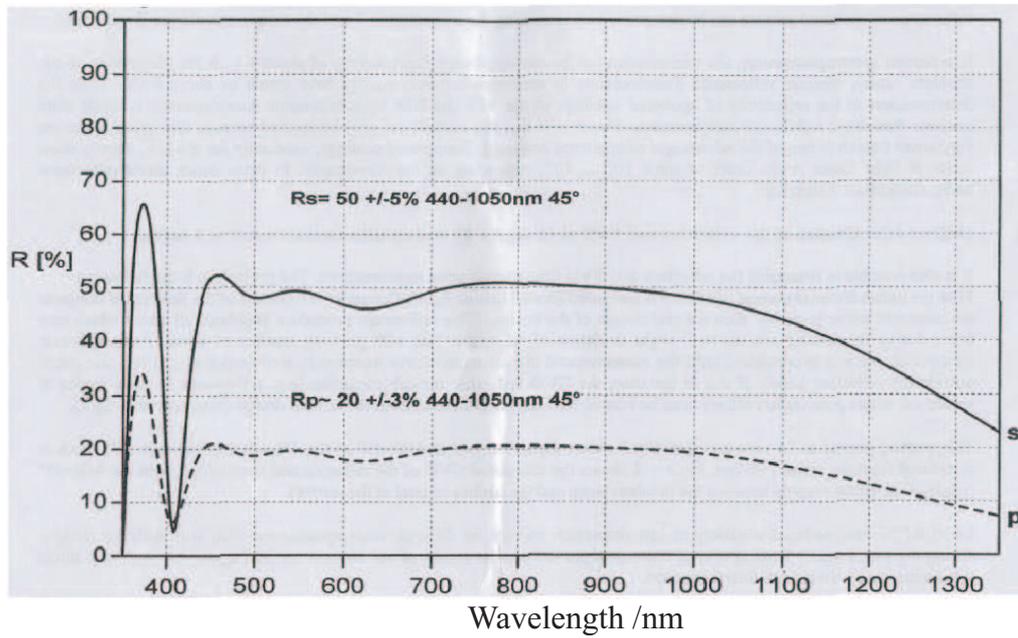


Figure B.1: Reflectivity curve for a broadband 50:50 beamsplitter.

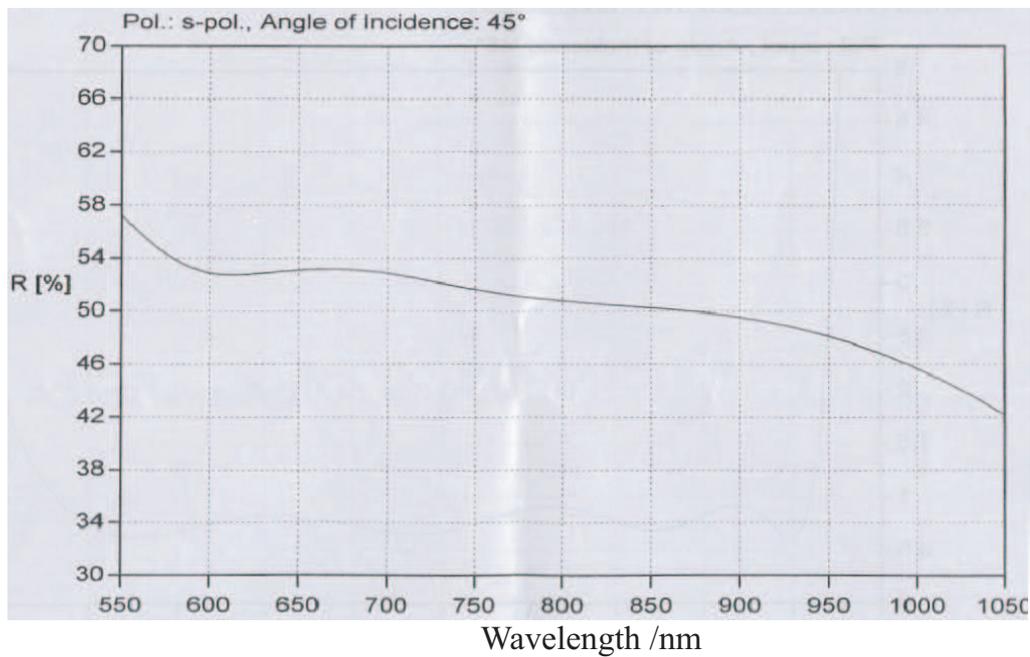


Figure B.2: Reflectivity curve for a broadband beamsplitter which is slightly linearly dependent on frequency.

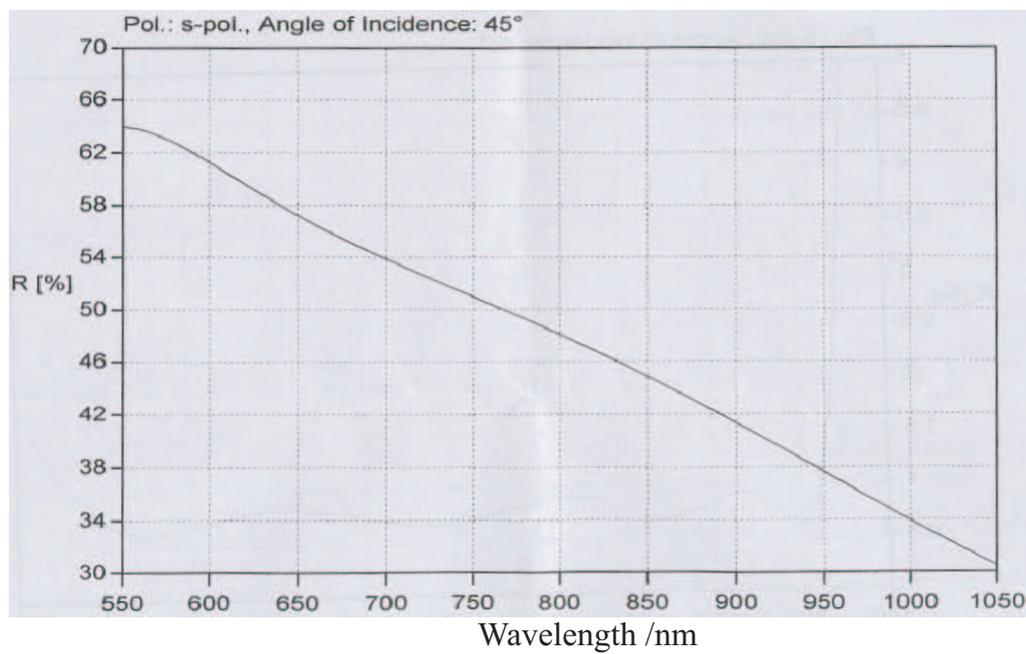


Figure B.3: Reflectivity curve for a broadband beamsplitter which is strongly linearly dependent on frequency.

