High-resolution Interferometric Diagnostics for Ultrashort Pulses
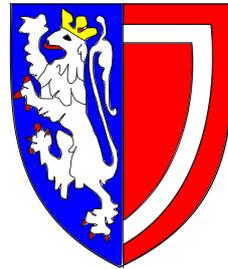
Dane Austin

Balliol College and Department of Physics

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Michaelmas term 2010

Abstract

I present several new methods for the characterisation of ultrashort pulses using interferometry.

A generalisation of the concatenation algorithm for spectral shearing interferometry enables interferograms taken at multiple shears to be combined. This improves the precision of the reconstructed phase in the presence of detector noise, and enables the relative phase between disjoint spectral components to be obtained without decreasing the spectral resolution. The algorithm is applied to experimental data from two different implementations of spectral shearing interferometry for ultrashort optical pulses. In one, the shears are acquired sequentially, and in the other they are acquired simultaneously.

I develop a form of spatio-temporal ultrashort pulse characterisation which performs both spatial and spectral shearing interferometry simultaneously. It requires a similar geometrical setup to common implementations of spectral phase interferometry for direct electric-field reconstruction, but provides complete amplitude and phase characterisation in time and one spatial dimension.

I develop the theory of lateral shearing interferometry for spectrally resolved wavefront sensing of extended ultraviolet and soft x-ray pulses generated using high-harmonic generation. A comprehensive set of wavefront measurements of harmonics 13–25 in Krypton show good agreement with theory, validating the technique.

I propose and numerically demonstrate quantum-path interferometry mediated by a weak control field for high harmonic generation. This is a general technique for measuring the amplitude and relative phases of each contributing quantum path. The control field perturbatively modulates the phase of each path. The differing sensitivity of each path to the parameters of the control field allows their contributions to be distinguished from one another.

Informal summary for nonexperts

Physics exposes some unlikely similarities. The unpleasant beating produced when the same note is struck on a pair of badly tuned musical instruments, the exquisite sensitivity of the LIGO gravitational wave detector, capable of detecting movements of less than 1000$^{\text{th}}$ of the diameter of a proton, and the shimmering colours of a thin layer of oil on water, are all manifestations of *interference*: the pattern of crests and troughs created when two or more waves combine. Whilst LIGO is one of the most sophisticated instruments ever devised, simpler examples of interferometry are everywhere: the rate of the beating between out-of-tune instruments provides the difference in their pitch, thus enabling musicians to tune them. The light reflected from the surface of the oil arrives at the eye earlier than light reflected from the interface between the oil and the water, resulting in an enhancement of certain colours and suppression of others.

In recent years, one important application of interferometry has been in the measurement of ultrashort laser pulses — bursts of light less than one picosecond (10^{-12} s) in duration. These pulses have become important tools in science and industry because of their ability to illuminate and induce processes on the femtosecond (10^{-15} s) timescale — considerably faster than the switching speed of modern electronics, despite all of its sophistication. However, this begs the question: how does one observe ultrashort pulses, when they are faster than any detector? In particular, the uncertainty principle requires that an ultrashort pulse be composed of a range of different colours, or frequencies. A complete characterisation of such a pulse therefore requires knowledge of the arrival time of each frequency component. Interferometry provides one answer: when two ultrashort pulses arrive at different times on a detector, interference enhances certain frequencies and suppresses others - similar to the oil-on-water effect. The spacing of these peaks and troughs tells us the difference in the arrival times of the frequency components, in the same way that the beating between two instruments indicates the difference in their pitch.

This thesis offers several relevant developments. One result concerns shearing interferometry, in which a pulse is interfered with a displaced or *sheared* version of itself. This enables a complete characterisation without needing a previously characterised reference pulse. However, one long-standing limitation of shearing interferometry has been its inability to precisely measure pulses

with many peaks and troughs in their intensity profile in time and space. This thesis resolves this problem by combining measurements taken with several different shears. These measurements may be acquired simultaneously or in sequence, and both approaches are demonstrated experimentally. This opens the door to characterizing extremely complex pulses. Another challenge in ultrashort pulse diagnostics is measuring the field as a function of space and frequency, rather than frequency alone. This is important because far from being perfectly round "balls of light", ultrashort pulses often have a complicated space-time structure which severely affects their applicability. Shearing interferometry offers a natural solution to this problem because one can combine two measurements: one with a shear in space, and the other with a shear in frequency. This thesis presents a novel experimental arrangement for acquiring both datasets on a single device.

Two other results of the thesis concerned *high-harmonic generation (HHG)* — the emission of ultrashort bursts of extended ultraviolet (XUV) radiation by the interaction of an intense laser field with matter. The characterization of HHG is a challenging task because of the complexity of the emitted radiation and the scarcity of optical elements suitable for XUV frequencies. This thesis presents a method for *frequency-resolved wavefront characterisation* of the harmonics — measuring not only where the light is, but where it is going.

Besides its potential as a novel source of XUV and soft x-ray radiation, HHG is a complex physical process which reveals the properties of matter at atomic length and time scales — a combination unattainable using other methods. In HHG, the ultra-intense laser field literally tears an electron away from an atom, before changing direction and smashing the electron and atom back together again. This process happens many times during the ultrashort pulse, and each of these *quantum paths* contribute to the emitted radiation. Just like light and sound waves, the contributions from the quantum paths interfere, which ultimately manifests in enhancement and suppression of different frequencies of the emitted radiation. With many quantum paths present, as is usually the case, the interference pattern becomes very complicated and it is difficult to isolate the contribution from each one — rather like trying to tune many different instruments at the same time. This thesis proposes an experimental method for disentangling the contributions by gently perturbing the high-harmonic generation process with a weak control field.

Acknowledgements

My supervisor Ian Walmsley has been a continual source of ideas and knowledge. He has created a group which is diverse, intellectual, eclectic, and hugely entertaining to be a part of. Despite his many responsibilities as administrator, diplomat, fund-raiser and safety guru he had retained a mastery of his subject which has guided all of the results of this dissertation, and provided much encouragement. For all of this he deserves great thanks.

The largest contributor to the experimental knowhow that I acquired throughout this degree is Tobias Witting, my laboratory partner for the first three years, and an honour and a pleasure to work with. Ultrafast metrologists Adam Wyatt, Antoine Monmayrant, and Andrea Schiavi also deserve special thanks for many stimulating discussions and practical assistance. One could not ask for a capable and amusing bunch of lab- and laser-mates than David McCabe, Duncan England, Phil Bustard. I must also thank the rest of the ultrafast group for a myriad of scientific matters as well as being such a enjoyable team of colleagues and friends.

Thanks go to Christopher Arrell, Felix Frank, Luke Chipperfield, Joe Robinson, and John Tisch at Imperial College for providing the high-harmonic generation setup and ensuring that visits to London were always great fun. Thanks also to Sébastien Weber, Ayhan Tajalli, Béatrice Chatel, and lately David McCabe for an enjoyable and continuing collaboration, and being great hosts in Toulouse.

Warm thanks to my effortlessly superior Balliol posse as well as the Oxford University Athletic Club crew. Silja Baller has also been incredibly supportive throughout the last year of this degree. Thanks also to Noni Austin for sisterly support from afar.

The greatest thanks of all go to my parents. Any wisdom contained herein is ultimately due to them; all errors are my own.

Author's contribution

The concept of using multiple shears, developed in chapter 3, has been suggested in the literature previously. The SEA-SPIDER and SEA-CAR-SPIDER described in chapter 4 were constructed with Tobias Witting (TW). Data were acquired using the former with his assistance. The concept of lateral-shearing interferometry for high-harmonic generation was developed in the Walmsley group before my arrival. The high-harmonic generation experiments, described in chapter 7, were also carried out with TW. The interferometer used in those experiments was constructed with the assistance of TW and Adam Wyatt. The high-harmonic generation equipment was constructed and maintained by the members of the Blackett Laboratory Laser Consortium at Imperial College. All other work presented in this dissertation is my own.

Contents

Publications

The following publications arose out of work connected to this dissertation:

D. R. Austin, T. Witting, C. A. Arrell, F. Frank, A. S. Wyatt, J. P. Marangos, J. W. G. Tisch and I. A. Walmsley, "Lateral shearing interferometry of high-harmonic wavefronts," submitted to Opt. Lett. 3 February 2011.

D. J. McCabe, D. R. Austin, A. Tajalli, S. Weber, I. A. Walmsley and B. Chatel, "Space–time coupling of shaped ultrafast ultraviolet pulses from an acousto-optic programmable dispersive filter," J. Opt. Soc. Am. B. **28**, 58–64 (2011).

D. R. Austin, T. Witting, and I. A. Walmsley, "Resolution of the relative phase ambiguity in spectral shearing interferometry of ultrashort pulses," Opt. Lett. **35**, 1971–1973 (2010).

D. R. Austin, T. Witting, A. S. Wyatt, and I. A. Walmsley, "Measuring sub-Planck structural analogues in chronocyclic phase space," Opt. Commun. **283**, 855–859 (2010). Quo vadis Quantum Optics?

T. Witting, D. R. Austin, and I. A. Walmsley, "Ultrashort pulse characterization by spectral shearing interferometry with spatially chirped ancillae," Opt. Express **17**, 18983–18994 (2009).

D. R. Austin, T. Witting, and I. A. Walmsley, "High precision self-referenced phase retrieval of complex pulses with multiple-shearing spectral interferometry," J. Opt. Soc. Am. B **26**, 1818–1830 (2009).

D. R. Austin, T. Witting, and I. A. Walmsley, "Broadband astigmatism-free Czerny-Turner imaging spectrometer using spherical mirrors," Appl. Opt. **48**, 3846–3853 (2009).

T. Witting, D. R. Austin, and I. A. Walmsley, "Improved ancilla preparation in spectral shearing interferometry for accurate ultrafast pulse characterization," Opt. Lett. **34**, 881–883 (2009).

List of acronyms

ADK Ammosov, Delone, and Kraïnov

CAR-SPIDER chirped-arrangement for SPIDER

CCD charge-coupled device

CEO carrier-envelope offset

CPA chirped-pulse amplifier

CRAB complete reconstruction of attosecond bursts

DFT discrete Fourier transform

FROG frequency-resolved optical gating

FTSI Fourier-transform spectral interferometry

FTSSI Fourier-transform spatio-spectral interferometry

FWHM full-width at half-maximum

LSI lateral shearing interferometry

HHG high-harmonic generation

HOT-SPIDER homodyne optical technique for SPIDER

LX-SPIDER long-crystal SPIDER

MCP micro-channel plate

MZI Mach-Zehnder interferometer

OPA optical parametric amplifier

PACER probing attosecond dynamics by chirp encoded recollision

PCGPA principal component generalized projections algorithm

PG-FROG polarisation-gating FROG

QPI quantum-path interferometry

RABBIT reconstruction of attosecond beating by interference of two-photon transitions

RMS root-mean-square

SD-FROG self-diffraction FROG

SEA-SPIDER spatially encoded arrangement SPIDER

SEA-CAR-SPIDER spatially encoded arrangement CAR-SPIDER

SEA-TADPOLE spatially encoded arrangement TADPOLE

SFA strong field approximation

SHG second-harmonic generation

SHG-FROG second-harmonic generation FROG

SNR signal-to-noise ratio

SPA stationary-phase approximation

SPIDER spectral phase interferometry for direct electric-field reconstruction

SPIRIT spectral interferometry resolved in time

SSI spectral shearing interferometry

ST-SPIDER space-time SPIDER

SWORD spectral wavefront optical reconstruction by diffraction

TADPOLE temporal analysis by dispersing a pair of e-fields

TOF time of flight

TDSE time-dependent Schrödinger equation

XFROG crosscorrelation FROG

XUV extended ultraviolet

ZAP-SPIDER zero-additional phase SPIDER

Definitions and symbols

All Fourier transforms in this dissertation are unitary and use angular frequency. The sign of the transform is usually determined by the physical context: a time-to-frequency transform has exponential factor $e^{i\omega t}$, whilst a space-to-wavenumber transform has exponential factor e^{-ixk_x}. In other cases, the sign is explicitly stated. When two quantities are related by an integral transform, such as the Fourier transform or the Hankel transform, they are distinguished by their argument e.g. $E(t)$ vs $E(\omega)$. Specifically, the Fourier transforms of the analytic electric field are

$$E(\omega) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} E(t) e^{i\omega t} \, \mathrm{d}t \qquad E(t) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} E(\omega) e^{-it\omega} \, \mathrm{d}\omega$$

$$E(k_x) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} E(x) e^{-ik_x x} \, \mathrm{d}x \qquad E(x) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} E(k_x) e^{ixk_x} \, \mathrm{d}k_x.$$

The zeroth-order Hankel transforms are

$$E(k_{\mathrm{T}}) = \int_{0}^{\infty} E(r) J_0(k_{\mathrm{T}} r) r \, \mathrm{d}r \qquad E(r) = \int_{0}^{\infty} E(k_{\mathrm{T}}) J_0(k_{\mathrm{T}} r) k_{\mathrm{T}} \, \mathrm{d}k_{\mathrm{T}}.$$

Several notational conventions are adopted throughout this dissertation. Boldface symbols denote vectorial quantities. Where both an italic roman and calligraphic version of a symbol exist they refer to the analytic (complex-valued) and physical (real-valued) signals respectively. The following table illustrates these two conventions:

	Real	Analytic
Scalar	$\mathscr{E}$	E
Vector	$\boldsymbol{\mathscr{E}}$	$\mathbf{E}$

1 Introduction

Ultrafast optics is a burgeoning area of science and technology, driven by the unique ability of *ultrashort pulses* to probe, control and modify matter on subpicosecond timescales. Through their brevity, these pulses enable unprecedented temporal resolution in the study of dynamic processes such as chemical reactions. The high intensity achievable by concentrating even a modest pulse energy into such a fleeting instant also provides access to optical nonlinearities, ranging from second harmonic generation all the way up to relativistic optics. Such nonlinearities reveal new properties of matter whilst dramatically increasing the capabilities of light-matter interactions. Ultrashort pulses also enable the highly space-time localised delivery of energy, enabling matter — including biological tissue — to be altered in unique ways.

The cost of the attractive properties of ultrashort pulses is that they are challenging to create, manipulate, and measure — three essential tasks in the development of successful ultrafast science and technology. This dissertation deals with the third of these problems. The principal difficulty in measuring the properties of ultrashort pulses is that both the shape of the pulse and the electric field oscillations within it occur on a significantly shorter timescale than the response of any electronic detector. In particular, detectors generally exhibit integrating square-law behaviour, and thus respond to the energy of incoming radiation, but not its phase. *Ultrashort pulse characterisation* is the large subfield of ultrashort optics which has arisen in search of practical solutions to this fundamental problem of performing a subpicosecond resolution measurement using a detector with significantly slower temporal resolution.

In parallel with the advances of the pulse sources themselves, great progress has been made in metrology. For example, measurement of the temporal electric field (that is, ignoring, averaging or otherwise neglecting the spatial variations) of simple ultrashort pulses at optical and infrared frequencies with durations ranging from a 10–100 fs is now routine, with commercial devices available. Close to the cutting edge, several groups have measured the temporal electric field of soft x-ray pulses of durations around 100 as. The demands of complex ultrashort pulse sources and experiments are driving metrology towards extremes of pulse brevity and complexity, applicability to different wavelengths, and the ability to gather more information about a pulse, such as its

profile in space as well as time.

This dissertation advances ultrashort pulse characterisation in several ways. *Shearing interferometry*, in which the phase is recovered by interfering a pulse with a displaced, or *sheared*, replica of itself, plays a key role in two of the results. This principle underlies a broad class of ultrashort pulse characterisation devices. I develop a generalisation to shearing interferometry which allows many different shears to be combined, providing greater precision and recovering properties of a pulse which are unobtainable using a single shear. I also develop a shearing interferometer capable of performing spatial and spectral shears simultaneously, enabling simple measurement of the spatio-temporal structure of optical pulses.

High-harmonic generation is a process by which intense laser pulses at optical and infrared wavelengths may be converted to extended ultraviolet and soft x-ray radiation, and is the subject of intensive study because of its potential as a short-wavelength light source of unprecedented brevity. However, adapting characterisation techniques designed for optical wavelengths to high-harmonic radiation is not straightforward because of the limited range of optical components available at the shorter wavelengths. I describe an adaption of shearing interferometry for frequency-resolved *wavefront sensing* of high-harmonic radiation. A comprehensive set of measurements show good agreement with theory.

Besides its potential as a light source, high-harmonic generation is a rich physical process which exposes the properties of matter at atomic length scales and ultrashort timescales. The electric field of an intense laser pulse ionises a target atom or molecule and accelerates the electron away, before reversing its direction and bringing the electron back with great velocity. In the resulting collision, the kinetic energy of the electron is released as a short-wavelength photon. Whilst this simple classical picture explains some of the properties of the generated radiation, there are in fact several possible trajectories that the electron could follow which lead to the same emitted photon. Such is the peculiarity of quantum mechanics that in fact all of these trajectories happen simultaneously; the emitted radiation corresponds to their coherent superposition. I propose and develop the theory of a technique for measuring the intensity and phase of the radiation corresponding to each trajectory, thus decomposing this quantum superposition. A "numerical

experiment" supports the plausibility of the method.