Bibliography

- [1] G.R. Fleming. *Chemical applications of ultrafast spectroscopy*. Oxford University Press New York, 1986.
- [2] J. Misewich, J.H. Glowia, J.E. Rothenberg, and P.P. Sorokin. Subpicosecond uv kinetic spectroscopy: Photolysis of thallium halide vapors. *Chemical Physics Letters*, 150(6):374–379, 1988.
- [3] M.J. Rosker, M. Dantus, and A.H. Zewail. Femtosecond real-time probing of reactions. I. The technique. *The Journal of Chemical Physics*, 89:6113, 1988.
- [4] A.H. Zewail. Laser Femtochemistry. *Science*, 242(4886):1645, 1988.
- [5] R.B. Cundall and R.E. Dale. *Time-resolved fluorescence spectroscopy in biochemistry and biology*. Plenum Press, 1983.
- [6] E.P. Diamandis. Immunoassays with time-resolved fluorescence spectroscopy: principles and applications. *Clin Biochem*, 21(3):139–50, 1988.
- [7] P. Kowalczyk, C. Radzewicz, J. Mostowski, and I.A. Walmsley. Time-resolved luminescence from coherently excited molecules as a probe of molecular wave-packet dynamics. *Phys. Rev. A*, 42(9):5622–5626, Nov 1990.
- [8] H. Ito, H. Yokoyama, S. Murata, and H. Inaba. Generation of picosecond optical pulses with highly RF modulated AlGaAs DH laser. *Quantum Electronics, IEEE Journal of*, 17(5):663–670, 1981.
- [9] M. Blanchet and J. P. Gex. Pockels cell shutter operating in the 100 picosecond range - Performances and applications to ultra high speed interferometry. *SMPTE JOURNAL*, 85(1):150–156, 1975.
- [10] Huston A.E. High-speed photography and photonic recording. *Journal of Physics E: Scientific Instruments*, 11(7):601–609, 1978.
- [11] H. Ihee, V. A. Lobastov, U. M. Gomez, B. M. Goodson, R. Srinivasan, C. Y. Ruan, and A. H. Zewail. Direct imaging of transient molecular structures with ultrafast diffraction. *Science*, 291(5503):458–462, January 2001.
- [12] R. J. Levis, G. M. Menkir, and H. Rabitz. Selective bond dissociation and rearrangement with optimally tailored, strong-field laser pulses. *Science*, 292(5517):709–713, April 2001.

-
- [13] Y. Hirayama and M. Obara. Heat effects of metals ablated with femtosecond laser pulses. *Applied Surface Science*, 197:741–745, September 2002.
- [14] A. Cavalleri, S. Wall, C. Simpson, E. Statz, D. W. Ward, K. A. Nelson, M. Rini, and R. W. Schoenlein. Tracking the motion of charges in a terahertz light field by femtosecond x-ray diffraction. *Nature*, 442(7103):664–666, August 2006.
- [15] A. Föhlisch, P. Feulner, F. Hennies, A. Fink, D. Menzel, D. Sanchez-Portal, P. M. Echenique, and W. Wurth. Direct observation of electron dynamics in the attosecond domain. *Nature*, 436(7049):373–376, July 2005.
- [16] C. Dorrer and M. Joffre. Characterization of the spectral phase of ultrashort light pulses. *Comptes Rendus De L Academie Des Sciences Serie Iv Physique Astrophysique*, 2(10):1415–1426, December 2001.
- [17] V. Wong and I. A. Walmsley. Analysis of ultrashort pulse-shape measurement using linear interferometers. *Optics Letters*, 19(4):287–289, February 1994.
- [18] V. Wong and I. A. Walsmley. Linear filter analysis of methods for ultrashort-pulse-shape measurements. *Journal Of The Optical Society Of America B-Optical Physics*, 12(8):1491–1499, August 1995.
- [19] A. S. Wyatt, I. A. Walmsley, G. Stibenz, and G. Steinmeyer. Sub-10 fs pulse characterization using spatially encoded arrangement for spectral phase interferometry for direct electric field reconstruction. *Optics Letters*, 31(12):1914–1916, June 2006.
- [20] G. Stibenz, C. Ropers, C. Lienau, C. Warmuth, A. S. Wyatt, I. A. Walmsley, and G. Steinmeyer. Advanced methods for the characterization of few-cycle light pulses: a comparison. *Applied Physics B-Lasers And Optics*, 83(4):511–519, June 2006.
- [21] E. M. Kosik, L. Corner, A. S. Wyatt, E. Cormier, I. A. Walmsley, and L. F. Dimauro. Complete characterization of attosecond pulses. *Journal Of Modern Optics*, 52(2-3):361–378, January 2005.
- [22] E. Cormier, I. A. Walmsley, E. M. Kosik, A. S. Wyatt, L. Corner, and L. F. DiMauro. Self-referencing, spectrally, or spatially encoded spectral interferometry for the complete characterization of attosecond electromagnetic pulses. *Physical Review Letters*, 94(3):033905, January 2005.
- [23] E. Cormier, L. Corner, E. M. Kosik, I. A. Walmsley, and A. S. Wyatt. Spectral phase interferometry for complete reconstruction of attosecond pulses. *Laser Physics*, 15(6):909–915, June 2005.
- [24] M. Lax, W.H. Louisell, and W.B. McKnight. From Maxwell to paraxial wave optics. *Physical Review A*, 11(4):1365–1370, 1975.
- [25] M.O. Scully and MS Zubairy. Simple laser accelerator: Optics and particle dynamics. *Physical Review A*, 44(4):2656–2663, 1991.
- [26] S. Quabis, R. Dorn, M. Eberler, O. Glöckl, and G. Leuchs. Focusing light to a tighter spot. *Optics Communications*, 179(1-6):1–7, 2000.
- [27] J. Paye. The Chronocyclic Representation of Ultrashort Light Pulses. *I.E.E.E. Journal of Quantum Electronics*, 28(10):2262–2273, october 1992.

- [28] S. Feng, H. G. Winful, and R. W. Hellwarth. Gouy shift and temporal reshaping of focused single-cycle electromagnetic pulses. *Optics Letters*, 23:385–387, March 1998.
- [29] S. Feng and H.G. Winful. Physical origin of the Gouy phase shift. *Opt. Lett.*, 26(8):485–487, 2001.
- [30] F. Lindner, G. G. Paulus, H. Walther, A. Baltuška, E. Goulielmakis, M. Lezius, and F. Krausz. Gouy phase shift for few-cycle laser pulses. *Physical Review Letters*, 92(11):113001, 2004.
- [31] A.E. Siegman. Hermite-Gaussian functions of complex argument as optical-beam eigenfunctions. *J. Opt. Soc. Am.*, 63:1093–1094, 1973.
- [32] P.A. Bélanger. Beam propagation and the ABCD ray matrices. *Z.*, 16(4):196–1, 1991.
- [33] M. Born and E. Wolf. *Principles of Optics: Electromagnetic Theory of Propagation, Interference and Diffraction of Light*. Cambridge University Press, 7 edition, October 1999.
- [34] R.W. Boyd. *Nonlinear Optics*. Academic Press, 2003.
- [35] K.W.D. Ledingham, P. McKenna, and R.P. Singhal. Applications for Nuclear Phenomena Generated by Ultra-Intense Lasers. *Science*, 300(5622):1107–1111, 2003.
- [36] P. Gires and C. R. Tournois. Interferometre utilisable pour la compression d’impulsions lumineuses modulees en frequence. *de l’Acad. Sci.*, 258:6112, 1964. An interferometer useful for pulse compression of a frequency modulated light pulse.
- [37] I. Walmsley, L. Waxer, and C. Dorrer. The role of dispersion in ultrafast optics. *Review Of Scientific Instruments*, 72(1):1–29, January 2001.
- [38] B. Tatian. Fitting refractive-index data with the Sellmeier dispersion formula. *Appl. Opt.*, 23(24):4477–4485, 1984.
- [39] C.G. Durfee, A.R. Rundquist, C. Backus, S. and Herne, M.M. Murnane, and H.C. Kapteyn. Phase matching of high-order harmonics in hollow waveguides. *Phys. Rev. Lett.*, 83(11):2187–2190, Sep 1999.
- [40] E. Matsubara, K. Yamane, T. Sekikawa, and M. Yamashita. Generation of 2.6 fs optical pulses using induced-phase modulation in a gas-filled hollow fiber. *Journal of the Optical Society of America B*, 24(4):985–989, 2007.
- [41] R.L. Fork, C.V. Shank, C. Hirlimann, R. Yen, and W.J. Tomlinson. Femtosecond white-light continuum pulses. *Optics Letters*, 8(1):1–3, 1983.
- [42] A. M. Weiner. Femtosecond pulse shaping using spatial light modulators. *Review of Scientific Instruments*, 71(5):1929–60, 2000.
- [43] L. Gallmann, G. Steinmeyer, D. H. Sutter, T. Rupp, C. Iaconis, I. A. Walmsley, and U. Keller. Spatially resolved amplitude and phase characterization of femtosecond optical pulses. *Optics Letters*, 26(2):96–98, January 2001.
- [44] Z. Bor. Distortion of femtosecond laser-pulses in lenses and lens systems. *Journal Of Modern Optics*, 35(12):1907–1918, December 1988.
- [45] J. Hartmann. Bemerkungen uber den bau und Die Justirung von Spektrographen. *Z Instrumentenknd*, 20:47, 1900.

-
- [46] B.C. Platt. History and principles of Shack-Hartmann wavefront sensing. *J. Refractive Surg*, 17:573–577, 2001.
- [47] R. N. Smartt and J. Strong. Point-Diffraction Interferometer. *Journal of the Optical Society of America (1917-1983)*, 62:737–+, 1974.
- [48] H. Medeck. Phase-shifting point diffraction interferometer, November 10 1998. US Patent 5,835,217.
- [49] A.K. Aggarwal and S.K. Kaura. Further applications of point diffraction interferometer. *Journal of Optics*, 17(3):135–137, 1986.
- [50] C. P. Hauri, J. Biegert, U. Keller, B. Schaefer, K. Mann, and G. Marowski. Validity of wavefront reconstruction and propagation of ultrabroadband pulses measured with a hartmann-shack sensor. *Optics Letters*, 30(12):1563–1565, June 2005.
- [51] W.J. Bates. A wavefront shearing interferometer. *Proceedings of the Physical Society*, 59(6):940–950–2, 1947.
- [52] M. Murty. The use of a single plane parallel plate as a lateral shearing interferometer with a visible gas laser source. *Appl. Opt*, 3(531):299, 1964.
- [53] C. Dorrer and I. A. Walmsley. Simple linear technique for the measurement of space-time coupling in ultrashort optical pulses. *Optics Letters*, 27(21):1947–1949, November 2002.
- [54] P. Hariharan and D. Sen. Radial shearing interferometer. *J. Sci. Instrum.*, 38:428–432, 1961.
- [55] T. Kohno, D. Matsumoto, T. Yazawa, and Y. Uda. Radial shearing interferometer for in-process measurement of diamond turning. *Optical Engineering*, 39(10):2696–2699, 2000.
- [56] I. A. Walmsley and V. Wong. Characterization of the electric field of ultrashort optical pulses. *Journal Of The Optical Society Of America B-Optical Physics*, 13(11):2453–2463, November 1996.
- [57] L. Cohen. The generalization of the Wiener-Khinchin theorem. *Acoustics, Speech, and Signal Processing, 1998. ICASSP'98. Proceedings of the 1998 IEEE International Conference on*, 3, 1998.
- [58] M. Beck, M. G. Raymer, I. A. Walmsley, and V. Wong. Chronocyclic tomography for measuring the amplitude and phase-structure of optical pulses. *Optics Letters*, 18(23):2041–2043, December 1993.
- [59] C. Dorrer and I. Kang. Simultaneous temporal characterization of telecommunication optical pulses and modulators by use of spectrograms. *Optics Letters*, 27(15):1315–1317, August 2002.
- [60] I. Kang, C. Dorrer, and F. Quochi. Implementation of electro-optic spectral shearing interferometry for ultrashort pulse characterization. *Optics Letters*, 28(22):2264–2266, November 2003.
- [61] C. Dorrer and I. Kang. Highly sensitive direct characterization of femtosecond pulses by electro-optic spectral shearing interferometry. *Optics Letters*, 28(6):477–479, March 2003.
- [62] C. Dorrer and I. Kang. Complete temporal characterization of short optical pulses by simplified chronocyclic tomography. *Optics Letters*, 28(16):1481–1483, August 2003.

- [63] C. Dorrer and I. Kang. Real-time implementation of linear spectrograms for the characterization of high bit-rate optical pulse trains. *Ieee Photonics Technology Letters*, 16(3):858–860, March 2004.
- [64] J. Bromage, C. Dorrer, I. A. Begishev, N. G. Usechak, and J. D. Zuegel. Highly sensitive, single-shot characterization for pulse widths from 0.4 to 85 ps using electro-optic shearing interferometry. *Optics Letters*, 31(23):3523–3525, December 2006.
- [65] P.T. Ho, L.A. Glasser, E.P. Ippen, and H.A. Haus. Picosecond pulse generation with a cw GaAlAs laser diode. *Applied Physics Letters*, 33:241, 1978.
- [66] R. Fork, B. Greene, and C. Shank. Generation of optical pulses shorter than 0.1 psec by colliding pulse mode-locking. *Quantum Electronics, IEEE Journal of*, 17(12 Part 1):2406–2406, 1981.
- [67] T. Andersson and S. T. Eng. Determination of the pulse response from intensity autocorrelation measurements of ultrashort laser pulses. *Optics Communications*, 47:288–290, September 1983.
- [68] J. Peatross and A. Rundquist. Temporal decorrelation of short laser pulses. *J. Opt. Soc. Am. B*, 15:216–222, 1998.
- [69] J.-H. Chung and A. M. Weiner. Ambiguity of ultrashort pulse shapes retrieved from the intensity autocorrelation and the power spectrum. *IEEE J. Select. Topics Quantum Electron.*, 7:656–666, 2001.
- [70] C. Ventalon, J. M. Fraser, J. P. Likforman, D. M. Villeneuve, P. B. Corkum, and M. Joffre. Generation and complete characterization of intense mid-infrared ultrashort pulses. *Journal Of The Optical Society Of America B-Optical Physics*, 23(2):332–340, February 2006.
- [71] J.C.M. Diels, J.J. Fontaine, I.C. Mcmichael, and F. Simoni. Control and measurement of ultrashort pulse shapes(in amplitude and phase) with femtosecond accuracy. *Applied Optics*, 24(9):1270–1282, 1985.
- [72] K. Naganuma, K. Mogi, and H. Yamada. General method for ultrashort light pulse chirp measurement. *IEEE Journal of Quantum Electronics*, 25(6):1225–1233, 1989.
- [73] F. Hache, T.J. Driscoll, M. Cavallari, and G.M. Gale. Measurement of ultrashort pulse durations by interferometric autocorrelation: influence of various parameters. *Applied Optics*, 35(18):3230–3236, 1996.
- [74] T. Feurer, A. Glass, and R. Sauerbrey. Two-photon photoconductivity in SiC photodiodes and its application to autocorrelation measurements of femtosecond optical pulses. *Applied Physics B: Lasers and Optics*, 65(2):295–297, 1997.
- [75] J.K. Ranka, A.L. Gaeta, A. Baltuska, M.S. Pshenichnikov, and D.A. Wiersma. Autocorrelation measurement of 6-fs pulses based on the two-photon-induced photocurrent in a gasp photodiode. *Optics Letters*, 22(17):1344–1346, 1997.
- [76] C. Spielmann, L. Xu, and F. Krausz. Measurement of interferometric autocorrelations: comment. *Applied Optics*, 36(12):2523–2525, 1997.
- [77] H. Mashiko, A. Suda, and K. Midorikawa. All-reflective interferometric autocorrelator for the measurement of ultra-short optical pulses. *Applied Physics B: Lasers and Optics*, 76(5):525–530, 2003.