This dissertation has two main themes. First, all of the results involve *interferometry* — inferring the properties of waves through the intensity of their superposition. Second, the results increase the *resolution* of ultrashort pulse metrology, where the word is taken in the general sense of meaning "the number of independent pieces of information that are gained by a measurement".

This dissertation has two parts. The first, consisting of chapters 2–5, concerns ultrashort pulses at optical and infrared wavelengths. Chapter 2 presents the background material for this part. Chapter 3 develops the theory of multiple shearing interferometry and then chapter 4 presents some experimental demonstrations. Chapter 5 describes and demonstrates the spatial and spectral shearing interferometer. Chapters 6–8 comprise the second part of this dissertation, which concerns high-harmonic generation (HHG). Chapter 6 presents the background for this part. Chapter 7 presents the theory of lateral shearing interferometry for high-harmonic generation and a proof-of-principle experiment. Chapter 8 introduces quantum-path interferometry mediated by a control field and describes the supporting numerical experiment. Chapter 9 presents a summary, conclusions, and outlook for both parts.

2 Background

This dissertation presents new interferometric methods for characterising ultrashort pulses, an essential task in ultrafast science. This chapter provides some motivation for pulse metrology specifically and ultrafast science in general. It presents the current state of relevant technologies and also introduces the necessary theory and mathematical formalism to appreciate the results.

A distinction is made between optical and infrared pulses, produced either directly by a laser or by low-order mixing processes, and extreme ultraviolet or soft x-ray pulses produced by high-harmonic generation (HHG). This is because fewer operations are available for pulses produced by HHG, and so generic approaches which can be adapted throughout the optical and infrared spectrum require substantial modification. In addition, one of the results pertaining to HHG, quantum-path interferometry (chapter 8), is dependent for its operation on a model of the generation process itself. This necessitates a presentation of the theory of HHG. By contrast, the optical and infrared techniques presented in part one of this dissertation are model-independent. They therefore apply in principle to pulses produced by any method. Therefore, the background material and new results concerning HHG are deferred to the latter part of this dissertation, chapters 6–8.

This chapter is organised as follows: section 2.1 introduces the field of ultrafast science and presents general arguments and historical evidence for the importance of pulse metrology. Section 2.2 presents the mathematical formalism used to describe ultrashort pulses. These preliminaries enable a detailed examination of existing methods for measurement of the time-dependence of the electric field of optical and infrared pulses in section 2.3. Existing methods for extending this to the full spatio-temporal field profile are described in section 2.4. Finally, section 2.5 provides an outline of the main body of the dissertation.

2.1 Metrology, ultrashort pulses, and ultrafast science

Ultrashort pulses are bursts of electromagnetic radiation with durations below one picosecond — amongst the briefest of all man-made events. Since light travels only 0.3 mm in this period, and the beam diameter is typically several millimetres, an ultrashort pulse may be envisaged as a

thin disc of light. By contrast, a flash of lightning, a commonly encountered optical "pulse", has a typical duration of 200 ms, during which light travels 60,000 km, or approximately one and a half times the circumference of the earth [1].

The restriction property of the Fourier transform requires an ultrashort pulse to span a range of frequencies inversely proportional to its duration, whilst electromagnetism prohibits a propagating pulse from having a zero-frequency (DC) component. Together, these constraints imply that the minimum duration of an electromagnetic pulse is the inverse of its central frequency. Durations below one picosecond therefore require frequencies above one terahertz. An ever-expanding array of source technologies allows ultrashort pulses to be produced at a variety of frequencies ranging from terahertz up to the soft x-ray regime. Furthermore, many sources are capable of producing *few-cycle* pulses, which contain only a few oscillations of the electromagnetic field and approach the aforementioned limit.

Ultrashort pulses occupy a unique place in science and technology because of several related characteristics. They offer the highest possible temporal resolution for the study of ultrafast processes. They correspondingly possess great bandwidth — for example, few-cycle pulses in the visible range span such a broad spectrum of frequencies that they appear white. This holds great potential for optical communications, and enables ultrashort pulses to simultaneously couple many different energy levels of an atomic or molecular system. Ultrashort pulses also offer high spatial resolution. Their longitudinal extent is given by their duration multiplied by the speed of light, whilst a high degree of spatial coherence enables focusing down to a diffraction-limited transverse area. Finally, ultrashort pulses offer the potential for high intensities at moderate pulse energies and average powers. These high intensities enable access to optical nonlinearities, in which the anharmonic motion of the driven electrons reveals properties of matter which are inaccessible at lower intensities. The delivery of spatially and temporally concentrated energy also enables matter to be modified in unique ways. The next section elaborates one some of the applications of ultrashort pulses.

2.1.1 Applications of ultrashort pulses

Perhaps the most scientifically revolutionary application of ultrashort laser pulses is *ultrafast spectroscopy*. Here, the ultrashort pulse illuminates a rapidly evolving sample, playing a similar role to the shutter in high-speed photography of such scenes as bullets slicing through cards. Of particular importance is the *pump-probe* arrangement, where an ultrafast process is initiated by a pump pulse. The state of the system is sampled some time later by a probe pulse. By repeating the operation with varying pump-probe delays, the evolution of the system can be traced with a resolution equal to the duration of the probe. Examples are excited electron and phonon states in semiconductors [2], ultrafast solvation dynamics [3] and transition states in chemical reactions [4].

Ultrashort pulses provide high intensities with only a modest deposition of energy per unit area upon the target and with only modest pulse energies. This avoids damage due to heating and obviates the need for large, complex and expensive lasers. Such intensities provide access to optical nonlinearities, in which the driven electron motion is anharmonic, and scattered light may be of a different wavelength to the incident. Nonlinear interactions provide many advantages over linear ones for interrogating the properties of matter. For example, nonlinear optical spectroscopy [5, 6], of which the pump-probe arrangement is an example, is widely used in chemistry to elucidate the dynamics of chemical reactions. In fact, the counterpart to studying such processes is controlling them [7, 8], and here again ultrashort pulses play a key role, particularly shaped pulses with specially tailored temporal structures. Another important application is nonlinear microscopy [9], which has high axial resolution because generation of the signal is constrained to the focal plane, the only place where the intensity is above the nonlinear threshold. It is therefore well suited to volumetric imaging of biological samples. Furthermore, the plethora of different fluorescence spectra which may be monitored [10] enables great chemical specificity.

Ultrashort pulse micromachining can either remove material or change its properties in novel ways [11]. Tight focusing, ultrashort duration and highly nonlinear absorption mean that energy deposition is highly localised in space and time. Specific sites can be altered, or vapourised, without substantial heating of the bulk material. This enables clean cuts and the machining of mi-

crometer scale features, and is finding applications in both technology and medicine.

Ultrashort pulses are generally produced as a periodic train, which can be engineered to exhibit exceptional timing stability. The Fourier spectrum of such a train consists of a series of precisely defined lines. Such *frequency combs* [12] have enabled revolutionary measurements of optical frequencies with a relative uncertainty of less than 10^{-15} [13, 14].

2.1.2 Need for ultrashort pulse characterisation

The measurement of ultrashort pulses, the topic of this dissertation, is a challenging task because the response time of electronics is tens of picoseconds at best, meaning that the temporal features of an ultrashort pulse, if measured directly, are blurred beyond recovery. However, ultrashort pulse characterisation is an essential component of ultrafast science and technology, because i) it illuminates the physics of ultrashort pulse generation and interactions with matter, and ii) even when the physics is understood in principle, the complexity of most ultrashort pulse generation methods combined with the sensitivity of the experiments and applications means that characterisation is a practical necessity if the entire system is to be understood, controlled and repeatable.

Metrology is essential for ultrashort source development. An estimate of the pulse duration is necessary to establish that the source is performing as desired, and gives a simple parameter to optimise. A more detailed description of the field provides insight into the workings of the source thus revealing the obstacles to further pulse shortening, enabling further improvement. Therefore, ultrashort sources and appropriate characterisation technology have progressed in tandem, each motivating and supporting the development of the other.

2.1.2.1 Metrology and the quest for shorter pulses

The first generation of sub-100 ps lasers, using Nd:glass, Nd:YAG, or ruby as gain media, stimulated the development of the picosecond streak camera, which records the temporal intensity $I(t)$ of the pulse [15]. The streak camera works by applying a time-dependent displacement to photoelectrons produced by the pulse using a rapidly varying electric field. An image of the photoelectrons therefore has a time-to-space mapping. Streak camera measurements revealed distributions of pulse durations and energies, as well as satellite pulses and substructure [16, 17], leading to im-

proved understanding.

The highest temporal resolution of a streak camera is approximately 2 ps [18]. The introduction of passively mode-locked dye lasers led to the first subpicosecond pulses [19]. Measurement of such pulses required the use of *nonlinear autocorrelation* methods, with the general feature that the detected signal is the result of the pulse nonlinearly interacting with a delayed replica of itself. By scanning the delay, an estimate of the pulse's properties is obtained. The intrinsic temporal resolution of such methods is the response time of optical nonlinearities, which in many cases is below one femtosecond. Nonlinear autocorrelations had previously been used in addition to streak cameras for longer pulse durations [20, 21], but became an essential technique in verifying and guiding improvements in laser technology below one picosecond. Besides estimating the pulse duration, intensity autocorrelations suggested a hyperbolic-secant pulse shape, consistent with theoretical models [22], and subtle dynamical effects [23] analogous to those observed in optical fibres.

Whilst nonlinear methods enabled subpicosecond time resolution, linear spectral measurements continued to provide useful information. Since the spectrum conveys the presence or absence of particular wavelengths, but not their arrival times, knowledge of the spectrum can only provide a lower bound on the pulse duration via the uncertainty property of the Fourier transform. Treacy [24, 25] observed that the duration of Nd:glass laser pulses, as inferred from nonlinear autocorrelations, was significantly longer than the minimum duration inferred from the spectrum — the pulses were not *Fourier transform-limited*. The pulses could be compressed to their transform limited duration by applying wavelength-dependent delay, and were thus diagnosed to possess *chirp* — different arrival times for different wavelengths. This was caused by *dispersion* — the wavelength-dependence of the group velocity of light — in the elements of the laser cavity [26, 27]. The shorter a pulse, the broader its spectrum, and the more sensitive to dispersion it becomes. Therefore, since its diagnosis in the 1960s, measurement of pulse chirp has been crucial in the quest for shorter pulse durations. For example, similar methods to Treacy [24] were used to measure and compensate for the chirp of passively mode-locked dye lasers [28]. In the colliding pulse technique [29, 30], which accessed the sub-100 fs regime for the first time, chirp measure-

ments [31, 32] led to the deliberate introduction of dispersive elements into the cavity leading to greater pulse compression [33, 34], eventually down to 27 fs [35, 36].

While ultrashort dye lasers were enjoying their zenith in the 1980s, investigations into solid-state materials revealed the potential for even greater bandwidths, potentially supporting pulses as short as 1 fs [37]. From a number of different technologies and pulse generation mechanisms, the Ti:Sapph oscillator emerged as a leading contender. By 1993, pulse durations down to about 10 fs [38–41] were being produced. However, in this regime of pulse duration, the enormous bandwidths possess a chirp which cannot be accurately characterised by a simple linear dependence of arrival time on wavelength. Instead, several orders of a Taylor polynomial expansion must be considered for good characterisation and compensation [42]. The combination of the spectrum with the simple nonlinear autocorrelations, as used previously, did not determine these higher order terms. This motivated the development of *complete characterisation methods*, capable of retrieving all information about the pulse up to a few well understood, and hopefully insignificant, ambiguities. These measurements guided improvements to the designs. More sensitive autocorrelation methods [43] guided the introduction of specially-designed chirped mirrors to compensate for intra-cavity dispersion [44] with corresponding pulse durations of 8 fs. However, the first complete characterisation methods such as frequency-resolved optical gating (FROG) [45] and spectral phase interferometry for direct electric-field reconstruction (SPIDER) [46], returned the chirp to all orders [47, 48, 48–50], assisting with design improvements which pushed pulse durations down to around 5 fs [51, 52].

2.1.2.2   Metrology and the manipulation of ultrashort pulses

Besides the drive towards shorter pulses, another development in ultrafast science is the ability to amplify, compress and shape ultrashort pulses. Metrology also plays a crucial role in these areas.

Pulse amplification systems are often limited by the peak intensity, which causes nonlinear distortions of the pulse or induces damage to the materials. Chirped pulse amplification [53, 54] circumvents this by stretching the pulse to reduce its peak intensity, then performing the amplification, then compressing the pulse by undoing the chirp. The stretching and compression steps

are usually accomplished by using a dispersive element to direct the light along a wavelength-dependent path of varying length. These two operations must be carefully balanced to produce transform-limited output. Accurate metrology enables such systems to be aligned and optimised [55, 56]. Because the wavelengths are temporarily separated in space, such devices can easily cause coupled spatio-temporal distortions, motivating the use of metrology with both spatial and temporal resolution [57–59].

Nonlinear compressors, which increase the bandwidth of the pulse, are often used to reduce the pulse duration below that which is emitted by the laser [42, 60]. This is particularly true after amplification, which often reduces the bandwidth of the pulse. However, nonlinear compressors usually involve several interacting processes, are sensitive to the details of the initial pulse shape, and intrinsically involve highly broadband pulses. Characterisation of the output is therefore essential [61–63]. A related topic is the study of nonlinear propagation effects, both in waveguides [64] and in bulk media [65–67] which possess rich dynamics that have analogies in several branches of physics [68].

Shaped pulses, with complex and deliberately crafted profiles, are used for coherent control of physical processes [7, 8]. Often, because of uncertainties in the shaping apparatus, knowledge of the shaped pulse profile can only be obtained through metrology. In addition, many common pulse shaper designs cause space-time distortions which may have deleterious effects on coherent control experiments [69, 70], necessitating careful monitoring [71, 72].

2.2   Formal introduction to ultrashort pulses

This section presents the mathematical description of ultrashort pulses used throughout this thesis.

2.2.1   Positive-frequency representation

Many situations in ultrafast optics exhibit *linearity* — for example, the actions of mirrors, lenses, waveguides, and even many active media such as low-gain amplifiers are all linear. Furthermore, many nonlinear effects are relatively weak and may be treated as perturbative extensions to linear

models. Therefore, the mathematics of linear systems play a fundamental role in ultrafast optics. In this regard, one key concept is the analytic representation of a real-valued signal, in which the negative-going frequencies of the signal are discarded since they are superfluous due to the Hermitian symmetry of their Fourier transform. Mathematically, if $\mathscr{E}(t)$ is the real-valued signal and $E(t)$ its analytic representation, then

$$\mathscr{E}(t) = E(t) + \text{c.c.} = 2\operatorname{Re} E(t) \tag{2.1}$$

and

$$E(t) = \frac{\mathscr{E}(t) + i\operatorname{HT}[\mathscr{E}(t)]}{2} \tag{2.2}$$

where HT is the Hilbert transform. One may move back and forth between the real and analytic representations with no loss of information. Further insight is given in the frequency domain. The Fourier transform of (2.2) is

$$E(\omega) = \mathscr{E}(\omega)u(\omega) \tag{2.3}$$

where $u(\omega)$ is the step function, showing that the analytic signal is simply the real signal with all negative frequencies set to zero. The Fourier transform of (2.1) is

$$\mathscr{E}(\omega) = E(\omega) + E^*(-\omega) \tag{2.4}$$

showing that the real signal is obtained by adding to the analytic signal a conjugated and inverted copy of itself. Throughout this thesis, I shall use $\mathscr{E}$ to denote the real-valued (physical) electric field and E for the analytic representation.

2.2.2 Definition of optical pulses

This section shows that an optical pulse is well defined by the time-dependent value of its electric field at all points on a plane perpendicular to its direction of propagation. This definition forms the starting point for the mathematical treatment of ultrafast optics.

One takes as given the fact that any solution $\mathbf{E}(\mathbf{x}, t)$ to Maxwell's equations in vacuum can be

written as an integral over plane waves

$$\mathbf{E}(\mathbf{x}, t) = \int \boldsymbol{\epsilon}(\mathbf{k}) e^{i(\mathbf{k} \cdot \mathbf{x} - \omega t)} \, \mathrm{d}\mathbf{k} \tag{2.5}$$

where $\mathbf{x} = (x, y, z)$ is position, $\omega = ck$ is the vacuum dispersion relation, and $\boldsymbol{\epsilon}(\mathbf{k})$ is the vectorial amplitude of the plane wave of wavenumber $\mathbf{k} = (k_x, k_y, k_z)$, satisfying the transversality condition $\boldsymbol{\epsilon}(\mathbf{k}) \cdot \mathbf{k} = 0$. The magnetic field is uniquely determined from the electric field via Faraday's law. I shall show that the electric field at reference plane $z = 0$ uniquely defines $\boldsymbol{\epsilon}(\mathbf{k})$, and hence the field at all points in space. One expresses the field at $z = 0$ as a Fourier transform in frequency and along the transverse axes x and y:

$$\mathbf{E}(x, y, 0, t) = \frac{1}{(2\pi)^{3/2}} \int \int \int \mathbf{E}(k_x, k_y, z = 0, \omega) e^{i(k_x x + k_y y - \omega t)} \, \mathrm{d}k_x \, \mathrm{d}k_y \, \mathrm{d}\omega. \tag{2.6}$$

Equating (2.6) with (2.5) at $z = 0$ gives

$$\boldsymbol{\epsilon}(\mathbf{k}) = \frac{1}{(2\pi)^{3/2}} \frac{c k_z}{\omega} \mathbf{E}(k_x, k_y, z = 0, \omega), \tag{2.7}$$

uniquely defining the amplitude of all plane waves with $k_z > 0$ where

$$k_z = \sqrt{\left(\frac{\omega}{c}\right)^2 - \left(k_x^2 + k_y^2\right)}. \tag{2.8}$$

Using (2.5) and (2.7) one can also write the field at any plane in terms of the field at $z = 0$:

$$\mathbf{E}(k_x, k_y, z, \omega) = E(k_x, k_y, z = 0, \omega) e^{i k_z z}. \tag{2.9}$$

Since the field at $z = 0$ uniquely defines the field at all points in space (in a vacuum), one can now drop the z-dependence and consider the pulse properties in this reference plane.