- [78] I.Z. Kozma, P. Baum, U. Schmidhammer, S. Lochbrunner, and E. Riedle. Compact auto-correlator for the online measurement of tunable 10 femtosecond pulses. *Review of Scientific Instruments*, 75:2323, 2004.
- [79] D. J. Kane and R. Trebino. Characterization of arbitrary femtosecond pulses using frequency-resolved optical gating. *Ieee Journal Of Quantum Electronics*, 29(2):571–579, February 1993.
- [80] K. W. DeLong, R. Trebino, and D. J. Kane. Comparison of ultrashort-pulse frequency-resolved-optical-gating traces for 3 common beam geometries. *Journal Of The Optical Society Of America B-Optical Physics*, 11(9):1595–1608, September 1994.
- [81] L. Gallmann, D. H. Sutter, N. Matuschek, G. Steinmeyer, and U. Keller. Techniques for the characterization of sub-10-fs optical pulses: a comparison. *Applied Physics B-Lasers And Optics*, 70:S67–S75, June 2000.
- [82] R. Trebino, A. Baltuska, M. S. Pshenichnikov, and D. A. Wiersma. Measuring ultrashort pulses in the single-cycle regime: Frequency-resolved optical gating. *Few-Cycle Laser Pulse Generation And Its Applications*, 95:231–263, 2004.
- [83] B. A. Richman, K. W. DeLong, and R. Trebino. Temporal characterization of the stanford mid-ir fel micropulses by frog. *Nuclear Instruments & Methods In Physics Research Section A-Accelerators Spectrometers Detectors And Associated Equipment*, 358(1-3):268–271, April 1995.
- [84] J. M. Dudley, X. Gu, L. Xu, M. Kimmel, E. Zeek, P. O’Shea, R. Trebino, S. Coen, and R. S. Windeler. Cross-correlation frequency resolved optical gating analysis of broadband continuum generation in photonic crystal fiber: simulations and experiments. *Optics Express*, 10(21):1215–1221, October 2002.
- [85] G. Stibenz and G. Steinmeyer. Interferometric frequency-resolved optical gating. *Optics Express*, 13(7):2617–2626, 2005.
- [86] K. W. DeLong, C. L. Ladera, R. Trebino, B. Kohler, and K. R. Wilson. Ultrashort-pulse measurement using noninstantaneous nonlinearities - raman effects in frequency-resolved optical gating. *Optics Letters*, 20(5):486–488, March 1995.
- [87] J. N. Sweetser, D. N. Fittinghoff, and R. Trebino. Transient-grating frequency-resolved optical gating. *Optics Letters*, 22(8):519–521, April 1997.
- [88] A. Kwok, L. Jusinski, M. A. Krumbugel, J. N. Sweetser, D. N. Fittinghoff, and R. Trebino. Frequency-resolved optical gating using cascaded second-order nonlinearities. *Ieee Journal Of Selected Topics In Quantum Electronics*, 4(2):271–277, March 1998.
- [89] D. Lee, S. Akturk, P. Gabolde, and R. Trebino. Experimentally simple, extremely broadband transient-grating frequency-resolved-optical-gating arrangement. *Optics Express*, 15(2):760–766, January 2007.
- [90] K. W. DeLong, R. Trebino, J. Hunter, and W. E. White. Frequency-resolved optical gating with the use of 2nd-harmonic generation. *Journal Of The Optical Society Of America B-Optical Physics*, 11(11):2206–2215, November 1994.
- [91] L. Gallmann, G. Steinmeyer, D. H. Sutter, N. Matuschek, and U. Keller. Collinear type II second-harmonic-generation frequency-resolved optical gating for the characterization of sub-10-fs optical pulses. *Optics Letters*, 25(4):269–271, February 2000.

- [92] P. O'Shea, M. Kimmel, X. Gu, and R. Trebino. Highly simplified device for ultrashort-pulse measurement. *Optics Letters*, 26(12):932–934, June 2001.
- [93] E. Zeek, A. P. Shreenath, P. O'Shea, M. Kimmel, and R. Trebino. Simultaneous automatic calibration and direction-of-time removal frequency-resolved optical gating. *Applied Physics B-Lasers And Optics*, 74:S265–S271, June 2002.
- [94] D. J. Kane and R. Trebino. Single-shot measurement of the intensity and phase of an arbitrary ultrashort pulse by using frequency-resolved optical gating. *Optics Letters*, 18(10):823–825, May 1993.
- [95] S. Akturk, M. Kimmel, P. O'Shea, and R. Trebino. Measuring pulse-front tilt in ultrashort pulses using grenouille. *Optics Express*, 11(5):491–501, March 2003.
- [96] S. Akturk, M. Kimmel, P. O'Shea, and R. Trebino. Measuring spatial chirp in ultrashort pulses using single-shot frequency-resolved optical gating. *Optics Express*, 11(1):68–78, January 2003.
- [97] R. Barakat and G. Newsam. Necessary conditions for a unique solution to two-dimensional phase retrieval. *J. Math. Phys.(NY)*, 23:3190–3193, 1984.
- [98] R.P. Millane. Phase retrieval in crystallography and optics. *J. Opt. Soc. Am. A*, 7(3):394–411, 1990.
- [99] D. J. Kane. Real-time measurement of ultrashort laser pulses using principal component generalized projections. *Ieee Journal Of Selected Topics In Quantum Electronics*, 4(2):278–284, March 1998.
- [100] D. J. Kane. Recent progress toward real-time measurement of ultrashort laser pulses. *Ieee Journal Of Quantum Electronics*, 35(4):421–431, April 1999.
- [101] D. J. Kane, J. Weston, and K. C. J. Chu. Real-time inversion of polarization gate frequency-resolved optical gating spectrograms. *Applied Optics*, 42(6):1140–1144, February 2003.
- [102] K. W. DeLong, D. N. Fittinghoff, R. Trebino, B. Kohler, and K. Wilson. Pulse retrieval in frequency-resolved optical gating based on the method of generalized projections. *Optics Letters*, 19(24):2152–2154, December 1994.
- [103] K. W. DeLong and R. Trebino. Improved ultrashort pulse-retrieval algorithm for frequency-resolved optical gating. *Journal Of The Optical Society Of America A-Optics Image Science And Vision*, 11(9):2429–2437, September 1994.
- [104] M. A. Krumbugel, C. L. Ladera, K. W. DeLong, D. N. Fittinghoff, J. N. Sweetser, and R. Trebino. Direct ultrashort-pulse intensity and phase retrieval by frequency-resolved optical gating and a computational neural network. *Optics Letters*, 21(2):143–145, January 1996.
- [105] L. Gallmann, D. H. Sutter, N. Matuschek, G. Steinmeyer, U. Keller, C. Iaconis, and I. A. Walmsley. Characterization of sub-6-fs optical pulses with spectral phase interferometry for direct electric-field reconstruction. *Optics Letters*, 24(18):1314–1316, September 1999.
- [106] C. Iaconis and I. A. Walmsley. Spectral phase interferometry for direct electric-field reconstruction of ultrashort optical pulses. *Optics Letters*, 23(10):792–794, May 1998.
- [107] C. Iaconis and I. A. Walmsley. Self-referencing spectral interferometry for measuring ultrashort optical pulses. *Ieee Journal Of Quantum Electronics*, 35(4):501–509, April 1999.

-
- [108] I. A. Walmsley. Characterization of ultrashort optical pulses in the few-cycle regime using spectral phase interferometry for direct electric-field reconstruction. *Few-Cycle Laser Pulse Generation And Its Applications*, 95:265–291, 2004.
- [109] G. Stibenz and G. Steinmeyer. Optimizing spectral phase interferometry for direct electric-field reconstruction. *Review Of Scientific Instruments*, 77(7):073105, July 2006.
- [110] C. Dorrer, P. Londero, and I. A. Walmsley. Homodyne detection in spectral phase interferometry for direct electric-field reconstruction. *Optics Letters*, 26(19):1510–1512, October 2001.
- [111] M. Hirasawa, N. Nakagawa, K. Yamamoto, R. Morita, H. Shigekawa, and M. Yamashita. Sensitivity improvement of spectral phase interferometry for direct electric-field reconstruction for the characterization of low-intensity femtosecond pulses. *Applied Physics B: Lasers and Optics*, 74:225–229, 2002.
- [112] C. Dorrer, N. Belabas, J. P. Likforman, and M. Joffre. Spectral resolution and sampling issues in fourier-transform spectral interferometry. *Journal Of The Optical Society Of America B-Optical Physics*, 17(10):1795–1802, October 2000.
- [113] C. Dorrer and I. A. Walmsley. Accuracy criterion for ultrashort pulse characterization techniques: application to spectral phase interferometry for direct electric field reconstruction. *Journal Of The Optical Society Of America B-Optical Physics*, 19(5):1019–1029, May 2002.
- [114] C. Dorrer and I.A. Walmsly. Precision and consistency criteria in spectral phase interferometry for direct electric-field reconstruction. *Journal Of The Optical Society Of America B-Optical Physics*, 19(5):1030–1038, May 2002.
- [115] M. E. Anderson, L. E. E. de Araujo, E. M. Kosik, and I. A. Walmsley. The effects of noise on ultrashort-optical-pulse measurement using spider. *Applied Physics B-Lasers And Optics*, 70:S85–S93, June 2000.
- [116] D. N. Fittinghoff, K. W. DeLong, R. Trebino, and C. L. Ladera. Noise sensitivity in frequency-resolved optical-gating measurements of ultrashort pulses. *Journal Of The Optical Society Of America B-Optical Physics*, 12(10):1955–1967, October 1995.
- [117] K. W. DeLong, D. N. Fittinghoff, and R. Trebino. Practical issues in ultrashort-laser-pulse measurement using frequency-resolved optical gating. *Ieee Journal Of Quantum Electronics*, 32(7):1253–1264, July 1996.
- [118] Z. Y. Wang, E. Zeek, R. Trebino, and P. Kvam. Determining error bars in measurements of ultrashort laser pulses. *Journal Of The Optical Society Of America B-Optical Physics*, 20(11):2400–2405, November 2003.
- [119] Z. Y. Wang, E. Zeek, R. Trebino, and P. Kvam. Beyond error bars: Understanding uncertainty in ultrashort-pulse frequency-resolved-opticalgating measurements in the presence of ambiguity. *Optics Express*, 11(26):3518–3527, December 2003.
- [120] D. Keusters, H. S. Tan, P. O’Shea, E. Zeek, R. Trebino, and W. S. Warren. Relative-phase ambiguities in measurements of ultrashort pulses with well-separated multiple frequency components. *Journal Of The Optical Society Of America B-Optical Physics*, 20(10):2226–2237, October 2003.
- [121] S. Yeremenko, A. Baltuka, M.S. Pshenichnikov, and D.A. Wiersma. The criterion of pulse reconstruction quality based on wigner representation. *Applied Physics B: Lasers and Optics*, 70(0):S109–S117, June 2000.

- [122] A.H. Zewail. Femtochemistry: Atomic-Scale Dynamics of the Chemical Bond Using Ultrafast Lasers(Nobel Lecture). *Angewandte Chemie International Edition*, 39(15):2586–2631, 2000.
- [123] R. Kienberger, E. Goulielmakis, M. Uiberacker, A. Baltuska, V. Yakovlev, F. Bammer, A. Scrinzi, T. Westerwalbesloh, U. Kleineberg, U. Heinzmann, et al. Atomic transient recorder. *Nature*, 427:817–821, 2004.
- [124] V. R. Bhardwaj, S. A. Aseyev, M. Mehendale, G. L. Yudin, D. M. Villeneuve, D. M. Rayner, M. Y. Ivanov, and P. B. Corkum. Few cycle dynamics of multiphoton double ionization. *Physical Review Letters*, 86(16):3522–3525, April 2001.
- [125] P. Haljan, M. Y. Ivanov, and P. B. Corkum. Laser control of electron localization in molecules and double quantum wells. *Laser Physics*, 7(3):839–843, May 1997.
- [126] D. J. Kane, A. J. Taylor, R. Trebino, and K. W. Delong. Single-shot measurement of the intensity and phase of a femtosecond uv laser-pulse with frequency-resolved optical gating. *Optics Letters*, 19(14):1061–1063, July 1994.
- [127] P. Londero, M. E. Anderson, C. Radzewicz, C. Iaconis, and I. A. Walmsley. Measuring ultrafast pulses in the near-ultraviolet using spectral phase interferometry for direct electric field reconstruction. *Journal Of Modern Optics*, 50(2):179–184, January 2003.
- [128] D. H. Sutter, L. Gallmann, N. Matuschek, F. Morier-Genoud, V. Scheuer, G. Angelow, T. Tschudi, G. Steinmeyer, and U. Keller. Sub-6-fs pulses from a sesam-assisted kerr-lens modelocked ti : sapphire laser: at the frontiers of ultrashort pulse generation. *Applied Physics B-Lasers And Optics*, 70:S5–S12, June 2000.
- [129] R. Ell, U. Morgner, F. X. K AÂrtner, J. G. Fujimoto, E. P. Ippen, V. Scheuer, G. Angelow, T. Tschudi, M. J. Lederer, A. Boiko, and B. Luther-Davies. Generation of 5-fs pulses and octave-spanning spectra directly from a ti:sapphire laser. *Opt. Lett.*, 26(6):373–375, 2001.
- [130] M. Nisoli, S. De Silveshi, and O. Svelto. Generation of high-energy 10-fs pulses by a new pulse compression technique. *Lasers and Electro-Optics, 1996. CLEO'96., Summaries of papers presented at the Conference on*, pages 189–190, 1996.
- [131] J. S. Robinson, C. A. Haworth, H. Teng, R. A. Smith, J. P. Marangos, and J. W. G. Tisch. The generation of intense, transform-limited laser pulses with tunable duration from 6 to 30 fs in a differentially pumped hollow fibre. *Applied Physics B-Lasers And Optics*, 85(4):525–529, December 2006.
- [132] C. P. Hauri, W. Kornelis, F. W. Helbing, A. Heinrich, A. Couairon, A. Mysyrowicz, J. Biegert, and U. Keller. Generation of intense, carrier-envelope phase-locked few-cycle laser pulses through filamentation. *Applied Physics B-Lasers And Optics*, 79(6):673–677, October 2004.
- [133] A. Couairon, J. Biegert, C. P. Hauri, W. Kornelis, F. W. Helbing, U. Keller, and A. Mysyrowicz. Self-compression of ultra-short laser pulses down to one optical cycle by filamentation. *Journal Of Modern Optics*, 53(1-2):75–85, January 2006.
- [134] P. Maine, D. Strickland, P. Bado, M. Pessot, and G. Mourou. Generation of ultrahigh peak power pulses by chirped pulseamplification. *Quantum Electronics, IEEE Journal of*, 24(2):398–403, 1988.
- [135] S. Backus, C.G. Durfee III, M.M. Murnane, and H.C. Kapteyn. High power ultrafast lasers. *Review of Scientific Instruments*, 69:1207, 1998.