Equation (2.7) gives a direct connection between the spatial and temporal Fourier transform of the electric field in the plane $z = 0$, and the plane wave solutions of the vacuum. The in-plane

Fourier component (k_x, k_y, ω) corresponds to a plane wave of wavenumber $\mathbf{k}$, with k_z defined by (2.8). This provides an intuitive interpretation of the field in either the wavenumber-temporal domain $E(k_x, k_y, t)$ or the wavenumber-spectral domain $E(k_x, k_y, \omega)$: wavenumber components (k_x, k_y) correspond to a wave propagating at an angle, projected onto the xz-plane, of $\theta_x = \sin^{-1}(k_x/k)$ to the z-axis (and likewise for y). Throughout this dissertation, a *paraxial approximation* $\sin \theta_x \approx \theta_x$ applies, giving a linear relationship between wavenumber and plane wave angle $\theta_x = k_x/k$. In a homogeneous dielectric this becomes $\theta_x = k_x/(n(\omega)k)$.

2.2.3 Describing optical pulses

2.2.3.1 Time, frequency and time-frequency picture

In many cases the spatial dependence of an ultrashort pulse is ignored because either a) only a single point in the transverse plane is being considered, b) it is sufficient to consider an average over the transverse plane, or c) the pulse does not possess *space-time coupling* and can be written as a product of temporal and spatial factors. In these cases, the pulse may then be described in either of the time or frequency domains, related bijectively by the Fourier transform, or using combined time-frequency distributions.

Time domain. The *analytic temporal field* $E(t)$ defines the actual electric field by (2.1). The power per unit area or *temporal intensity* is given by $I(t) = 2|E(t)|^2/Z_0$ in SI units, where Z_0 is the vacuum impedance. For many-cycle pulses it may be interpreted as the cycle-averaged intensity. For shorter durations the definition of a single cycle loses its precision.

Writing the temporal analytic field in polar notation, as $E(t) = A(t)e^{i\phi(t)}$ defines the *temporal phase $\phi(t)$* and *temporal amplitude $A(t) > 0$*. The *carrier-envelope offset* (CEO) phase is the phase of the field oscillations at the intensity peak of the pulse. If without loss of generality one places the peak of the pulse at the time origin, then the CEO phase is $\phi(0)$. The CEO phase is largely irrelevant for pulses longer than a few optical cycles, but for shorter pulses undergoing highly nonlinear interactions, it becomes crucial.

Time-derivatives of the phase are commonly encountered. The *instantaneous frequency* is defined as $-\partial \phi(t)/\partial t$. In some cases, consideration of the instantaneous frequency provides a very

intuitive description of a pulse: one common example is *linear temporal chirp* in which the instantaneous frequency varies linearly with time. A pulse is said to have *positive temporal chirp* if the instantaneous frequency increases with time i.e. the colour changes from red to blue; the opposite situation is *negative temporal chirp*. The sign and degree of temporal chirp is characterised by the second order temporal phase $\partial^2 \phi(t)/\partial t^2$.

Sometimes, it is useful to remove the rapidly varying *carrier frequency* ω_c from the analytic electric field. One writes $\bar{E}(t) = E(t) \exp(i \omega_c t)$ as the *baseband* analytic field which varies on a much slower timescale than the underlying carrier frequency oscillations.

Frequency-domain. Fourier transforming in time, one obtains the spectral analytic field $E(\omega)$. The energy per unit frequency per unit area, or *spectral intensity*, is given by $\tilde{I}(\omega) = 2|E(\omega)|^2/Z_0$ in SI units.

Using polar notation, $E(\omega) = A(\omega) e^{i\phi(\omega)}$, defines the *spectral amplitude* $A(\omega)$ and the *spectral phase* $\phi(\omega)$. As with the time domain, derivatives of the spectral phase carry intuitive significance. The *group delay* $\partial \phi(\omega)/\partial \omega$ is the arrival time of a particular frequency. Analogous to the instantaneous frequency, the frequency-dependence of the group delay provides a useful interpretation of *spectral chirp* — if lower frequencies arrive before higher ones, the pulse has *positive spectral chirp*. The sign and magnitude of the spectral chirp is given by *group-delay dispersion* $\partial^2 \phi(\omega)/\partial \omega^2$. The spectral phase is commonly written as a Taylor expansion at a nominal centre frequency ω_c:

$$\phi(\omega) = \sum_{n=0}^{\infty} \frac{\partial^n \phi}{\partial \omega^n}\bigg|_{\omega_c} \frac{(\omega - \omega_c)^n}{n!}. \tag{2.10}$$

If the second order and higher terms are zero, then the pulse is said to be *transform-limited*. A transform-limited pulse has the shortest root-mean-square (RMS) duration possible for a given spectral intensity. The zeroth-order or *absolute phase* $\phi(\omega_c)$ plays a role in determining the CEO phase by affecting the optical carrier, but does not affect the temporal intensity. As such, the comments concerning the CEO phase above also apply to the absolute phase. A change in the absolute phase is reflected exactly in the CEO phase, and vice-versa.

It is often useful to work in terms of the baseband frequency $\bar{\omega} = \omega - \omega_c$. The *baseband spectral*

amplitude $\bar{E}(\bar{\omega}) = E(\bar{\omega} + \omega_c)$ *is a shifted version of the spectral amplitude.*

Time-frequency domain. Together, the amplitude and phase, in either of the time- or frequency-domains, uniquely define the pulse. However they do not necessarily provide an intuitive representation, especially for complex pulses. In such cases, *time-frequency functions*, which represent the pulse on the two-dimensional (t, ω) domain, are often helpful, despite the fact that the conjugate nature of time and frequency causes fundamental resolution and blurring issues related to the uncertainty principle. Often, not just a single pulse but the fluctuations of an ensemble of pulses may also be described in the time-frequency domain, in which case a *time-frequency distribution* is used.

Many time-frequency distributions may be defined. They may be categorised according to their theoretical significance, their intuitiveness and the practicality of measuring them. The Wigner function [73–75] has great theoretical significance because the equations of motion in optics and quantum mechanics can be re-formulated to describe its evolution. It also has the property that interference fringes appear between well separated subpulses in the time-frequency domain. For an ensemble, the fringes in the Wigner distribution represent the coherence between the subpulses, whilst for a single pulse they show the relative phase of the subpulses. This is not necessarily intuitive information, and the fringes become very complex when many subpulses are present. There are currently few means of directly observing the Wigner function.

The spectrogram is a gated Fourier transform:

$$B(\tau, \omega) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} g(t - \tau) E(t) e^{i\omega t} \, \mathrm{d}t. \tag{2.11}$$

It does not feature the interference fringes of the Wigner function, and can be measured directly for ultrashort pulses.

2.2.3.2   Introducing spatial dependence

The electric field of a pulse propagating along the z-axis is written as $E(x, y, t)$ or $E(x, y, \omega)$. Some useful definitions are the fluence, which is the energy deposited per unit area

$$F(x, y) = \int_{-\infty}^{\infty} I(x, y, t) \, \mathrm{d}t \tag{2.12}$$

(and equivalently in the frequency domain via Parseval's theorem), the pulse power

$$P(t) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} I(x, y, t) \, \mathrm{d}x \, \mathrm{d}y \tag{2.13}$$

and the spectrum

$$S(\omega) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \tilde{I}(x, y, \omega) \, \mathrm{d}x \, \mathrm{d}y. \tag{2.14}$$

A complete description of the pulse is greatly simplified if can be written as a product of a temporal/spectral factor, and a spatial factor:

$$E(x, y, t) = E(t)E(x, y). \tag{2.15}$$

Such *space-time factorable* pulses are often desirable from both a theoretical and practical perspective. They are easier to understand and measure and also can be focused to the smallest spot sizes. For such pulses, all of the temporal and spectral properties defined in section 2.2.3.1 are spatially invariant. Furthermore, all spatial properties of the beam are temporally and spectrally invariant, and may be described using the same terms that apply to CW lasers, such as beam centroid, beam waist, wavefront curvature and M^2 beam quality.

The definition of space-time factorable implies a choice of co-ordinate system. In paraxial optics, a small rotation θ about the y-axis, as effected by a rotation $\theta/2$ of a mirror, transforms a pulse as $E'(x, y, t) = E(x, y, t - \theta x/c)$, or, in the frequency domain, $E'(x, y, \omega) = E(x, y, \omega) \exp(i\theta \omega x/c)$. It therefore transforms a factorable pulse into a nonfactorable one according to the above definition. An appropriate choice of co-ordinates is generally assumed in discussions of space-time

factorability.

2.2.3.3 Space-time coupled pulses

If the pulse cannot be written in the form (2.15) for any rotation θ, it is said to possess *space-time coupling*, and must be described in a combined space/wavenumber and time/frequency picture. For this discussion, coupling with only one spatial dimension, x, is considered. Fourier transformation may be performed along either or both of the x and the t axes, resulting in four possible domains in which the two-dimensional field of an ultrashort pulse may be described: spatio-temporal (x, t), spatio-spectral (x, ω), wavenumber-temporal (k_x, t), and wavenumber-spectral (k_x, ω).

For cylindrically symmetric profiles $E(r)$, the appropriate conjugate representation is given by the Hankel transform $E(k_\mathrm{T})$, where k_T is the transverse wavenumber.

There are a vast range of possible space-time couplings. One useful starting point is to consider the action of the lowest order couplings of the phase. There is one term for each of the four domains. For example, in the spatio-spectral domain, the lowest-order coupling term is $\exp(i\alpha x\bar{\omega})$ where α is the coupling coefficient. Since the expansions are around the centre frequency, I use baseband signals for this section. The physical significance of each coupling can be found by performing a Fourier transform to the other domains. For example,

$$\frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} \bar{E}(x, \bar{\omega}) e^{\alpha x \bar{\omega}} e^{-i\bar{\omega}t} \, \mathrm{d}\bar{\omega} = \bar{E}(x, t - \alpha x) \tag{2.16}$$

$$\frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} \bar{E}(x, \bar{\omega}) e^{\alpha x \bar{\omega}} e^{-i x k_x} \, \mathrm{d}x = \bar{E}(k_x - \alpha\bar{\omega}, \bar{\omega}) \tag{2.17}$$

showing that a spatio-spectral coupling phase causes a shear along the time-axis in the spatio-temporal domain, and a shear along the wavenumber axis in the wavenumber-frequency domain. A second Fourier transform produces a more complicated expression. Table 2.1 shows the mathematical forms of each of the four couplings. Assuming that the original pulse $\bar{E}(x, t)$ is space-time factorable, each phase coupling in a given domain has a simple interpretation in the two other domains related by a single Fourier transform. These are given in table 2.1.

$(x, \bar{\omega})$	(x, t)	$(k_x, \bar{\omega})$	(k_x, t)
$\bar{E}(x, \bar{\omega}) \exp(i\alpha x \bar{\omega})$ spatio-spectral phase	$\bar{E}(x, t - \alpha x)$ position-dependent arrival time	$\bar{E}(k_x - \alpha \bar{\omega}, \bar{\omega})$ frequency-dependent tilt	—
$\bar{E}(x, \bar{\omega} + \beta x)$ position-dependent centre frequency	$\bar{E}(x, t) \exp(i\beta x t)$ spatio-temporal phase	—	$\bar{E}(k_x - \beta t, t)$ time-dependent tilt
$\bar{E}(x + \gamma \bar{\omega}, \bar{\omega})$ frequency-dependent beam centroid	—	$\bar{E}(k_x, \bar{\omega}) \exp(i\gamma k_x \bar{\omega})$ wavenumber-spectral phase	$\bar{E}(k_x, t - \gamma k_x)$ wavenumber-dependent arrival time
—	$\bar{E}(x + \delta t, t)$ time-dependent beam centroid	$\bar{E}(k_x, \bar{\omega} + \delta k_x)$ wavefront-dependent centre frequency	$\bar{E}(k_x, t) \exp(i\delta k_x t)$ wavenumber-temporal phase

Table 2.1: Lowest-order space-time phase couplings. Each row corresponds to a type of coupling, and each column corresponds to one of the four domains.

Three types of space-time coupling are particularly common. *Pulse-front tilt* refers to any situation in which the arrival time is position dependent. Both a spatio-spectral phase (row 1 of table 2.1) and a wavenumber-temporal phase (row 4) produce pulse-front tilt. *Angular dispersion* refers to any situation in which the tilt of the beam depends on frequency. Again, both a spatio-spectral and a wavenumber-temporal phase coupling produce angular dispersion, and any operation which induces pulse-front tilt on a space-time factorable pulse will result in angular dispersion, and vice versa [58, 76, 77]. However, for general pulses, one does not necessarily imply the other [78]. The spatio-spectral phase coefficient is related to the angular dispersion $\partial\theta/\partial\omega$ by $\alpha = (n\omega_c/c)(\partial\theta/\partial\omega)$.

Note the distinction between a tilt and a pulse-front tilt. Applying a tilt (e.g. a small rotation) to a pulse rotates the phase and group fronts by equal amounts, whereas the essence of pulse-front tilt is that the phase front (the plane of constant phase at the centre frequency) and group fronts (the plane of constant group delay at the centre frequency) are different. Mathematically, a tilt is represented as a spatio-spectral phase of $\theta\omega x/c$ whereas the simplest form of pulse-front tilt is represented by $\alpha\bar{\omega}x$ i.e. in a tilt, the phase is proportional to the passband frequency whereas

for a pulse-front tilt it is proportional to the baseband frequency. In the wavenumber-spectral domain, a tilt is represented as $E(k_x - \theta \omega/c, \omega)$, showing that the displacement in wavenumber is proportional to the passband frequency and hence represents a constant angle θ. On the other hand the spatio-spectral phase is represented as $\bar{E}(k_x - \alpha\bar{\omega}, \bar{\omega})$, showing that the frequencies are spread around the axis. This difference introduces a subtlety into the measurements of chapter 5.

Spatial chirp refers to any situation in which the central frequency of the pulse varies with position. A spatio-temporal phase (row 2) and a wavenumber-spectral phase (row 3) both produce spatial chirp. Furthermore, a pulse with angular dispersion will acquire spatial chirp after propagation.

For concreteness, I shall now discuss some physical operations which introduce space-time coupling.

Non-normal incidence at a dielectric boundary produces angular dispersion via Snell's law. If the central frequency of a pulse, is incident with angle θ_1 on a boundary between media with refractive indices n_1 and n_2, then the resulting angular dispersion is

$$\frac{\partial \theta_2}{\partial \omega} = \frac{\frac{\partial n_1}{\partial \omega}\sin\theta_1 - \frac{\partial n_2}{\partial \omega}\sin\theta_2}{n_2 \cos\theta_2}. \tag{2.18}$$

In most situations, this effect is quite small. For example, for fused silica at 800 nm, $\partial n/\partial \omega \approx 0.005\,\mathrm{fs}$, so a 5 fs pulse incident at 45deg experiences an angular spread of ≈ 0.1 mrad across its full-width at half-maximum bandwidth due to refraction. The equivalent pulse-front tilt is $\approx 2.1\,\mathrm{fs/mm}$. Furthermore, for a parallel glass plate, the angular dispersion is completely reversed upon exiting, so that only a small spatial chirp acquired during propagation results. On the other hand, dispersive prisms typical produce an angular dispersion of $\partial \theta/\partial \omega = 0.01\,\mathrm{fs}$, giving a pulse-front tilt of $100\,\mathrm{fs/mm}$, significant for femtosecond pulses.

Frequency-dependent mode area. The focus of a space-time factorable pulse produced by a perfect lens has a frequency-dependent mode area because the mapping $x' = f k_x/k$ between wavenumber and position x' in the focus is frequency-dependent. This dependence becomes significant for broadband pulses. Few-cycle oscillators can also produce such an effect [79].

Self-action via nonlinearities. In unguided propagation, or multi-mode propagation in waveguides, self-action nonlinearities cause space-time coupling. For example, propagation of a distance L in a material with intensity-dependent refractive index $n = n_0 + n_2 I$ causes a spatio-temporal *self-phase modulation* $E(x, t, L) = E(x, t, 0) \exp[i n_2 k_0 L I(x, t, 0)]$. In the frequency domain this causes a position-dependent spectral width and chirp, and in the wavenumber domain the distribution is broadened and chirped in a time-dependent manner. Combined with diffraction, the pulse experiences time-dependent self-focusing (for $n_2 > 0$) or defocusing ($n_2 < 0$). At higher intensities, ionisation produces a free-electron plasma which reduces the refractive index, inducing a spatio-temporal phase. A position-dependent blue-shift of the beam results. Combined with diffraction, the result is time-dependent defocusing.