- [136] A. Shirakawa, I. Sakane, M. Takasaka, and T. Kobayashi. Sub-5-fs visible pulse generation by pulse-front-matched noncollinear optical parametric amplification. *Applied Physics Letters*, 74:2268, 1999.
- [137] T. Kobayashi and A. Baltuska. Sub-5 fs pulse generation from a noncollinear optical parametric amplifier. *Measurement Science and Technology*, 13(11):1671–1682, 2002.
- [138] S. Witte, R. Zinkstok, W. Hogervorst, and K. Eikema. Generation of few-cycle terawatt light pulses using optical parametric chirped pulse amplification. *Opt. Express*, 13(13):4903–4908, 2005.
- [139] C. Teisset, N. Ishii, T. Fuji, T. Metzger, S. Köhler, R. Holzwarth, A. Baltuška, A. Zheltikov, and F. Krausz. Soliton-based pump-seed synchronization for few-cycle OPCPA. *Optics Express*, 13(17):6550–6557, 2005.
- [140] K. Yamane, T. Kito, R. Morita, and M. Yamashita. 2. 8-fs clean single transform-limited optical-pulse generation and characterization. *Springer series in chemical physics*, pages 13–15, 2004.
- [141] S. Luan, M.H.R. Hutchinson, and R. Smith. High dynamic range third-order correlation measurement of picosecond laser pulse shapes. *Measurement Science and Technology*, 4(12):1426–9, 1993.
- [142] D. N. Fittinghoff, J. L. Bowie, J. N. Sweetser, R. T. Jennings, M. A. Krumbugel, K. W. DeLong, R. Trebino, and I. A. Walmsley. Measurement of the intensity and phase of ultraweak, ultrashort laser pulses. *Optics Letters*, 21(12):884–886, June 1996.
- [143] G. Steinmeyer, L. Gallmann, F. W. Helbing, and U. Keller. New directions in sub-10-fs optical pulse generation. *Comptes Rendus De L Academie Des Sciences Serie Iv Physique Astrophysique*, 2(10):1389–1406, December 2001.
- [144] C. Dorrer, E. M. Kosik, and I. A. Walmsley. Spatio-temporal characterization of the electric field of ultrashort optical pulses using two-dimensional shearing interferometry. *Applied Physics B-Lasers And Optics*, 74:S209–S217, June 2002.
- [145] C. Dorrer. Influence of the calibration of the detector on spectral interferometry. *Journal Of The Optical Society Of America B-Optical Physics*, 16(7):1160–1168, July 1999.
- [146] C. Dorrer. Implementation of spectral phase interferometry for direct electric-field reconstruction with a simultaneously recorded reference interferogram. *Optics Letters*, 24(21):1532–1534, November 1999.
- [147] L. Xu, C. Spielmann, A. Poppe, T. Brabec, F. Krausz, and TW Hansch. Route to phase control of ultrashort light pulses. *Opt. Lett*, 21(24):2008–2010, 1996.
- [148] D.J. Jones, S.A. Diddams, J.K. Ranka, A. Stentz, R.S. Windeler, J.L. Hall, and S.T. Cundiff. Carrier-Envelope Phase Control of Femtosecond Mode-Locked Lasers and Direct Optical Frequency Synthesis. *Science*, 288(5466):635–639, 2000.
- [149] A. Baltuška, M. Uiberacker, E. Goulielmakis, R. Kienberger, V.S. Yakovlev, T. Udem, T.W. Hänsch, and F. Krausz. Phase-controlled amplification of few-cycle laser pulses. *Selected Topics in Quantum Electronics, IEEE Journal of*, 9(4):972–989, 2003.
- [150] H. R. Telle, G. Steinmeyer, A. E. Dunlop, J. Stenger, D. H. Sutter, and U. Keller. Carrier-envelope offset phase control: A novel concept for absolute optical frequency measurement and ultrashort pulse generation. *Applied Physics B-Lasers And Optics*, 69(4):327–332, October 1999.

- [151] F. W. Helbing, G. Steinmeyer, and U. Keller. Carrier-envelope offset phase-locking with attosecond timing jitter. *Ieee Journal Of Selected Topics In Quantum Electronics*, 9(4):1030–1040, July 2003.
- [152] P. Dombi, A. Apolonski, C. Lemell, GG Paulus, M. Kakehata, R. Holzwarth, T. Udem, K. Torizuka, J. Burgdörfer, TW Hänsch, et al. Direct measurement and analysis of the carrier-envelope phase in light pulses approaching the single-cycle regime. *Simulation*, 6(39), 2004.
- [153] M. Kress, T. Loeffler, M.D. Thomson, R. Doerner, H. Gimpel, K. Zrost, T. Ergler, R. Moshhammer, U. Morgner, J. Ullrich, et al. Determination of the Carrier-Envelope Phase of Few-Cycle Laser Pulses with Terahertz-Emission Spectroscopy. *Arxiv preprint physics/0603134*, 2006.
- [154] D. Milošević, G. Paulus, and W. Becker. High-order above-threshold ionization with few-cycle pulse: a meter of the absolute phase. *Optics Express*, 11(12):1418–1429, 2003.
- [155] M F Kling, J Rauschenberger, A J Verhoef, E Hasovic, T Uphues, D B Milosevic, H G Muller, and M J J Vrakking. Imaging of carrier-envelope phase effects in above-threshold ionization with intense few-cycle laser fields. *New Journal of Physics*, 10(2):025024, 2008.
- [156] C. A. Haworth, L. E. Chipperfield, J. S. Robinson, P. L. Knight, J. P. Marangos, and J. W. G. Tisch. Half-cycle cutoffs in harmonic spectra and robust carrier-envelope phase retrieval. *Nature Physics*, 3(1):52–57, January 2007.
- [157] E. M. Kosik, A. S. Radunsky, I. A. Walmsley, and C. Dorrer. Interferometric technique for measuring broadband ultrashort pulses at the sampling limit. *Optics Letters*, 30(3):326–328, February 2005.
- [158] J.R. Birge, R. Ell, and F.X. Kärtner. Two-dimensional spectral shearing interferometry for few-cycle pulse characterization. *Optics Letters*, 31(13):2063–2065, 2006.
- [159] P. Baum, S. Lochbrunner, and E. Riedle. Zero-additional-phase SPIDER: full characterization of visible and sub-20-fs ultraviolet pulses. *Optics Letters*, 29(2):210–212, 2004.
- [160] T.W. Hänsch. A proposed sub-femtosecond pulse synthesizer using separate phase-locked laser oscillators. *Optics communications*, 80(1):71–75, 1990.
- [161] N. Zhavoronkov and G. Korn. Generation of Single Intense Short Optical Pulses by Ultrafast Molecular Phase Modulation. *Physical Review Letters*, 88(20):203901, 2002.
- [162] A. Nazarkin, G. Korn, M. Wittmann, and T. Elsaesser. Generation of multiple phase-locked stokes and anti-stokes components in an impulsively excited raman medium. *Phys. Rev. Lett.*, 83(13):2560–2563, Sep 1999.
- [163] M. Wittmann, A. Nazarkin, and G. Korn. fs-pulse synthesis using phase modulation by impulsively excited molecular vibrations. *Phys. Rev. Lett.*, 84(24):5508–5511, Jun 2000.
- [164] B.W. Shore and P.L. Knight. Enhancement of high optical harmonics by excess-photon ionisation. *Journal of Physics B: Atomic, Molecular, and Optical Physics*, 20(2):413–423, 1987.
- [165] A. McPherson, G. Gisbon, H. Jara, U. Johann, and TS Luk. Studies of multiphoton production of vacuum-ultraviolet radiation in the rare gases. *Optical Society of America, Journal, B: Optical Physics*, 4:595–601, 1987.