With these mathematical preliminaries established, the next section introduces ultrashort pulse characterisation, the topic of this dissertation.

2.3   Introduction to ultrashort pulse metrology

This section introduces the main concepts of ultrashort pulse metrology. It focuses on characterisation in the temporal and spectral domains; spatial and coupled spatio-temporal measurements are deferred until section 2.4.

2.3.1   Opening comments

There exists a great diversity of ultrashort pulse characterisation techniques. Nonetheless, some general comments and categorisations apply almost universally, and assist with a comprehensible discussion.

One basic distinction is that between self-referenced and nonself-referenced methods. Nonself-referenced methods require a previously well-characterised pulse; self-referenced methods do not require this assistance. In general, a well-characterised reference is a valuable asset in ultrashort metrology.

Characterisation methods vary in the degree of information they provide about the pulse. Methods that return the electric field $E(t)$ up to a few well-defined *ambiguities* are said to be

complete. Methods which only return limited information on the pulse, or equivalently possess a massive and usually infinite number of ambiguities represent *partial* characterisations. There is some slackness in this distinction — a partial method combined with some assumptions is often sufficient to define the pulse with a precision that qualifies as complete. For example, the spectrum is a partial characterisation, but if the pulse is known to be transform-limited then the characterisation can be said to be complete.

The topic of ambiguities is often nebulous and occasionally controversial. An almost universal ambiguity of self-referenced methods is that of the absolute phase — this can only be determined using ultrabroadband pulses such that different orders from a nonlinear interaction overlap in the frequency domain. One example is f-$2f$ interferometry used in CEO phase stabilisation [80]. Another common ambiguity is the arrival time of the pulse. *Exact ambiguities*, such as those of the absolute phase and arrival time, can be proven exactly given the mathematical form of the acquired signal. The standard diagnosis of such an ambiguity is the substitution of two different pulses into the mathematical expression for the signal; if the results are identical, then an ambiguity is present.

Other *approximate* ambiguities are harder to classify because they may depend on experimental issues: the signal-to-noise ratio (SNR), discretisation of the data and the algorithm which is being used. Examples of this class are the *relative phase ambiguities* [81] between components of the pulse which are separate in the time and/or frequency domains. For small separations, the relative phase is usually returned correctly, but as the separation increases the estimated value becomes increasingly susceptible to detector noise. A more general issue is the precision with which the spectral phase is retrieved in the low-intensity tails of the spectrum. The notion of an ambiguity then folds into more general questions of *accuracy* and *precision*. For a given measurement protocol, accuracy quantifies the similarity between the retrieved pulse and the physical pulse, whilst precision quantifies the spread of pulses retrieved from repeated measurements of the same physical pulse. Things which commonly affect accuracy and precision include detector noise, vibrations and instability of the device, incorrect calibrations, violation of the assumptions required of a pulse for a procedure to be valid, misalignment and general user error.

The next section introduces the key obstacles to ultrashort pulse characterisation.

2.3.2 The challenge of ultrashort pulse metrology

The first difficulty of ultrashort pulse metrology arises from the fact that even state-of-the-art detectors are too slow to respond to details of the pulse profile. Instead, they produce a time-integrated signal proportional to the incident fluence or energy. One must therefore perform operations on the pulse which cause the necessary information to be encoded into the fluence which falls onto the detector.

Several widely used definitions are worth restating at this point. An operation f on the field is *linear* if it is additive and homogeneous: that is if $f[E_1 + E_2] = f[E_1] + f[E_2]$ and $f[\alpha E_1] = \alpha E_1$, where E_1 and E_2 are arbitrary fields and α an arbitrary scalar. An operation is *stationary* along a particular dimension, such as time or space, if the result is invariant under a shift in the coordinate system along that dimension. For example, a *time-stationary*[1] operation obeys the property $f[E(t + \tau)](t) = f[E(t)](t + \tau)$.

The next difficulty arises from the fact that most common operations — reflection from mirrors, diffraction from gratings or prisms, focusing through lenses — are linear and time-stationary. Synthesising time-nonstationary elements using electronics is challenging on an ultrafast timescale because subpicosecond response times push the limits of state-of-the-art modulators. A fundamental result of the theory of linear time-stationary systems is that any such system is completely characterised by an impulse response $H(t)$ in the time domain, or its Fourier transform the transfer function $H(\omega)$ in the frequency domain. Any assembly of linear time-stationary operations is itself a linear time-stationary operation; therefore the signal produced by any sequence of passive linear optical elements applied to a field $E(\omega)$ followed by a photodetector can be reduced to the form

$$B = \int_{-\infty}^{\infty} |H(\omega)E(\omega)|^2 \, \mathrm{d}\omega, \tag{2.19}$$

which is independent of the spectral phase. The best that can be hoped for is to produce a tunable

---

[1]A synonym for time-stationary, commonly used outside the field of ultrashort pulse characterisation, is *time invariant*.

spectrally selective response $H(\omega) = \delta(\omega - \omega_0)$; this represents an ideal *spectrometer*. Ultrashort pulse characterisation therefore requires at least one nonlinear or time-nonstationary element.

Because a spectrometer can be produced using linear time-stationary elements, knowledge of the spectrum is generally taken as a given when discussing ultrashort pulse characterisation methods, and the problem is reduced to finding the unknown spectral phase $\phi(\omega)$. Similarly, most characterisation methods can be viewed as a means of encoding the profile of the unknown pulse into the spectrum.

A variety of physical processes are used as nonlinear or time-nonstationary elements. The most common is a nonlinear optical response, almost universally of second or third order. Since nonlinear optics arises from the anharmonic motion of bound electrons, it benefits from an effectively instantaneous response, at least for pulse durations above one femtosecond. The disadvantage is a lack of sensitivity. Electro-optic modulators provide a stronger signal, but struggle with timescales below 200 fs. For extended ultraviolet (XUV) and soft x-ray pulses, which are capable of photoionising atoms, electron momentum streaking can be performed with an ancillary infrared pulse, and will be described in chapter 6. Regardless of the particular method, the process is often represented as a temporal gate at time τ

$$E'(t) = E(t)G(t - \tau). \tag{2.20}$$

It must be emphasised that this is in many cases a mathematical idealisation. For example the limited phase-matching bandwidth of nonlinear optical processes means that the signal involves a convolution, or smoothing operation. Furthermore, in many cases the gate is derived from the pulse itself, and in some cases is even equal to it, so that the assignment of the roles of pulse and gate are arbitrary. Nonetheless, the "gated-pulse" picture is useful for describing many characterisation techniques.

This concludes the generic discussion of ultrafast metrology. The rest of this section covers specific concepts and implementations. The discussion is organised in rough order of complexity, starting with the incomplete or nonself-referenced methods and moving through to complete,
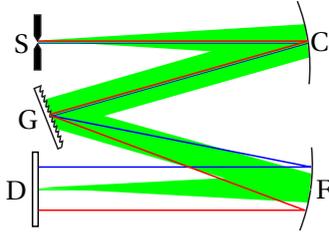
Figure 2.1: Czerny-Turner spectrometer; the symbols are defined in the text.
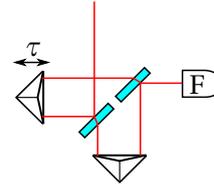
Figure 2.2: Apparatus for measuring the field autocorrelation using a balanced interferometer.

self-referenced methods.

2.3.3 Linear time-stationary measurements

Measurements performed using linear time-stationary elements and slow integrating detectors are a fundamental building block in most forms of ultrashort pulse characterisation.

2.3.3.1 Single-pulse spectral intensity measurements

All measurements of a single pulse using linear time-stationary elements and slow integrating detectors are of the form (2.19) — they are some approximation to the spectral intensity.

A typical Czerny-Turner [82] spectrometer design is shown in Fig. 2.1. Light spreads out from the input slit S and is collimated by the curved mirror C. It is angularly dispersed by diffraction grating G, and during propagation from the grating to the curved focusing mirror F the wavelengths become spread. The focused beam falls upon the spatially resolving detector D with a space-to-frequency mapping which enables $\tilde{I}(\omega)$ to be inferred.

An alternative to the spectrometer is measurement of the field autocorrelation, defined as

$$B_{\mathrm{AC1}}(\tau) = \int_{-\infty}^{\infty} E(t)E^*(t-\tau)\,\mathrm{d}t. \tag{2.21}$$

This can be obtained using a balanced interferometer as shown in Fig. 2.2. The signal obtained on the photodetector is

$$\int_{-\infty}^{\infty} |E(t)+E(t-\tau)|^2\,\mathrm{d}t \tag{2.22}$$

from which the field autocorrelation may be inferred using the fact that $E(\omega)$ contains only positive frequencies. Elementary properties of the Fourier transform show that the spectrum is the Fourier transform of the field autocorrelation.

2.3.3.2 Interferometry

Where multiple fields $E_i(\omega)$ are involved, one can generalise (2.19) to show that the signal acquired on a slow integrating detector following a series of linear time-stationary operations can be reduced to a linear superposition of bilinear terms $E_i(\omega)E_j^*(\omega)$ for all pairs of fields (i, j). Under certain circumstances it is possible to obtain each of these terms individually. Taking the complex argument of each term $\phi_i(\omega) - \phi_j(\omega)$, one obtains the *spectral phase difference* between pairs of fields. This is the basis of *interferometry*. For simplicity, I shall consider two fields from now on.

Spectral phase differences acquired through interferometry can be put to several uses.

1. If the two pulses originate from the same pulse, the difference in optical path lengths between the beam splitter and the detector is obtained. This is a widely used means of measuring the dispersion of transparent materials. Whilst not technically a pulse characterisation method, it is very closely related and uses a similar setup [83–86].

2. If the spectral phase difference is related to the spectral phase of an unknown pulse by some invertible mathematical operation, then the phase of the unknown pulse can be inferred. This represents self-referenced pulse characterisation, and according to the restrictions above, requires nonlinear or time-nonstationary elements to prepare. This will be discussed in the section on complete characterisation.

3. If one of the fields is a well-characterised reference, then the spectral phase of the other can be inferred.

The third application represents a powerful externally referenced characterisation technique [87, 88], with one specific implementation known as temporal analysis by dispersing a pair of *e*-fields (TADPOLE). One of its main advantages is sensitivity; the unknown pulse heterodynes with the reference, and detection is completely linear. Provided the reference is sufficiently intense, a

train of pulses with an energy corresponding to less than one photon per pulse can be measured. The main constraint is that the spectral intensity of the reference pulse must be nonzero wherever the spectral intensity of the unknown pulse is nonzero. This implies that the reference pulse must have greater bandwidth than the test pulse. It is suitable for characterising fields produced by pulse shapers, which despite their potential complexity, cannot have a bandwidth exceeding that of the unshaped input pulse due to the linearity of the device. The unshaped pulse is generally simple to characterise using self-referenced methods. By the same token however, reference-based interferometry cannot determine the phase of the new frequencies generated in nonlinear spectral broadening processes.

Interferometry is sensitive to all orders of the phase difference between the test and reference pulses, including the zeroth order, or absolute phase. Whilst this may be advantageous in certain situations, it also necessitates single-shot acquisition when the absolute phase of one or both of the pulses is fluctuating.

There are several specific means of obtaining $E_1(\omega)E_2^*(\omega)$.

Field crosscorrelation. The field autocorrelation may be generalised to a crosscorrelation by sending different beams onto the detector [89].

$$B_{\text{CC2}}(\tau) = \int_{-\infty}^{\infty} E_1(t)E_2^*(t-\tau)\,\mathrm{d}t. \tag{2.23}$$

The Fourier transform of the field crosscorrelation is then $E_1(\omega)E_2^*(\omega)$. This is *time-domain interferometry* [90, 91].

Fourier-transform spectral interferometry (FTSI) begins by measuring the spectrum of the coherent sum of the two fields. The intensity of the superposition depends on the relative phase of the fields due to the presence of constructive or destructive interference. However, without further assumptions this procedure does not uniquely define the relative phase, as shown by the following argument: let the two complex-valued fields be E_1 and E_2, and the superposition $E_S = E_1 + E_2$. Intensity measurements provide the intensities of all three fields (individual intensities can be obtained by blocking each field in turn). The cosine rule of triangles then relates the relative phase

$\phi_2 - \phi_1$ between E_1 and E_2 to the intensities as

$$\cos(\phi_2 - \phi_1) = \frac{|E_S|^2 - |E_1|^2 - |E_2|^2}{2|E_1||E_2|}. \tag{2.24}$$

However, the inverse cosine is not single-valued, producing a *sign ambiguity* in the relative phase. Another concern with this direct approach is that it requires precise separate measurements of the individual intensities and is sensitive to imperfections (such as nonlinearity) in the detector response.

In FTSI, this problem is solved by the use of a *carrier phase* caused by a time delay τ between the pulses, which in the frequency domain amounts to an additive phase $i\omega\tau$. Figure 2.3 illustrates this. The superposition is of the form

$$B(\omega) = |E_1(\omega) + E_2(\omega)e^{i\omega\tau}|^2 \tag{2.25}$$

$$= |E_1(\omega)|^2 + |E_2(\omega)|^2 + 2|E_1(\omega)||E_2(\omega)|\cos[\Gamma(\omega) + \omega\tau] \tag{2.26}$$

where $\Gamma(\omega) = \phi_2(\omega) - \phi_1(\omega)$ is the relative phase. The cosine term produces fringes of nominal period $2\pi/\tau$ in the frequency domain, whose spacing is modulated by the relative phase. For the fringes to be well-defined, they must be on a finer scale than any phase or amplitude variation of the two fields. The spirit of this condition is expressed by the following equation

$$\frac{1}{|E_1(\omega)|}\left|\frac{\partial E_1(\omega)}{\partial \omega}\right| \ll \tau \tag{2.27}$$

and likewise for $E_2(\omega)$. Absent in (2.27) is a quantitative specification of the required separation of scales — this depends on the desired accuracy. A useful practical criterion, which naturally arises from the time-domain processing procedure, shall be described below.

Given a well-defined fringe pattern, various algorithms exist for extracting the relative phase $\Gamma(\omega)$. Intuitively the most direct is to pick the peaks and compare their positions with those expected when $\Gamma(\omega) = 0$. However a more robust and widely used method which is adopted here is Fourier-domain filtering [92], which utilises the separation of the fringe pattern from the individ-
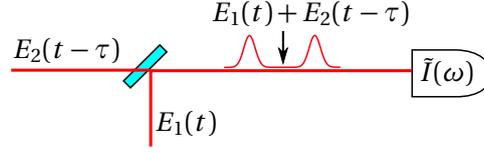
Figure 2.3: Apparatus for performing FTSI.
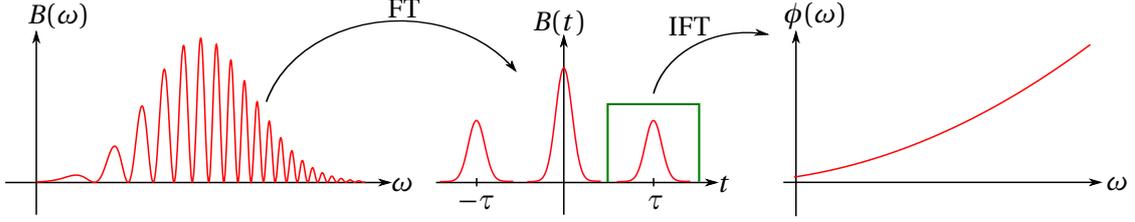
Figure 2.4: Fourier transform interferometry; in this example one of the pulses is transform-limited and one has a linear chirp. The filter in the quasi-time domain is shown in green.

ual intensities $|E_1(\omega)|^2$ and $|E_2(\omega)|^2$ upon Fourier transformation of (2.26). Rewriting (2.26), this can be made explicit:

$$B(\omega) = |E_1(\omega)|^2 + |E_2(\omega)|^2 + \left[E_2(\omega) E_1^*(\omega) \exp(i\omega\tau) + \text{c.c.} \right] \tag{2.28}$$

where c.c. denotes complex conjugate. The first two terms, which are real and slowly varying, are located at $t = 0$ upon Fourier transformation to the time domain. These form the *baseband* component. The first bracketed term is located at $t = \tau$, and its complex conjugate is located at $t = -\tau$. These are the positive and negative *sidebands*. Multiplication by a rectangularly-shaped passband filter and inverse Fourier transforming allows one of the sidebands to be isolated. The process is illustrated in Fig. 2.4. The time-domain filtering step is only possible if the sidebands are separate from the baseband component. As has been mentioned, there is a degree of arbitrariness since the components generally decay smoothly. However a useful practical criterion is for the signals to drop below the background noise level in the quasi-time domain.