- [166] M. Ferray, A. L'Huillier, X. F. Li, L. A. Lompre, G. Mainfray, and C. Manus. Multiple-harmonic conversion of 1064-nm radiation in rare-gases. *Journal Of Physics B-Atomic Molecular And Optical Physics*, 21(3):L31–L35, February 1988.
- [167] P. Antoine, A. L'Huillier, and M. Lewenstein. Attosecond pulse trains using high-order harmonics. *Phys. Rev. Lett.*, 77(7):1234–1237, Aug 1996.
- [168] D. A. G. Deacon, L. R. Elias, J. M. J. Madey, G. J. Ramian, H. A. Schwettman, and T. I. Smith. First operation of a free-electron laser. *Phys. Rev. Lett.*, 38(16):892–894, Apr 1977.
- [169] E.L. Saldin, E.A. Schneidmiller, and M.V. Yurkov. Scheme for attophysics experiments at a X-ray SASE FEL. *Optics Communications*, 212(4-6):377–390, 2002.
- [170] E.L. Saldin, E.A. Schneidmiller, and M.V. Yurkov. A new technique to generate 100 GW-level attosecond X-ray pulses from the X-ray SASE FELs. *Optics Communications*, 239(1-3):161–172, 2004.
- [171] G. Penn and AA Zholents. Technique for the Generation of Attosecond X-Ray Pulses Using an FEL. *Particle Accelerator Conference, 2005. PAC 2005. Proceedings of the*, pages 2506–2508, 2005.
- [172] P. Emma, K. Bane, M. Cornacchia, Z. Huang, H. Schlarb, G. Stupakov, and D. Walz. Femtosecond and subfemtosecond x-ray pulses from a self-amplified spontaneous-emission-based free-electron laser. *Physical Review Letters*, 92(7):074801, 2004.
- [173] P. Antoine, D. B. Milosevic, A. L'Huillier, M. B. Gaarde, P. Salieres, and M. Lewenstein. Generation of attosecond pulses in macroscopic media. *Physical Review A*, 56(6):4960–4969, December 1997.
- [174] I.P. Christov, M.M. Murnane, and H.C. Kapteyn. High-Harmonic Generation of Attosecond Pulses in the Single-Cycle Regime. *Physical Review Letters*, 78(7):1251–1254, 1997.
- [175] M. Drescher, M. Hentschel, R. Kienberger, G. Tempea, C. Spielmann, G. A. Reider, P. B. Corkum, and F. Krausz. X-ray pulses approaching the attosecond frontier. *Science*, 291(5510):1923–1927, March 2001.
- [176] G. Sansone, E. Benedetti, F. Calegari, C. Vozzi, L. Avaldi, R. Flammini, L. Poletto, P. Villoresi, C. Altucci, R. Velotta, S. Stagira, S. De Silvestri, and M. Nisoli. Isolated Single-Cycle Attosecond Pulses. *Science*, 314(5798):443–446, 2006.
- [177] K.T. Kim, C.M. Kim, M. Baik, G. Umesh, and C.H. Nam. Single sub-50-attosecond pulse generation from chirp-compensated harmonic radiation using material dispersion. *Physical Review A (Atomic, Molecular, and Optical Physics)*, 69(5):051805, 2004.
- [178] P. B. Corkum. Plasma perspective on strong-field multiphoton ionization. *Physical Review Letters*, 71(13):1994–1997, September 1993.
- [179] M. V. Ammosov, N. B. Delone, and V. P. Krainov. Tunnel ionization of complex atoms and atomic ions in a varying electromagnetic-field. *Zhurnal Eksperimental noi I Teoreticheskoi Fiziki*, 91(6):2008–2013, December 1986.
- [180] P. B. Corkum, N. H. Burnett, and F. Brunel. Above-threshold ionization in the long-wavelength limit. *Physical Review Letters*, 62(11):1259–1262, March 1989.
- [181] S. Augst, D. Strickland, D. D. Meyerhofer, S. L. Chin, and J. H. Eberly. Tunneling ionization of noble gases in a high-intensity laser field. *Phys. Rev. Lett.*, 63(20):2212–2215, Nov 1989.

- [182] F.A. Ilkov, J.E. Decker, and S.L. Chin. Ionization of atoms in the tunnelling regime with experimental evidence using Hg atoms. *Journal of Physics B: Atomic, Molecular, and Optical Physics*, 25(19):4005–4020, 1992.
- [183] P. Dietrich, N. H. Burnett, M. Ivanov, and P. B. Corkum. High-harmonic generation and correlated 2-electron multiphoton ionization with elliptically polarized-light. *Physical Review A*, 50(5):R3585–R3588, November 1994.
- [184] D. N. Fittinghoff, P. R. Bolton, B. Chang, and K. C. Kulander. Observation of nonsequential double ionization of helium with optical tunneling. *Phys. Rev. Lett.*, 69(18):2642–2645, Nov 1992.
- [185] J.L. Krause, K.J. Schafer, and K.C. Kulander. High-order harmonic generation from atoms and ions in the high intensity regime. *Phys. Rev. Lett.*, 68(24):3535–3538, Jun 1992.
- [186] M. Lewenstein, P. Balcou, M. Y. Ivanov, A. L’Huillier, and P. B. Corkum. Theory of high-harmonic generation by low-frequency laser fields. *Physical Review A*, 49(3):2117–2132, March 1994.
- [187] VA Borovikov. *Uniform Stationary Phase Method*. Iet, 1994.
- [188] L. E. Chipperfield, P. L. Knight, J. W. G. Tisch, and J. P. Marangos. Tracking individual electron trajectories in a high harmonic spectrum. *Optics Communications*, 264(2):494–501, August 2006.
- [189] D. Waxman. Numerical determination of bound states without matrix diagonalization. *Journal of Physics A: Mathematical and General*, 31(4):1329–1339, 1998.
- [190] N. Watanabe and M. Tsukada. Fast and stable method for simulating quantum electron dynamics. *Phys. Rev. E*, 62(2):2914–2923, Aug 2000.
- [191] M.D. Feit and J.A. Fleck. Solution of the Schroedinger equation by a spectral method. *Journal of Computational Physics*, 47(3):412–433, 1982.
- [192] K. Burnett, V. C. Reed, J. Cooper, and P. L. Knight. Calculation of the background emitted during high-harmonic generation. *Phys. Rev. A*, 45(5):3347–3349, Mar 1992.
- [193] M. Geissler, G. Tempea, A. Scrinzi, M. Schnürer, F. Krausz, and T. Brabec. Light propagation in field-ionizing media: Extreme nonlinear optics. *Phys. Rev. Lett.*, 83(15):2930–2933, Oct 1999.
- [194] T. Brabec and F. Krausz. Nonlinear optical pulse propagation in the single-cycle regime. *Phys. Rev. Lett.*, 78(17):3282–3285, Apr 1997.
- [195] B.L. Henke, E.M. Gullikson, and J.C. Davis. X-Ray Interactions: Photoabsorption, Scattering, Transmission, and Reflection at $E= 50\text{--}30,000$ eV, $Z= 1\text{--}92$. *Atomic Data and Nuclear Data Tables*, 54(2):181–342, 1993.
- [196] P. Balcou, P. Salieres, A. LHuillier, and M. Lewenstein. Generalized phase-matching conditions for high harmonics: The role of field-gradient forces. *Physical Review A*, 55(4):3204–3210, April 1997.
- [197] M. Lewenstein, P. Salieres, and A. Lhuillier. Phase of the atomic polarization in high-order harmonic-generation. *Physical Review A*, 52(6):4747–4754, December 1995.