The carrier amplitude required for separateness of the baseband and sidebands depends on their spread in the time domain, which can be derived using the correlation properties of the Fourier transform of (2.28). In the quasi-time domain, the baseband terms are the field auto-correlation function, the duration of which is a factor of order unity times the transform-limited

pulse duration. For a Gaussian pulse the factor is $\sqrt{2}$. The sideband terms are the field crosscorrelation function, with a duration which depends on the transform-limited duration and relative amount of chirp of the two pulses. In general, the greater the difference in chirp, the longer the crosscorrelation. The support of the crosscorrelation (down to some nominal noise level) does not exceed the sum of the supports of the two pulses.

Regardless of the details, a well-resolved fringe pattern must be sampled at least three times more finely than required for the pulse spectra themselves. This challenges the resolution of many spectrometers, and is the principal limitation of FTSI. For example, one immediate application is the characterisation of the output of femtosecond pulse shapers. However, these typically have around 512 independent channels, implying that an effective resolution of around 1500 independent channels would be required on the spectrometer. This is one motivation for using a spatially resolved spectrometer, as I describe next.

Fourier-transform spatio-spectral interferometry (FTSSI) adds a spatial dimension to the purely spectral form [71, 93–96]. The beams are brought to interference on a spatially resolved spectrometer, such as an imaging spectrometer, with a tilt θ which produces a spatially varying linear phase $\theta k y$ between the beams. The interference fringes so produced are approximately parallel to the frequency axis. Figure 2.5 illustrates the concept. The Fourier-domain filtering algorithm is identical except that a two-dimensional transform is used. The advantages of achieving spatial resolution will be discussed in section 2.4. However, even for purely spectral measurements, this approach has the advantage of alleviating the spectral resolution requirement of spectral interferometry because it obviates the need for a temporal carrier. Although this in turn introduces a corresponding spatial sampling requirement, this is less commonly a problem because complex spatial structure is usually rare and undesired in an ultrashort pulse.

Multiple-trace interferometry. Whilst the use of a carrier is common and used throughout this thesis, it is also possible to perform interferometry without one. In general, carrier-free interferometry relaxes the required resolution at the expense of requiring multiple shots and/or detectors. The procedure involves obtaining several interferograms, each with a different phase shift applied to one of the fields. The interferometric contribution to the j-th interferogram is proportional to
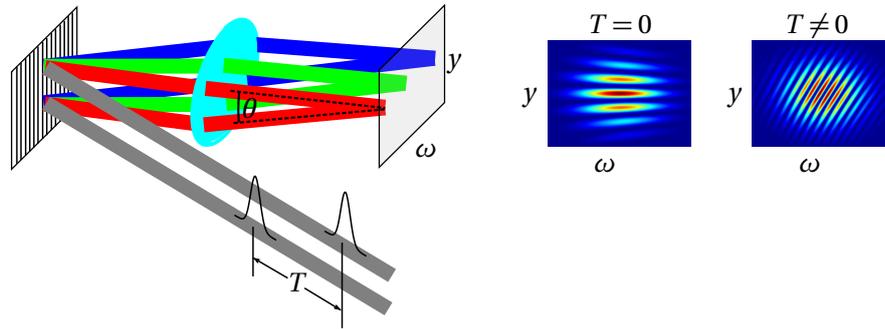
Figure 2.5: Fourier-transform spatio-spectral interferometry. The beams are dispersed and focused in the horizontal plane. In the vertical plane they impinge on the detector at an angle. Typical fringe patterns with and without a time delay are shown.

$\cos(\phi_2 - \phi_1 + \alpha_j)$ where α_j is the small phase shift applied at that shot. Judicious choice of α_j enables unique determination of $\phi_2 - \phi_1$ — for example, if the phase shifts are 0 and $\pi/2$ radians, then the cosine and sine of the phase difference is obtained, enabling unambiguous and robust quadrature detection [88]. An alternative is to scan the phase shift continuously, acquiring a two dimensional data set [97]. A Fourier-domain filtering step similar to the carrier case yields the phase of the interferometric term.

2.3.3.3 Summary — linear time-stationary operations

Linear-time stationary operations and slow photodetectors can recover the spectral intensity $|E(\omega)|^2$ of a single pulse, and pairwise interferometric products $E_i(\omega)E_j^*(\omega)$ of multiple pulses. When a suitable reference is available, interferometry is one of the most powerful characterisation methods available.

2.3.4 Nonlinear correlations

Nonlinear correlations represent the next step in complexity, and involve mixing the unknown pulse with a gate pulse, and recording the energy of the result as a function of the time delay of the gate pulse. If the gate is the unknown pulse itself, the method is an autocorrelation; otherwise it is a crosscorrelation. In general, nonlinear correlations are not considered complete characterisations unless additional information or assumptions are made. Nonetheless they are important historically, and remain in use for their relative simplicity.

2.3.4.1 Second-order correlations

Second-order autocorrelations are obtained by mixing the unknown pulse and the time-delayed replica in a material with a second order response. The detected signal may be light produced through second-harmonic generation, two-photon absorption on a photodetector, or two-photon fluorescence, and in the most general case is of the form

$$B_{\text{AC2}}(\tau) = \int_{-\infty}^{\infty} |E(t) + E(t - \tau)|^4 \tag{2.29}$$

$$= \int_{-\infty}^{\infty} \left| E^2(t) + 2E(t)E(t - \tau) + E^2(t - \tau) \right|^2 \tag{2.30}$$

From (2.30), the second-order autocorrelation is seen to be the coherent sum of individual second harmonics of the pulse and its delayed replica (first and third summands) and a pulse resulting from sum-frequency mixing (second summand).

Noncollinear arrangement. Figure 2.6 shows two possible noncollinear arrangements for acquiring the second-order autocorrelation. In Fig. 2.6(a), the signal is produced by second-harmonic generation. The individual second-harmonic beams propagate in different directions and do not arrive at the detector, and hence only the mixing term in (2.30) appears in the signal. The remaining signal is the *intensity autocorrelation* [20, 98–101]:

$$B_{\text{AC2}}(\tau) = \int_{-\infty}^{\infty} I(t)I(t - \tau)\,\mathrm{d}t. \tag{2.31}$$

In Fig. 2.6(b), the signal is produced by a two-photon absorbing detector. Although all terms in (2.30) appear in the signal, the two beams have a spatially varying time delay. Averaged across the area of the detector, interference between the terms is washed out. The detected signal is then equal to (2.31) except for the presence of constant background terms.

The intensity autocorrelation is the autoconvolution of the temporal intensity. As such it does contain some information on the spectral phase via that quantity's influence on the pulse duration. However, there are many significant ambiguities [102, 103]. One common approach to esti-
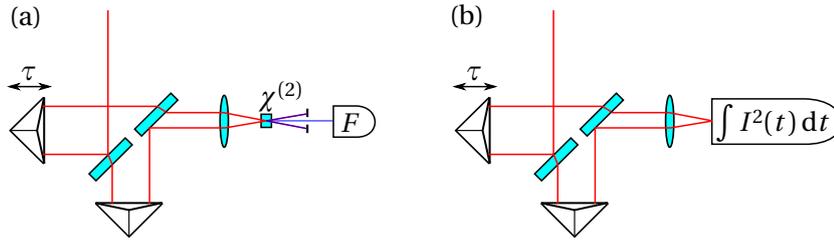
Figure 2.6: Arrangements for measuring the intensity autocorrelation using a balanced interferometer. The detected signal is (a) sum-frequency mixing of the replicas, and (b) two-photon absorption.

mating the pulse duration is to assume a certain functional form for the pulse, such as a linearly chirped Gaussian, and fit the expected intensity autocorrelation to the observed one, thus inferring a pulse duration. Combining this information with a measurement of the spectrum provides an estimate of the proximity of the pulse to the transform limit. Additional information may be inferred by performing a series of intensity autocorrelations with different amounts of dispersion applied [104].

Collinear arrangement. If the arrangement is collinear, as depicted in Fig. 2.7, then the entire signal in (2.30) is detected, giving the *interferometric autocorrelation* [105–114]

$$B_{\text{IAC}}(\tau) = \int_{-\infty}^{\infty} 2I(t) + 2\operatorname{Re}\left[E(t)E^*(t-\tau)\right]^2 +$$

$$4\operatorname{Re} I(t)E(t)\left[E(t-\tau) + E(t+\tau)\right]^* + 4I(t)I(t-\tau)\,\mathrm{d}t. \tag{2.32}$$

The second summand oscillates (in τ) at frequency $2\omega_0$, whilst the third summand oscillates with frequency ω_0. The fringe pattern so formed is phase sensitive, and so the interferometric autocorrelation provides more information about the pulse than the intensity autocorrelation. In general, fringe visibility decreases as the spectral phase moves away being from transform-limited. An iterative deconvolution algorithm exists for retrieving the spectral phase given an interferometric autocorrelation combined with a measurement of the spectrum [107]. This algorithm has at least one ambiguity: that of time reversal. Additionally, since interferometric autocorrelation encodes details of the pulse shape onto weak fringes on a high background, an extremely high SNR is often required. No general theory of the ambiguities of the interferometric autocorrelation exists.
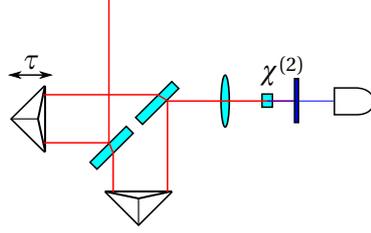
Figure 2.7: Arrangement for measuring the interferometric autocorrelation using a balanced interferometer.

Crosscorrelations. The noncollinear second-order crosscorrelation [25]

$$\int_{-\infty}^{\infty} I_1(t) I_2(t-\tau) \, \mathrm{d}t \tag{2.33}$$

is useful if one of the pulses is much shorter than the other; in that limit, it yields the temporal intensity. This is particularly useful when shapers are used to produce complex temporal profiles such as pulse sequences from near transform-limited input pulses [115].

Self-referenced crosscorrelations, in which the second pulse is derived from the first by applying some known operation such as dispersion, have also been performed. Combined with the spectrum, a complete characterisation using an iterative error-minimisation algorithm has been demonstrated; this has been called the PICASO method [116, 117]. No general theory of the ambiguities of this method exists.

2.3.4.2 Higher-order correlations.

Higher-order correlations yield signals of the form

$$\int_{-\infty}^{\infty} I_{\mathrm{pr}}^n(t-\tau) I(t) \, \mathrm{d}t. \tag{2.34}$$

The *probe* pulse $I_{\mathrm{pr}}^n(t)$ is itself nonlinearly generated from either the test pulse, for an autocorrelation, or another pulse, for a crosscorrelation. Its nonlinear dependence on intensity gives it higher contrast and shorter duration, and so higher-order correlations yield less blurred versions

of the test pulse. If n is large enough, $I_{\text{pr}}^n(t)$ approaches a delta function, and the signal is a close approximation to $I(t)$.

Third-order correlations have been implemented using Kerr-induced birefringence [118]. This is particularly useful for characterising ultraviolet pulses because no upconversion of the test pulse is required [119]. Two stages of sum-frequency generation can also be used [120]; the first stage producing the probe at frequency $2\omega_0$ from the test pulse at ω_0, and the second stage mixing the probe and test pulse to produce the signal at $3\omega_0$. This arrangement has a low background because scattering from the probe and test pulses can be blocked from the signal. Parametric gain is another possible nonlinearity [121]. Single-shot geometries are available [122, 123].

A version of PICASO using a third-order correlation has been demonstrated [124]. This removes some of the ambiguities of the second order version, but no general theory of ambiguities exists.

Finally, the temporal intensity may be obtained exactly from the triple intensity correlation [125], measured as a function of the two relative delays between the three replicas of the pulse combined with a third-order nonlinearity.

2.3.4.3   Single-shot geometries

Implementations which depend on scanning a delay stage involve averaging over many shots, making the implicit assumption that the pulse train is stable. For low repetition-rate amplifiers, which also often have poor shot-to-shot stability, single-shot geometries are preferable. Most single-shot correlators use the space-to-time mapping $\tau = x\theta/c$ created between two beams which cross at an angle θ in the nonlinear medium [126–130]. Spatially resolving the signal onto a detector array enables the correlation signal at a range of time delays to be acquired simultaneously. However, the spatial profile of the beam applies an envelope to the signal which must be considered.

2.3.4.4   Summary — nonlinear correlations

Nonlinear correlations provide information on the pulse duration. With additional measurements and assumptions, some techniques for complete characterisation have been demonstrated but

these have not found widespread use and are difficult to analyse for ambiguities.

2.3.5 Time-frequency distributions

Direct measurement of time-frequency distributions constitute the one of the three broad classes of self-referenced complete characterisation techniques. Historically, they were the first to be developed and remain in widespread use. This dissertation is about interferometry, and to the extent that this dissertation concerns purely temporal and spectral characterisation — that is, chapters 3 and 4 — time-frequency methods represent the major "competition". It is therefore worth discussing their strengths and weaknesses.

Practically measurable distributions can be generally defined as the energy of the optical pulse after the application of a *temporal gate* and a *spectral gate*. The temporal gate selects portions of the pulse near a tunable time τ, and the spectral gate selects frequencies around a tunable passband of centre frequency ω_{p}. In all experiments to date, the spectral gate has been a linear time-stationary filter, such as a tunable bandpass filter or the various channels of a spectrometer. There is greater variety in implementations of the temporal gate; nonlinear mixing, electronic sampling, optical streaking and even streaking of photoelectrons have been used, and these shall be discussed below. Apparatus for acquiring time-frequency distributions can be broadly classified by the order of the gates: if the temporal gate follows the spectral gate, the technique is one of *sonography*; otherwise it is *spectrography*.

Time-frequency distributions are often a very intuitive representation of a pulse and contain redundant information which is useful for error detection and self-consistency checking. Apparatus for acquiring time-frequency distributions are also amongst the simplest of the complete characterisation methods. On the other hand, the inversion of time-frequency distributions to yield the unknown pulse is mathematically complex and generally requires a computationally intensive iterative algorithm. Additionally, the indirect relationship between data and reconstructed pulse complicates analysis of the effects of experimental error. The need to acquire two-dimensional data leaves no dimension "free" for spatial resolution, at least on conventional detectors.

Most time-frequency distributions have a few well-known exact ambiguities. These can often
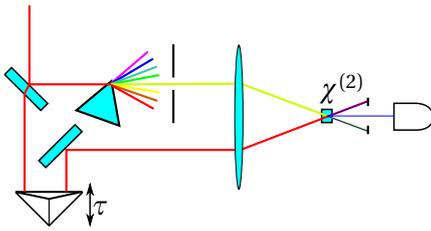
Figure 2.8: Typical apparatus for measuring the sonogram; the angularly dispersive element is drawn as a prism but could equally be a grating.
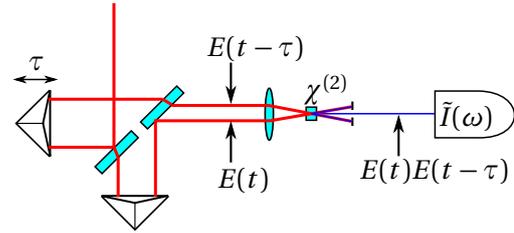
Figure 2.9: Typical apparatus for measuring an SHG-FROG.

be resolved by incorporating additional information, such as a measurement of the pulse spectrum, or performing some additional measurements. General proofs of the nonexistence of additional exact ambiguities are rare. The action of noise and other experimental distortions, which may introduce approximate ambiguities, is also rather difficult to analyse. In general, the probability of encountering an ambiguity increases with pulse complexity, particularly when multiple subpulses are present [131, 132]. Also, the presence of ambiguities appears to depend on numerical parameters, particularly the discretisation of data [133–136].

2.3.5.1 Sonography

Sonograms consist of a frequency gate followed by a time gate; the frequency gate is a tunable bandpass filter, whilst the time gating may be accomplished using a streak camera [137–139], electronic sampling of the intensity [138], or a nonlinear optical interaction [140–142] with a gate pulse. In the well-known methods considered here, all the temporal gates may be represented by (2.20), where the gate function $G(t)$ determines the temporal resolution of the measurement, which for streak cameras and electronic sampling is 10 ps and 50 ps respectively. The recorded signal is

$$B(\tau, \omega_\mathrm{p}) = \int_{-\infty}^{\infty} \left| G(t-\tau) \int_{-\infty}^{\infty} E(\omega) R(\omega; \omega_\mathrm{p}) e^{-i\omega t} \, \mathrm{d}\omega \right|^2 \, \mathrm{d}t \qquad (2.35)$$

where $R(\omega; \omega_\mathrm{p})$ is the response of a tunable bandpass filter with centre frequency ω_p. One implementation of a sonogram is shown in Fig. 2.8.

Sonograms are intuitive time-frequency representations. In particular, by taking moments

37

over τ, the frequency-dependent group-delay can be directly extracted [141, 142]. Whilst this is often a useful quantity, for example in measurements of fibre dispersion, it must be integrated to obtain the complete spectral phase, an approach which struggles when the spectrum of the pulse contains nulls or if the phase contains discontinuities. General phase retrieval from the sonogram requires an iterative algorithm [140, 143]. In contrast with many spectrographic distributions, the sonogram does not possess the time-reversal ambiguity and the sign of any group-delay dispersion is directly reflected in the data.

Single-shot sonograms may also be acquired. Using the angularly-varying phase matching response of upconversion in a nonlinear crystal, it is possible to achieve a frequency-to-space mapping which serves as the spectral gate [144], whilst a noncollinear second order interaction provides the time gate. Alternatively, the pulse may be dispersed using a prism or grating, and then crosscorrelated in noncollinear single-shot fashion using two-photon absorption on a CCD camera [145, 146].

2.3.5.2 Spectrography

In spectrograms, the gating order is reversed with respect to sonograms; the temporal gate is applied before the frequency gate. The latter is usually a spectrometer with enough resolution to resolve all the spectral features of the signal; the spectrogram is therefore the spectral intensity of the gated unknown pulse. Early usage of spectrograms [147–149] was restricted to interpretations of the chirp and did not attempt full retrieval of the amplitude and phase. The development of an algorithm for inverting the spectrogram to retrieve the unknown pulse [45] heralded the introduction of FROG. This approach has become popular, and many variations have been developed.

Second-harmonic generation (SHG)-FROG. In SHG-FROG the unknown pulse undergoes sum-frequency generation with a delayed replica in a second-order nonlinear medium [150–152]. The measured trace is then

$$B(\tau, \omega) = \left| \int_{-\infty}^{\infty} E(t)E^*(t - \tau)e^{i\omega t}\, \mathrm{d}t \right|^2. \tag{2.36}$$

An apparatus for obtaining an SHG-FROG trace is shown in Fig. 2.9.

Since it uses a second order nonlinearity, it is the most sensitive of the FROG family, and has

been used to characterise nanojoule pulses from an oscillator [49] and in telecommunication systems [153–155]. Enhanced sensitivity can be achieved by conducting the interaction in waveguides [156]. By taking into account various spectral factors arising from the response of the optical elements and the phase-matching, and temporal factors arising from the crossing angle of the beams, few-cycle optical pulses can be measured using SHG-FROG [61, 157, 158]. Characterisation of mid-IR free electron laser pulses [159] also suggest the method's versatility.

The aforementioned SHG-FROG geometries were noncollinear — essentially a spectrally resolved intensity autocorrelator. A collinear geometry is also possible, avoiding temporal blurring artefacts arising from the beam crossing angle. Using type II sum-frequency generation, in which the polarizations of the replicas are mutually orthogonal [160], the acquired signal is the same as the noncollinear case [161–163]. Using a type I interaction, in which the polarizations are parallel, one obtains a spectrally resolved interferometric autocorrelation. Although the interference terms can be numerically removed to retrieve a normal SHG-FROG trace [161], they can be exploited for an independent reconstruction [164, 165].

Whilst SHG-FROG is sensitive and versatile, it has the major drawback of possessing a time reversal ambiguity, stemming from the fact that the temporal gate is multiplication by the unknown pulse itself. In the frequency domain, this corresponds to a phase sign ambiguity; positive and negative dispersion cannot be distinguished. This is resolveable by applying recognisable features to the pulse — introducing glass of known dispersion or using multiple reflections to produce trailing satellite pulses [166].

Third-order nonlinearities. Self-diffraction FROG (SD-FROG) is one of a number of FROG geometries based on the third order nonlinearity. Whilst generally possessing a lower sensitivity than SHG-FROG, third-order spectrograms are often more intuitive [167] and lack the reversal of time ambiguity. In SD-FROG, the test pulse and its delayed replica cross in a third-order medium, forming a refractive index grating via the Kerr effect. The test pulse and replica scatter from this

grating. Spectrally resolving the intensity of one of the diffracted beams gives the SD-FROG signal

$$B(\tau, \omega) = \left| \int_{-\infty}^{\infty} E^2(t) E^*(t - \tau) e^{i\omega t} \, \mathrm{d}t \right|^2. \tag{2.37}$$

In polarisation-gating FROG (PG-FROG), the delayed replica induces a time-dependent bire-fringence in a third order medium [45]. The test pulse, polarised at 45°to the replica, experiences a polarisation rotation which depends on the temporal intensity of the delayed replica $I(t - \tau)$. An analysing polariser, placed after the nonlinear medium, selects the component of the test pulse which has been rotated by 90°. For a small birefringence, this is proportional to $E(t)I(t - \tau)$, so the measured signal is

$$B(\tau, \omega) = \left| \int_{-\infty}^{\infty} E(t) I(t - \tau) e^{i\omega t} \, \mathrm{d}t \right|^2. \tag{2.38}$$

Because the gate is real-valued, PG-FROG has a particularly intuitive trace and, for simple reference pulses without significant satellite structure, conforms to the technical mathematical notion of a spectrogram.

Other forms of the third-order nonlinearity used for FROG include third-harmonic generation [168–170], transient grating [171], and four-wave mixing [172, 173].

Crosscorrelation frequency-resolved optical gating. (XFROG) is the nonself-referenced version of FROG, in which the temporal gating is performed using a well-characterised reference pulse [174–178]. It is potentially more sensitive than self-referenced versions, because a strong reference pulse may be used to characterise a weak unknown. Because the reference may be at a different frequency to the unknown, a wider range of phase-matched nonlinear processes are available. If the reference pulse is simple, then the acquired signal is similar to the gated Fourier transform, eq. (2.11), which has an intuitive interpretation. It also lacks the arrival time ambiguity. Notable uses for XFROG include spectrograms of supercontinuum [64, 179] and soliton dynamics [180] in photonic crystal fibres.

Single-shot geometries of FROG employ the same principle as single-shot autocorrelators, with the beams crossing at an angle in the nonlinear medium. The gated pulse is then spectrally dis-

persed in the plane perpendicular to the crossing angle. This can be achieved by re-imaging the plane of the nonlinear medium onto the entrance slit of an imaging spectrometer [181], or using the angular dispersion produced by phase-matching in a thick nonlinear crystal [182, 183]. Either way, time delay and frequency are mapped to two dimensions of a detector array, and a single-shot FROG trace acquired.

2.3.5.3 Summary — time-frequency distributions

The measurement and inversion of time-frequency distributions is a general approach to the complete characterisation of ultrashort pulses. FROG and its variants are particularly widespread. The dataset is inherently two-dimensional and its inversion requires an iterative algorithm.

2.3.6 Self-referenced interferometry

Self referencing is achieved in spectral interferometry when the two pulses have some relation to one another so that knowledge of their spectral phase difference permits unique determination of their spectral phase.

By far the most common approach is spectral shearing interferometry (SSI), in which the pulses being interfered are related by a translation, or *shear*, along the frequency axis. This underlies a broad category of ultrashort pulse characterisation devices, the most longstanding and well known of which is SPIDER. The spatial analogue, lateral shearing interferometry (LSI) has a long history. Besides wavefront characterisation it is extensively used in the testing and characterisation of optical components. It plays a fundamental role in this dissertation and is discussed in section 2.4.4.2. However, most of its features are completely analogous to the spectral case.

Several important features of shearing interferometry deserve emphasis up-front. It is a phase-only measurement — the intensity must be acquired separately, although in many cases this is conveniently accomplished using only a minor adjustment to the apparatus. There is a computationally efficient and algebraically simple mapping between the phase and the data, which simplifies processing and error analysis. Finally, an n-dimensional field can be reconstructed[1] using n n-dimensional datasets. Applying this rule to the most common case, a one dimensional phase

---

[1] except for ambiguities which I shall discuss in a subsequent section

measurement can be performed using only a single one-dimensional dataset.

The principle of spectral shearing interferometry is as follows: two replicas of the unknown pulse are produced, each identical to the original except for frequency shift of ω_s and $\omega_s - \Omega$ respectively. The replicas are therefore $E(\omega - \omega_s)$ and $E(\omega - \omega_s + \Omega)$. The shear Ω determines the spectral resolution of the measurement, and without loss of generality I take $\Omega > 0$. Interferometry between the replicas reveals their phase difference

$$\Gamma(\omega - \omega_s) = \phi(\omega + \Omega) - \phi(\omega). \tag{2.39}$$

Several procedures exist for reconstructing $\phi(\omega)$ from $\Gamma(\omega)$, each making subtly different assumptions about the unknown field. The most straightforward procedure is *concatenation*, which takes as its premise the finite support of the field in the time domain. That is, the temporal field is assumed to be nonzero only for a duration $T = 2\pi/\Omega$. Then, by the sampling theorem [184] the field is uniquely defined by its values at a regularly spaced set of frequencies $\omega_n = \omega_0 + n\Omega$, where $n = 0, 1, \ldots$ and ω_0 is an arbitrary frequency origin. The unknown phase can be recovered at these points by repeated invocation of (2.39), yielding

$$\phi(\omega_n) = \phi(\omega_0) + \sum_{j=0}^{n-1} \Gamma(\omega_j - \omega_s). \tag{2.40}$$

2.3.6.1 Ambiguities in shearing interferometry

The presence of the arbitrary phase offset $\phi(\omega_0)$ on the right-hand side of (2.40) shows that shearing interferometry does not recover the absolute phase.

Spectral shearing interferometry often has an arrival time ambiguity as well, although this is not shown in (2.40). The linear component of the reconstructed spectral phase, and hence the arrival time of the reconstructed pulse, is determined by the constant component of $\Gamma(\omega)$, which in turn is given by the difference in the absolute phases of the replicas. If the frequency shift changes these absolute phases in a way which is difficult to predict, then the reconstructed arrival time is effectively random.
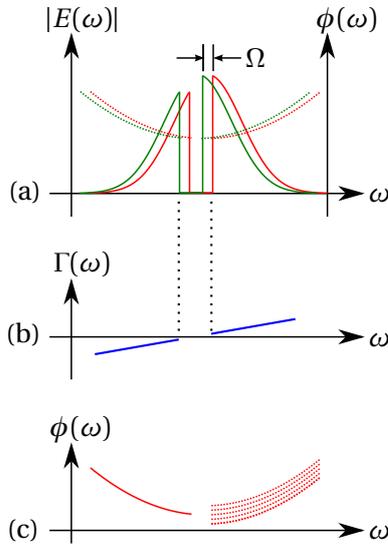
Figure 2.10: The relative phase ambiguity in spectral shearing interferometry; (a) original (red) and sheared replica (green) of a pulse with quadratic spectral phase and a spectral null; the shear Ω is less than the width of the null. (b) Phase difference between the replicas and original, blanked out where there is zero overlap and phase difference is undefined. (c) Reconstructed phase; the phase of the low-frequency lobe has been set to an arbitrary value but the relative phase of the high-frequency lobe is unconstrained by the data.

Another ambiguity in shearing interferometry occurs when the spectrum has two or more disjoint components, separated by a *spectral null* region wider than the shear. In this case, the phase differences are undefined at one or more points between the components, and will take a completely random value in an actual reconstruction. The concatenation algorithm will "blindly" add this random value as it steps across the spectrum. Whilst the individual components are still reconstructed correctly, their relative phase is not determined. More generally, if the spectrum is written as the sum of two components $E(\omega) = E_1(\omega) + E_2(\omega)$, and the two components are separated by more than the spectral shear, then the measured interferometric product $E(\omega+\Omega)E^*(\omega)$ is completely independent of the relative phase of the components. Simply increasing the shear may not be an option because this decreases the spectral resolution, or equivalently the time window of the reconstructed pulse; indeed setting the shear to be larger than a spectral null is somewhat contradictory since the shear should in principle be fine enough to resolve all spectral features.

One of the results of this thesis is a means of alleviating this ambiguity — a reconstruction algorithm which permits shears of different sizes to be combined. It is presented, along with some

specific motivations, in chapter 4.

2.3.6.2 Specific implementations of spectral shearing interferometry

Spectral phase interferometry for direct-electric field reconstruction (SPIDER) is a popular technique, with many variations [185, 186], based on Fourier-transform spectral interferometry between spectrally sheared replicas produced by nonlinear optical mixing with a quasi-monochromatic *ancilla* field.

A standard SPIDER setup is shown in Fig. 2.11 [46, 50, 187]. The test pulse is incident upon an etalon, which produces two reflections separated by a time delay τ. The transmission through the etalon is passed through a *stretcher*, which applies a large spectral chirp $\phi^{(2)}$ to produce the ancilla. The degree of the chirp of the ancilla is such that it is quasi-monochromatic over the duration of the test pulse. The ancilla and the two reflections of the test pulse are focused into a $\chi^{(2)}$ crystal, where sum-frequency generation between them produces two spectrally shifted replicas with spectral shear $\Omega = \tau/\phi^{(2)}$. Fourier-transform spectral interferometry is then performed between the upconverted replicas. The measured signal is obtained making the appropriate substitutions into (2.26):

$$B(\omega + \omega_0) = |E(\omega) + E(\omega + \Omega)e^{i\omega\tau}|^2 \tag{2.41}$$

$$= |E(\omega)|^2 + |E(\omega + \Omega)|^2 + 2|E(\omega)||E(\omega + \Omega)|\cos\left[\phi(\omega + \Omega) - \phi(\omega) + \omega\tau\right] \tag{2.42}$$

where ω_0 and $\omega_0 - \Omega$ are the frequencies of the ancilla which are contemporaneous with the test pulse replicas in the crystal. The Fourier-domain filtering procedure of section 2.3.3.2 is performed to obtain the phase differences $\Gamma(\omega) = \phi(\omega + \Omega) - \phi(\omega)$. These form the input to one of the reconstruction procedures discussed in Section 2.3.6, yielding $\phi(\omega)$.

Important features of SPIDER are its one-dimensional encoding and computationally simple and noniterative reconstruction algorithm. This lends itself to single-shot measurement with rapid update rates [56, 188, 189]. SPIDER, along with other techniques that perform spectral shearing using nonlinear optics, also has less stringent phase-matching requirements than spectrographic-
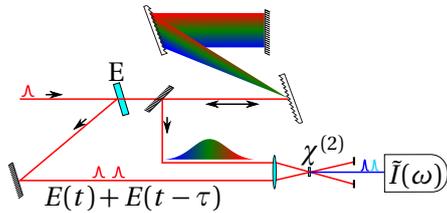
Figure 2.11: Schematic of a temporally encoded SPIDER device; E is an etalon.
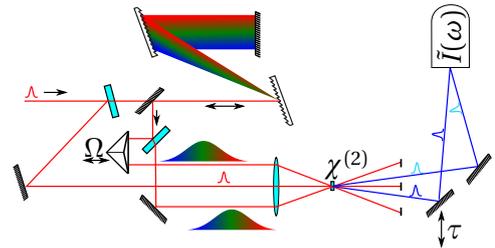
Figure 2.12: Schematic of a ZAP-SPIDER — see text for an explanation.

or correlation-based methods because of the quasi-monochromaticity of the ancilla; a phase-matching response that encompasses the test pulse is only required for one of the two fields involved in the nonlinear interaction. Furthermore, since the spectral phase information is encoded in the fringe spacing, it is invariant to the spectral response of the apparatus. This simplifies implementations of SPIDER for broadband and few-cycle pulses [50, 190, 191]. Reduced to the barest of essentials, SPIDER is only a spectral phase measurement; external measurement of the spectral intensity is necessary for a complete characterisation. However in several common cases it is possible to extract the spectral intensity from the same data trace as the phase [192, 193].

The cost of these advantages is, in general, additional complexity. Equation (2.42) makes evident three parameters ω_0, τ and Ω, which must be calibrated. A pulse stretcher for the ancilla with an appropriate chirp must be designed and constructed, unless a chirped pulse happens to be available. Because the test pulse is mixed with a stretched pulse, the signal intensity is generally lower than in spectrographic methods. This does not in itself imply that SPIDER is less sensitive, because its single-shot nature means that for a given total measurement time one may integrate the signal for longer than in a multi-shot method. However, it does introduce the practical difficulty of aligning a weak signal.

There is a menagerie of variations on SPIDER, with several playing an important role in this dissertation. It should be noted that the various features discussed below are not necessarily mutually exclusive.

Zero-additional phase SPIDER (ZAP-SPIDER) has the useful property that the test pulse does not pass through any transmissive optics before the nonlinear mixing step; only reflections are used
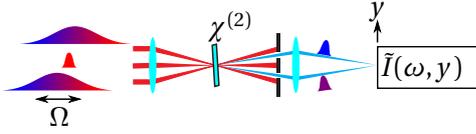
Figure 2.13: SEA-SPIDER concept — two chirped ancillae mix with the test pulse at an angle. The crystal is re-imaged onto the entrance slit of an imaging spectrometer.
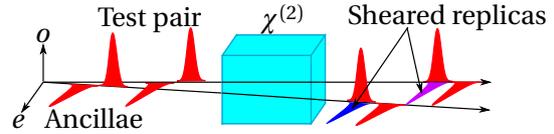
Figure 2.14: Producing spectrally sheared replicas using group-velocity mismatch in a long crystal.

[194, 195]. This obviates the need for calibration of the dispersion of these optics, which would otherwise be necessary for accurate characterisation of pulses with durations below 10 fs. In zero-additional phase designs, the test pulse is mixed with two ancillae propagating in different directions; the two independent mixing processes produce two upconverted replicas of the test pulse. The spectral shear results from the time delay between the ancillae. The replicas are then interfered with time delay τ on the spectrometer. A typical ZAP-SPIDER design is shown in Fig. 2.12. Besides the zero-additional phase property, this configuration allows the temporal carrier to be tuned independently of the spectral shear.