- [198] C. Corsi, A. Pirri, E. Sali, A. Tortora, and M. Bellini. Direct Interferometric Measurement of the Atomic Dipole Phase in High-Order Harmonic Generation. *Physical Review Letters*, 97(2):23901, 2006.
- [199] A. Zaïr, M. Holler, A. Guandalini, F. Schapper, J. Biegert, L. Gallmann, U. Keller, A.S. Wyatt, A. Monmayrant, I.A Walmsley, E. Cormier, T. Auguste, J.P. Caumes, and P. Salières. Quantum path interferences in high-order harmonic generation. Submitted to *Physical Review Letters*, 2007.
- [200] E. Goulielmakis, G. Nersisyan, N.A. Papadogiannis, D. Charalambidis, G.D. Tsakiris, and K. Witte. A dispersionless michelson interferometer for the characterization of attosecond pulses. *Applied Physics B: Lasers and Optics*, 74(3):197–206, 2002.
- [201] M. Kovacev, H. Merdji, E. Priori, P. Monchicourt, P. Salieres, B. Carre, J. Norin, A. Persson, A. L’Huillier, and C. G. Wahlstrom. Measurement of autocorrelation in uv-x-ray with high-order harmonics. *Journal De Physique Iv*, 108:101–104, June 2003.
- [202] K.J. Schafer, B. Yang, L.F. DiMauro, and K.C. Kulander. Above threshold ionization beyond the high harmonic cutoff. *Physical Review Letters*, 70(11):1599–1602, 1993.
- [203] F. Quere, Y. Mairesse, and J. Itatani. Temporal characterization of attosecond xuv fields. *Journal Of Modern Optics*, 52(2-3):339–360, January 2005.
- [204] F. Quere, J. Itatani, G. L. Yudin, and P. B. Corkum. Attosecond spectral shearing interferometry. *Physical Review Letters*, 90(7):073902, February 2003.
- [205] J. Itatani, F. Quere, G. L. Yudin, M. Y. Ivanov, F. Krausz, and P. B. Corkum. Attosecond streak camera. *Physical Review Letters*, 88(17):173903, April 2002.
- [206] H.G. Muller. Reconstruction of attosecond harmonic beating by interference of two-photon transitions. *Applied Physics B: Lasers and Optics*, 74:17–21, 2002.
- [207] Y. Mairesse and F. Quere. Frequency-resolved optical gating for complete reconstruction of attosecond bursts. *Physical Review A*, 71(1):011401, January 2005.
- [208] P. Antoine, M. Gaarde, P. Salieres, B. Carre, A. L’Huillier, and M. Lewenstein. On the phase of atomic polarization in harmonic generation processes. *Multiphoton Processes 1996*, 154(154):142–159, 1997.
- [209] P. Salieres, L. Le Deroff, T. Auguste, P. Monot, P. d’Oliveira, D. Campo, J. F. Hergott, H. Merdji, and B. Carre. Frequency-domain interferometry in the xuv with high-order harmonics. *Physical Review Letters*, 83(26):5483–5486, December 1999.
- [210] H. Merdji, J. F. Hergott, M. Kovacev, E. Priori, P. Salieres, and B. Carre. Manipulating intense xuv coherent light in the temporal domain. *Laser And Particle Beams*, 22(3):275–278, September 2004.
- [211] J. F. Hergott, P. Salieres, H. Merdji, L. Le Deroff, B. Carre, T. Auguste, P. Monot, P. D’Oliveira, D. Descamps, J. Norin, C. Lynga, A. L’Huillier, C. G. Wahlstrom, M. Bellini, and S. Huller. Xuv interferometry using high-order harmonics: Application to plasma diagnostics. *Laser And Particle Beams*, 19(1):35–40, March 2001.
- [212] J. F. Hergott, T. Auguste, P. Salieres, L. Le Deroff, P. Monot, P. d’Oliveira, D. Campo, H. Merdji, and B. Carre. Application of frequency-domain interferometry in the extreme-ultraviolet range by use of high-order harmonics. *Journal Of The Optical Society Of America B-Optical Physics*, 20(1):171–181, January 2003.

- [213] D. Descamps, C. Lynga, J. Norin, A. L'Huillier, C. G. Wahlstrom, J. F. Hergott, H. Merdji, P. Salieres, M. Bellini, and T. W. Hansch. Extreme ultraviolet interferometry measurements with high-order harmonics. *Optics Letters*, 25(2):135–137, January 2000.
- [214] T. Ditmire, E. T. Gumbrell, R. A. Smith, J. W. G. Tisch, D. D. Meyerhofer, and M. H. R. Hutchinson. Spatial coherence measurement of soft x-ray radiation produced by high order harmonic generation. *Phys. Rev. Lett.*, 77(23):4756–4759, Dec 1996.
- [215] L. Le Deroff, P. Salieres, B. Carre, D. Joyeux, D. Phalippou, P. Monot, P. D'Oliveira, T. Auguste, H. Merdji, and J. F. Hergott. Temporal and spatial coherence of high order harmonics. *Laser Physics*, 10(1):294–302, January 2000.
- [216] H. Merdji, P. Salieres, L. Le Deroff, J. F. Hergott, B. Carre, D. Joyeux, D. Descamps, J. Norin, C. Lynga, A. L'Huillier, C. G. Wahlstrom, M. Bellini, and S. Huller. Coherence properties of high-order harmonics: Application to high-density laser-plasma diagnostic. *Laser And Particle Beams*, 18(3):495–502, September 2000.
- [217] B. Carre, P. Salieres, L. Le Deroff, H. Merdji, J. F. Hergott, T. Auguste, P. Monot, P. D'Oliveira, and D. Joyeux. Coherence properties of high harmonics and their applications. *Journal De Physique Iv*, 11(PR2):355–364, July 2001.
- [218] M. Kovacev, S. V. Fomichev, E. Priori, Y. Mairesse, H. Merdji, P. Monchicourt, P. Breger, J. Norin, A. Persson, A. L'Huillier, C. G. Wahlstrom, B. Carre, and P. Salieres. Extreme ultraviolet fourier-transform spectroscopy with high order harmonics. *Physical Review Letters*, 95(22):223903, November 2005.
- [219] T. Kita, T. Harada, N. Nakano, and H. Kuroda. Mechanically ruled aberration-corrected concave gratings for a flat-field grazing-incidence spectrograph. *Applied Optics*, 22(4):512–513, 1983.
- [220] T. Sekikawa, T. Katsura, S. Miura, and S. Watanabe. Measurement of the intensity-dependent atomic dipole phase of a high harmonic by frequency-resolved optical gating. *Phys. Rev. Lett.*, 88(19):193902, Apr 2002.
- [221] S. A. Aseyev, Y. Ni, L. J. Frasinski, H. G. Muller, and M. J. J. Vrakking. Attosecond angle-resolved photoelectron spectroscopy. *Phys. Rev. Lett.*, 91(22):223902, Nov 2003.
- [222] Y. Mairesse, O. Gobert, P. Breger, H. Merdji, P. Meynadier, P. Monchicourt, M. Perdrix, P. Salieres, and B. Carre. High harmonic xuv spectral phase interferometry for direct electric-field reconstruction. *Physical Review Letters*, 94(17):173903, May 2005.
- [223] J. Mauritsson, P. Johnsson, R. Lopez-Martens, K. Varju, W. Kornelis, J. Biegert, U. Keller, M. B. Gaarde, K. J. Schafer, and A. L'Huillier. Measurement and control of the frequency chirp rate of high-order harmonic pulses. *Physical Review A*, 70(2):021801, August 2004.
- [224] M. Guizar-Sicairos and J. C. Gutierrez-Vega. Computation of quasi-discrete hankel transforms of integer order for propagating optical wave fields. *Journal Of The Optical Society Of America A-Optics Image Science And Vision*, 21(1):53–58, January 2004.