Spatially encoded arrangement SPIDER (SEA-SPIDER) uses a spatial carrier for the Fourier-transform interferometry [196, 197]. This improves the spectral resolution by a factor of ≈ 4 for any given spectrometer. The cost is that a spatially resolving spectrometer must be used. A SEA-SPIDER design is shown in Fig. 2.13. Several of the devices used in this thesis build upon SEA-SPIDER.

Homodyne optical technique SPIDER (HOT-SPIDER) is based on the realisation that sensitive measurement of the phase difference between two fields can be achieved by sequentially interfering each of them with a strong *local oscillator*. The local oscillator need not be characterised at all — the only requirement made upon it is that it stays constant between the two measurements, and encompasses the spectra of the two fields under test. The two measurements yield $\phi_1 - \phi_{\text{LO}}$ and $\phi_2 - \phi_{\text{LO}}$. Upon taking their difference, the phase of the local oscillator cancels out, leaving $\phi_2 - \phi_1$. In HOT-SPIDER, the replica with frequency shifts ω_0 and $\omega_0 - \Omega$ is sequentially interfered with the local oscillator, which has been produced using the second harmonic of the test pulse [198]. In chapter 7, the homodyne concept is applied in the context of high-harmonic generation.

Long-crystal arrangement SPIDER (LX-SPIDER). Spectrally sheared replicas may be produced without chirping the ancillae using the phase matching properties of a long crystal [199]. Several crystals exhibit the property that when oriented for type II sum-frequency generation, the acceptance bandwidth is much greater in the ordinary polarisation than the extraordinary polarisation. The asymmetry is due to group-velocity matching between the incident o-wave and the generated e-wave at the sum frequency, but a group-velocity mismatch between the incident and generated e-waves. In the time domain, the generated e-wave pulse travels at the same group velocity as the incident o-wave pulse but the incident e-wave pulse walks through them. (This description applied to a negative uniaxial crystal; for a positive crystal the o and e are reversed.)

Figure 2.14 illustrates the use of a long crystal to produce spectrally sheared replicas. The incident e-wave plays the role of the ancilla, with phase-matching selecting only a narrow portion of its spectrum for upconversion. The test pulse experiences the broad acceptance bandwidth of the o-ray, providing accurate upconversion. The ancilla pulse is sent into the crystal behind the test pulse and walks through it during propagation. Crucially, the phase-matched frequency is angle-dependent, so sending two pairs of beams at different angles achieves the desired spectral shear.

Chirped-Arrangement SPIDER (CAR-SPIDER) uses upconverted replicas with spatially varying upconversion frequency. The long-crystal concept may be used to this end. Instead of sending two noncollinear beams into the long crystal, as in an LX-SPIDER, a single tightly focused beam can be used since it contains a range of wavevectors [200]. Figure 2.15 illustrates the principle. Each wavevector experiences a different upconversion frequency. After subsequent propagation and then collimation, the generated pulse has an upconverted frequency which varies spatially: $E_A(\omega, x) = E(\omega - \omega_{up} + \alpha x)$. The spatially chirped replica is passed through a Mach-Zehnder interferometer which performs a spatial flip in one arm. This is accomplished by having an odd number of reflections (in the horizontal plane) in one arm and an even number in the other. The pulse from the flipped arm is therefore $E_B(\omega, x) = E(\omega - \omega_{up} - \alpha x)$. Interfering the two arms on an imaging spectrometer yields spectral shearing interferometry traces with a spatially varying shear $\Omega = 2\alpha x$.
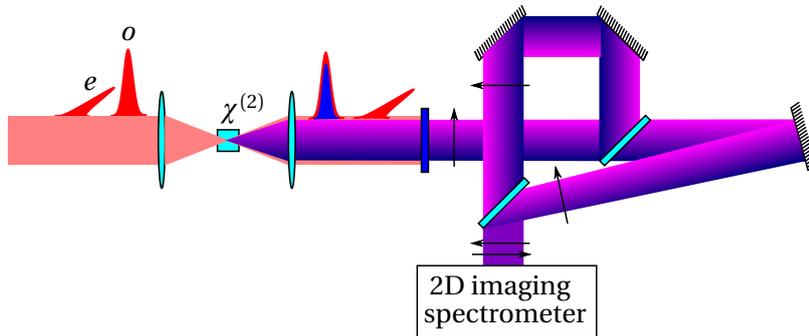
Figure 2.15: Principle of long-crystal CAR-SPIDER: the test pulse forms the *o*-wave and the ancilla forms the *e*-wave. The black arrows across the beams indicate the direction of increasing upconversion frequency.

Multiple-trace interferometry. It is possible to measure the phase difference between the sheared replicas using multiple trace methods similar to those discussed in section 2.3.3.2. These obviate the need for a temporal carrier, so that the replicas are not separated in time. This is essential, for example, with high-duty cycle arbitrary waveforms. A form of quadrature detection [201], in which the interferogram is sequentially measured with absolute phase of $0, \pi/2, \pi$, and $3\pi/2$, has applied to such cases. Another approach, two-dimensional spectral shearing interferometry [97], uses a design similar to a ZAP- or SEA-SPIDER, in which two chirped ancillae are mixed with a single test pulse. However, the upconverted replicas are detected colinearly and cotemporally on a one-dimensional spectrometer. The phase is uniquely determined by scanning the time delay of one of the ancilla over a few optical cycles (not enough to appreciably change the shear), producing a phase shift of the corresponding replica.

Spectral interferometry resolved in time (SPIRIT). A different approach is to spatially disperse the test pulse, as in a spectrometer, and then interfere two slightly displaced replicas of the dispersed pulse on a linear detector array. The field incident on any given pixel of the detector is $E(\omega) + E(\omega + \Omega)$, where the shear Ω is produced by the displacement. Of course, under normal circumstances interference between these two different frequencies does not occur. If however, the resulting signal is then time-gated with a duration less than the temporal stretching caused by the dispersion (the inverse of the spectrometer resolution), then the phase of the temporal beating produced by the frequency-shifted fields may be sampled. This is the basis of SPIRIT [202–204].

Self-referencing without a spectral shear. A displacement along the frequency axis is the most common operation used in self-referenced spectral interferometry. However, a recent development is the use of a nonlinear pulse shortening operation — essentially a pinhole in time [205]. This has the effect of flattening phase variations, producing a "reference" which is closer to the transform limit than the test pulse. Interferometric measurement of the phase difference between the test pulse and its shortened replica provides an estimate of the unknown pulse.

2.3.6.3   Summary — self-referenced interferometry

Spectral shearing is almost unique in enabling spectral interferometry to become self-referenced. Implementations have been demonstrated for wavelengths from the ultraviolet to the infrared, and for pulse durations down to the few-cycle regime. Its advantages are a one-dimensional encoding and a direct and easily invertible relation between the signal and the phase of the unknown pulse. However it requires careful calibration of the delay between the sheared pulse replica and synthesis of a shearing operation of reasonable fidelity.

2.3.7   Tomography

Tomographic methods operate by principles analogous to spatial imaging, based on the correspondence between the pair of transverse space and wavenumber co-ordinates (x, k_x) and time-frequency co-ordinates (t, ω). Paraxial diffraction over a distance L is a quadratic phase factor in transverse wavenumber $\exp[i L/(2 k_0) \, k_x^2]$, exactly the same mathematical form as second-order dispersion $\exp[i \phi_2/2 \, \omega^2]$. An aberration-free lens of focal length f introduces a spatial phase $\exp[-i k_0/(2 f) \, x^2]$, of the same form as the action of a quadratic temporal phase modulator $\exp[i \psi/2 \, t^2]$. By associating dispersion with diffraction and a lens with a quadratic temporal phase modulator, spatial imaging systems may be transformed to the time-frequency domain. However, achieving suitable quadratic phase modulation for subpicosecond pulses is difficult, and the best temporal resolutions that have been achieved using tomographic methods are around 200 fs. For this reason, tomographic methods do not at present represent a viable alternative to the techniques discussed in this thesis and so will be discussed briefly.

2.3.7.1 Time magnifier

A number of tomographic methods can be understood using familiar concepts in imaging. A single-lens imaging system has as its temporal analogy second-order dispersion $\phi_2^{(o)}$, a quadratic temporal phase ψ, and further second-order dispersion $\phi_2^{(i)}$, related by Newton's thin lens equation

$$\frac{1}{\phi_2^{(o)}} + \frac{1}{\phi_2^{(i)}} = \psi. \tag{2.43}$$

The temporal magnification is $m = -\phi_2^{(i)}/\phi_2^{(o)}$. The signal as detected on a photodiode is therefore a temporally magnified replica of the temporal intensity $I_{\text{sig}}(t) \approx I_{\text{orig}}(t/m)$. Ultrafast temporal profiles can be "slowed down" to the extent that they can be detected on a photodiode. Magnifications of 100 with 300 fs resolution have been achieved [206–209].

2.3.7.2 Time-to-frequency converter

Another well-known spatial operation that has been reproduced in the time-frequency domain is the spatial Fourier transform, yielding a time-to-frequency converter [210–214]. In space, this is achieved by propagation over a distance L, a lens of focal length $f = L$, and then another length L of propagation. The analogous processes are quadratic dispersion ϕ_2, quadratic temporal phase $\psi = -1/\phi_2$, and another application of dispersion ϕ_2. The field is then detected on a spectrometer; the spectrum is related to the temporal profile of the original pulse by $\tilde{I}_{\text{sig}}(\omega) \approx I_{\text{orig}}(-\phi_2 t)$. The best temporal resolutions achieved with time-to-frequency converters is about 220 fs [214].

Both the time-magnifier and the time-to-frequency converter obtain the temporal intensity. An amplitude and phase which is consistent with both the temporal intensity and the spectrum can be retrieved using the Gerchberg-Saxton algorithm [215, 216]. There are several exact or approximate ambiguities in such a procedure.

2.3.7.3 Tomography

The time-magnifier and time-to-frequency converter achieve their results through specific combinations of dispersion and quadratic phase. Arbitrary arrangements of these operations can be understood in terms of Wigner functions, also known as the chronocyclic representation. If $W(t, \omega)$

is the Wigner function of pulse $E(t)$, then application of a quadratic spectral phase modulation shears the Wigner function along the time axis yielding $W(t - \phi\omega, \omega)$, whilst a quadratic temporal phase modulation shears the Wigner function along the frequency axis yielding $W(t, \omega + \psi t)$. Any area and orientation preserving linear transformation in chronocyclic space can be composed of a series of such operations[1]. Temporal intensity is the projection of the Wigner distribution onto the time axis $I(t) = (2\pi)^{-1} \int_{-\infty}^{\infty} W(t, \omega) \, d\omega$; the spectrum is the projection of the Wigner distribution onto the frequency axis $\tilde{I}(\omega) = (2\pi)^{-1} \int_{-\infty}^{\infty} W(t, \omega) \, dt$.

A particularly important case is when the shears are combined to produce a series of rotations. For each rotation, a projection of the Wigner function is obtained. The reconstruction of an arbitrary n-dimensional function from a series of projections onto lower dimensional datasets is called tomography [218] and is of fundamental importance in medical imaging [219]. When the function in question is a Wigner distribution, the procedure is *chronocyclic tomography* [220]. This has not been experimentally demonstrated because of the difficulty in performing accurate and adjustable spectral and temporal phase modulations. Nonetheless, it would have an advantage over other methods because it directly measures the Wigner distribution, which is the lowest-order description of shot-to-shot fluctuations in a partially coherent pulse train. Other methods assume a coherent pulse train with no shot-to-shot fluctuation.

The assumption of a coherent pulse train provides the additional constraint that the unknown is a Wigner function of a single pulse. This enables reconstruction of the Wigner function from a limited set of projections [221, 222]. Such *simplified chronocyclic tomography* [223] has been demonstrated by performing two small rotations with an electro-optic phase modulator and measuring the resulting spectra.

2.3.7.4   Limits to temporal resolution

The resolution of a spatial imaging system is determined by the focal length f and the diameter D of its lens. Together, these form the numerical aperture $\mathrm{NA} = D/f$ which determines the maximum

---

[1]This result comes from group theory: the matrix representation of area- and orientation-preserving linear transformations is *special linear group* $\mathbf{SL}_2(\mathbb{R})$ of all 2×2 real matrices with determinant one. Shear matrices form a generating set of this group — all elements of $\mathbf{SL}_2(\mathbb{R})$ can be written as a product of shear matrices [217].

wavenumber $k_x^{\max} = kNA$ and hence the spatial resolution $1/k_x^{\max}$. Analogously, the temporal resolution of tomographic methods is limited by the quadratic phase modulation amplitude ψ and the period T over which it can be sustained. The analogue of the numerical aperture is the bandwidth $\omega^{\max} = T\psi$ of the modulation, which determines the temporal resolution $1/\omega^{\max}$.

The temporal quadratic phase may be applied using an electro-optic phase modulator, which imparts a quadratic temporal phase modulation by synchronising the test pulse with the peak of the sinusoidal refractive index modulation driven by a narrowband radio-frequency wave. The phase modulation is

$$\exp[iA\cos\omega_{\mathrm{m}}t] \approx \exp[iA(1-\omega_{\mathrm{m}}^2/2)t] \tag{2.44}$$

where A is the maximum phase modulation, ω_{m} is the modulation frequency and the quadratic expansion is performed around $t = 0$. For an electro-optic modulator one therefore has $\psi = A\omega_{\mathrm{m}}^2$. The sinusoidal phase modulation is quadratic for about 10% of one optical period, giving $T = 2\pi/(10\omega_{\mathrm{m}})$. The maximum temporal resolution available from this time lens is therefore $(0.1 \cdot 2\pi A\omega_{\mathrm{m}})^{-1}$. Modern electro-optic modulators achieve a bandwidth-modulation product $A\omega_{\mathrm{m}} \approx 100$ GHz, implying a resolution of approximately 500 fs.

To increase the speed, the quadratic temporal phase is applied by a nonlinear interaction, such as cross-phase modulation [213, 224], sum-frequency generation [208, 210], and four-wave mixing [214]. The bandwidth required from the auxiliary pulse must be at least the bandwidth of the test pulse, and the auxiliary pulse must be carefully prepared so as to introduce a quadratic temporal phase. Such an implementation is not self-referenced. Alternatively, the auxiliary pulse may be derived from the test pulse itself, in which case the phase modulation is unknown and an iterative retrieval algorithm is needed. Such a technique is tomographic only in a loose sense. Self-phase modulation has been used for this purpose [225, 226], achieving a temporal resolution of 20 fs. Some ambiguities were noticed in the initial investigation and no general analysis of ambiguities has been performed.

2.3.7.5   Summary — tomography

Tomography of ultrashort pulses exploits powerful analogies with imaging, is of great interest for telecommunications applications and has potential advantages over other characterisation techniques due to its ability to recover the Wigner distribution directly. However, the difficulty of applying the necessary temporal phase modulation currently limits its applicability to pulses above about 200 fs.

2.3.8   Summary — one-dimensional ultrashort pulse metrology

Linear time-stationary optics and slow photodetectors can measure the spectral intensity of a single pulse and, via interferometry, the spectral phase differences between pulses. Nonlinear correlations provide information on the pulse duration. There are three broad approaches to self-referenced complete characterisation: spectrography, shearing interferometry, and tomography; only the first two are generally suitable for pulses below 200 fs. Whilst spectrography has several advantages, such as experimental simplicity, self-referenced interferometry is the method of choice in this dissertation because of its compact one-dimensional encoding.

2.4   Extending ultrashort metrology to the space-time domain

2.4.1   Monochromatic spatial phase measurement

Several mature methods exist for obtaining the spatially resolved amplitude and phase of a monochromatic beam. These may be employed to obtain the spatial phase profile of a factorable ultrashort pulse — the factor $s(x, y)$ in (2.15) — simply by passing the pulse through a narrowband spectral filter. They also form the basis of many extensions of one-dimensional ultrashort characterisation methods to the spatial domain. This section provides a brief overview.

It is useful to compare the spatial and temporal problems. Section 2.3.2 discussed how the principal difficulties in ultrashort temporal characterisation were i) that photodetectors are slow integrators, reporting only the total incident energy, and ii) common optical elements are linear and time-stationary, so that only the spectral intensity could be detected. Spatial metrology shares the first problem, but not the second — operations which are nonstationary in the transverse po-
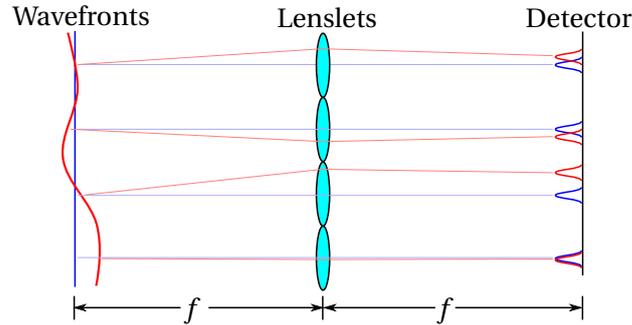
Figure 2.16: Shack-Hartmann wavefront sensor. Measurement of a flat wavefront (blue) and a distorted wavefront (red) is shown.

sition or wavenumber are trivially accomplished. Every single pixel on a detector array performs a space-nonstationary selection, whilst a spatial filter selects one particular wavenumber. This changes the game somewhat, allowing the use of techniques whose temporal analogue would be completely impractical.

2.4.1.1 Shack-Hartmann wavefront sensor

The Shack-Hartmann wavefront sensor [227–229], depicted in Fig. 2.16, is widely used. An array of lenslets each sample a small portion of the beam, producing an array of foci on a detector placed in the common focal plane. The foci are displaced according to the local wavefront tilt $f k_x / k$.

The combination of a Shack-Hartmann sensor with FROG to measure a space-time factorable pulse is called Shackled-FROG [230].

2.4.1.2 Interferometry

Interferometry is also widely used to obtain the spatial phase difference between two fields. The three main applications of this information mentioned in the spectral phase context (section 2.3.3.2) also apply. Comparison of the phase of a beam before and after the application of an optical instrument, such as a lens, provides the phase of the transfer function of the instrument, and is the basis of *optical testing*. Where the phase difference of the beams is related to the phase of an unknown beam by some invertible operation, then self-referenced phase measurement is possible. Finally, if the phase of one of the beams is known then the other may be inferred.

Many of the spatial interferometric methods used in this thesis can be traced back to optical
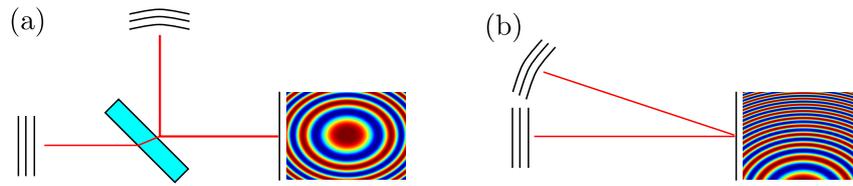
Figure 2.17: Spatial interferometry with (a) and without (b) a carrier, and the resulting fringe patterns. The wavefronts of the two beams are indicated; one is flat and the other is spherical.

shop testing methods developed in the 1950s. For convenient evaluation of optical aberrations in the workshop, the early devices were often designed to produce fringe patterns that were easily interpretable by eye. An example is shown in Fig. 2.17(a). The spherical wavefront is interfered co-linearly with a flat wavefront, and the resulting fringe pattern readily reveals the difference in their curvature. However, the sign ambiguity discussed in the context of spectral interferometry section 2.3.3.2, applies equally to the spatial case. In Fig. 2.17(a) the sign of the curvature cannot be determined, in the same way as one cannot distinguish a depression from a hill on a topographic map using the contour lines only.

Analogous to the spectral case, unambiguous phase recovery can be performed in several different ways. Multiple-trace methods require the phase of one of the arms to be modulated, enabling both the in-phase (cosine) and quadrature (sine) components to be obtained. Here, I concentrate on the only known single-trace method, which employs a rapidly varying carrier, in the form of a tilt between the beams. The tilt is analogous to the time delay in Fourier-transform spectral interferometry. For beams inclined at an angle θ in the yz plane incident on a detector perpendicular to the z-axis, the relative phase is $\theta k y$. The local fringe phase is therefore $\phi_2(x, y) - \phi_1(x, y) + \theta k y$. Before the development of Fourier-transform interferometry [92], the fringe phase was extracted by a peak searching algorithm [231]. However, this has been almost completely superseded by Fourier-transform spatial interferometry, which uses the same algorithm as the spectral case described in section 2.3.3.2 except that a two-dimensional Fourier transform is used.

One difference between spectral and spatial interferometry is the relative ease with which a well-characterised reference may be produced in the latter using spatial filtering. A pinhole may

be used, which projects the beam onto an approximation of a spatial Dirac delta function. Alternatively, a single-mode fibre projects the beam onto the fundamental mode of the fibre. From these well-defined modes, diffraction produces a well-characterised spreading of the beam to any desired spatial extent. A specific implementation of this is *point diffraction interferometry* [232].

2.4.1.3 Spatial shearing interferometers

Spatial shearing interferometry involves applying a small geometric transform to a beam, and then interfering the transformed beam with the original. The measured phase is a finite difference, which is commonly treated as a derivative along the shear direction. Concatenation or integration, appropriately generalised to two dimensions, then recovers the phase.

Lateral shearing interferometry (LSI) [233–237] is the most common form of spatial shearing interferometry, and is an exact analogue of spectral shearing interferometry. The beam is interfered with a displaced version of itself, yielding the finite difference $\phi(x+X, y) - \phi(x, y) \approx X \frac{\partial \phi}{\partial x}$, where for simplicity I take the shear to be along the x-axis. Integration along the x-axis gives $\phi(x, y) + f(y)$, where $f(y)$ is an unknown function resulting from the constant of integration of each line. To uniquely define the phase, two linearly independent shears are therefore needed [238].

Methods of implementing the shear include a Mach-Zehnder [233, 238] or Michelson interferometer [235] with displaced output beams, a parallel or slightly wedged plate [234], or a diffraction grating [239]. Implementations have been proposed using atom interferometry [240].

Radial shearing interferometry. In LSI, the interferogram does not cover the whole area of the wavefront. Furthermore two shears are required for a complete characterisation. These limitations led to the development of radial shearing interferometers, in which the beams being interfered are replicas of the original with a slightly different magnification. In a circular co-ordinate system (r, θ), phase difference is therefore $\phi(mr, \theta) - \phi(r, \phi) \approx (m-1)r \frac{\partial \phi}{\partial r}$. Excluding difficulties at the origin, where the shear drops to zero, the two-dimensional phase may be uniquely reconstructed by integrating along radial lines, since each share the origin at a common reference point. Implementations involve placing a telescope in either a Mach-Zehnder or Sagnac interferometer [241–243].

Rotational shearing interferometry [244] can also be performed, yielding the azimuthal phase derivative. A continuously adjustable rotation is provided by a Dove prism.

2.4.2 Spatially resolved temporal measurements

A limited form of spatio-temporal characterisation involves performing many one-dimensional temporal characterisations on different regions of the beam. One may simply scan a pinhole (with significant loss of sensitivity) in the transverse plane, performing temporal characterisations on the spatially selected region. Alternatively, if the dataset for each temporal characterisation is one-dimensional, the other dimension may provide spatial resolution, enabling temporal characterisations to be performed in parallel at each spatial point. Spectral shearing interferometry is well suited for this context [58, 59, 196, 197].

The limitation of all such methods is that the ambiguities of the temporal reconstruction become unknown functions of space. For example, in self-referenced methods the near-universal absolute phase ambiguity means that the recovered phase $\phi(x, y, \omega)$ contains an unknown additive term $f(x, y)$, providing no information on spatial wavefronts and preventing numerical propagation of the pulse. The arrival time ambiguity prevents the diagnosis of pulse-front tilt. Specific examples of this approach include spatially resolved crosscorrelations for the characterisation of multi-dimensional pulse shapers [72, 245], and SPIDER measurements revealing position-dependent duration and temporal profiles of few-cycle pulses obtained through filamentation [66].

Another approach, free from the aforementioned ambiguities, is spatially encoded arrangement TADPOLE (SEA-TADPOLE) [246–248]. The temporal measurement is externally referenced FTSI, which is free of ambiguities provided the reference contains every frequency of the unknown pulse. SEA-TADPOLE therefore represents a complete but nonself-referenced spatio-temporal characterisation. A schematic is shown in Fig. 2.18. The reference pulse, which must also be characterised by a separate device, is coupled into a single mode fibre. The unknown pulse is also coupled into a separate single-mode fibre. However, the mode area of the fibre is much smaller than the beam size. The fibre therefore selects a point in the beam, which can be scanned in the trans-
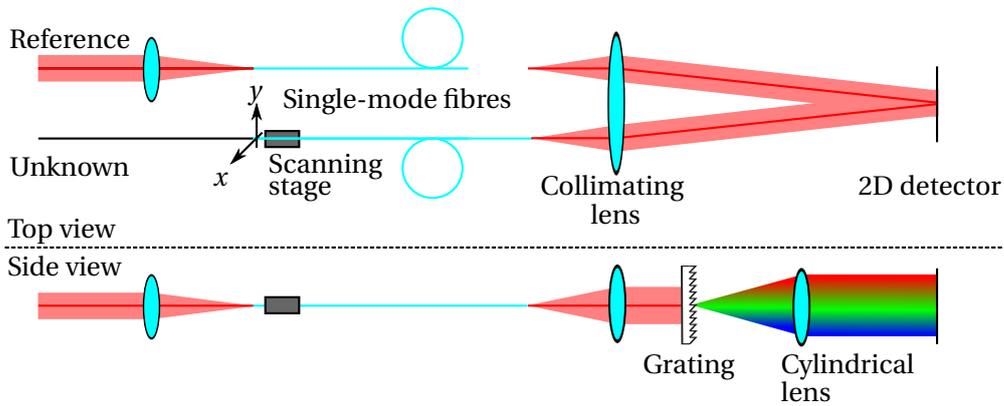
Figure 2.18: SEA-TADPOLE apparatus.

verse plane. The output of the fibres is directed to a spatially resolved spectrometer, producing spatial fringes which encode the relative phase of the reference and unknown pulses as described in section 2.3.3.2. Interferograms are recorded from a sampling of points in the unknown beam. Some specific issues with SEA-TADPOLE are calibrating the dispersion and the group delay difference of the two arms. For example, during the transverse scan, erroneous longitudinal wobble of the stage will introduce a spatially varying error in the group delay of the reconstructed pulse.

2.4.3 Incomplete spatial characterisations

Several modifications to the intensity autocorrelator provide limited space-time coupling information. A spatio-temporal intensity autocorrelator, in which the beams are combined with a lateral displacement and a time delay, provides the shape of the pulse front [249]. The pulse-front tilt may also be inferred using a single-shot autocorrelator which flips the pulse in both the horizontal and vertical planes [250].

Spatial chirp results in a sheared trace in single-shot FROG implementations [251], and can therefore be diagnosed with such devices.

2.4.4 Spectrally resolved spatial phase measurement

2.4.4.1 Spatio-spectral interferometry

Analogous to spectral interferometry (section 2.3.3.2) and spatial interferometry (section 2.4.1.2), the phase difference between two pulses may be obtained on a combined space-frequency do-

main.

As mentioned in the discussion of the purely spatial case (section 2.4.1.2), a well-characterised spatial reference may be produced by spatial filtering. Applied to a space-time coupled pulse, this produces a factorable pulse with well-characterised spatial profile. Temporal characterisation of this factorable pulse yields a fully characterised reference, against which the phase of an unknown pulse can be measured using spatio-spectral interferometry [67, 252]. There are however, two caveats to this approach.

- The spectrum of the factorable pulse must contain all the frequencies present in the unknown pulse. Using linear time-stationary optics, it impossible to produce a spatial filtering system which satisfies this requirement for all pulses. As a concrete example, first consider a beam with spatial chirp. A pinhole in the collimated beam will block certain frequencies, so the filtering must be performed in the Fourier domain, using a pinhole placed at the focus of a lens. If the pulse is now changed to exhibit some angular dispersion, then in the focal plane the frequencies will be dispersed and some will be blocked by the pinhole.

- The spatial filtering must be more aggressive than in the purely spatial case. In the spatial case, one requires a beam which approximately resembles a plane wave. However, the wavevector of this plane wave is not important, since in externally referenced interferometry one is almost always unconcerned by the tilt between the beams, or equivalently the linear component of the phase. In the spatio-spectral case, one requires each frequency of the reference to have the same wavevector, or else erroneous angular dispersion will be diagnosed. Adhering to this requirement reduces the throughput of the spatial filter.

There have been two main implementations of spatio-spectral interferometry.

Fourier-transform interferometry. Many implementations use some form of two-dimensional spectrometer, in which frequency is mapped to one axis of a two-dimensional detector array leaving the other free to resolve one spatial dimension. In this discussion, I shall associate the x-axis with frequency. A particular example is an *imaging spectrometer*, which performs re-imaging of its entrance slit perpendicular to the plane of dispersion. In any case, the signal is proportional to

$|E(y,\omega)|^2$. Structure in the x-axis may be simply averaged over, or a specific x co-ordinate may be selected for examination. In many cases, for instance that of cylindrical symmetry, the loss of one spatial dimension is perfectly acceptable.

The most common approach to resolving the sign-ambiguity problem is via a rapidly-varying carrier, which in general is a combination of time delay $\omega\tau$ which produces a spectral carrier, and tilt $\theta k y$ which gives a spatial carrier. The carrier amplitudes $(\tau, \theta k)$ locate the sidebands in the (t, k_y) Fourier domain, and judicious selection of their values can enable optimal usage of the spectrometer's resolution. Combined with the Fourier-filtering phase extraction procedure [92], *spatio-spectral Fourier-transform interferometry* [253] has been extensively used to characterise pulse propagation through plasmas [254], high-numerical aperture lenses [255, 256], and pulse shapers [71, 96].

If required, there are several ways of simultaneously resolving two spatial dimensions as well as frequency. *Integral field spectroscopy* [257, 258] is the name given to this task in astronomy, but no applications to ultrashort pulse characterisation have been demonstrated to date. A different approach is STRIPED-FISH [259, 260], a form of Fourier-transform interferometry with a spatial carrier. The test and reference beams pass through a coarse grating, which produces many diffracted orders. The orders then pass through an interference filter. Since each order has a slightly different wave-vector and the passband of the filter is angularly dependent, the passband frequency seen by each order is different. Each order therefore contains a quasi-monochromatic spatially encoded interferogram, and these are simultaneously recorded on a two-dimensional detector. Combining the interferograms yields a three-dimensional dataset, albeit sampled coarsely in frequency.

Time-domain interferometry. The field crosscorrelation, given by equation 2.23, may be obtained on a spatially resolved detector, enabling direct measurement of the interferometric overlap $E_2(x,y,\omega)E_1^*(x,y,\omega)$ [76, 261].

2.4.4.2   Self-referenced spectrally resolved wavefront characterisation

The methods of this section return the two or three-dimensional spatio-spectral phase up to an unknown function of frequency $f(\omega)$, avoiding the limitations inherent in producing a spatially

uniform reference that were discussed in the previous section. To complete the spatio-temporal characterisation, one or more spectral phase measurements at specific spatial locations are required. In fact, there are several advantages to measuring the spectral phase at all spatial points using one of the methods described in section 2.4.2, and this is one motivation for the work presented in chapter 5 of this dissertation.

A Shack-Hartmann wavefront sensor may be placed in the detector plane of an imaging spectrometer [262], a highly practical solution because such sensors are a mature and robust technology. Care must be taken to calibrate out the quadratic spatial phase caused by the imaging system. In ref. [262] this approach was combined with FROG to perform complete characterisation in two dimensions.

Frequency-resolved tomographic imaging is performed by making intensity measurements in the near and far field. For each frequency, the spatial phase profile is then retrieved using the Gerchberg-Saxton algorithm [215, 216]. This has been combined with FROG to form CROAK [263], a complete characterisation method.

Spectrally resolved lateral shearing interferometry is the method of choice in this dissertation. The principle is identical to the spatial case discussed in section 2.4.1.3, except that an imaging spectrometer is used so that one of the spatial dimensions is replaced by frequency. The beams are combined with a time delay and a tilt, and Fourier-transform filtering is used to extract the fringe phase [57, 264]. As with all forms of one-dimensional shearing interferometry, there is a relative phase ambiguity across null regions.

The combination of lateral and spectral shearing interferometry is space-time SPIDER (ST-SPIDER) [58, 59, 265].

2.4.5   Summary

The most common approach to extending ultrashort phase measurement to the spatial domain is using interferometry, although the use of tomography and wavefront sensors has also been demonstrated. There are only a handful of spatio-temporal characterisation methods which can be considered self-referenced and complete. Many of these involve using spatial filtering or selec-

tion to produce a spatially uniform reference pulse which must contain all frequencies present in the test pulse.

2.5   Outline of this dissertation

The work presented in chapters 3–5 can now be placed in context.

All forms of shearing interferometry (sections 2.3.6 and 2.4.1.3) suffer from relative phase ambiguities across *spectral nulls* — regions of zero intensity. In chapter 3, I present a means of overcoming this ambiguity: an algorithm for combining interferograms taken at different shears. A small shear may sample fine spectral details, achieving high resolution, simultaneously with a large shear which obtains the phase difference across regions of zero intensity. I also show that the use of multiple shears improves the precision even if spectral nulls are absent.

Chapter 4 presents two experimental implementations of multiple spectral shearing interferometry. The first uses a SEA-SPIDER to acquire multiple shears in sequence. The second uses a newly developed implementation of CAR-SPIDER to acquire a range of shears simultaneously.

Section 2.4 showed that only a few methods for complete spatio-temporal characterisation exist. Of these, most rely on the production of a reference pulse with properties that may not be achieveable in practice. Most approaches are also the combination of two separate instruments — one for the spatial and one for the temporal characterisation, with an associated increase in complexity and number of components. In chapter 5, I present a method for two-dimensional spatio-temporal pulse measurement which addresses both of these concerns. It is based on combined lateral and spectral shearing interferometry but performs both of these using a single interferometer.

3   Phase reconstruction algorithm for multiple spectral shearing interferometry

This chapter presents a method for reconstructing the spectral phase of an optical pulse using several spectral shearing interferometry (SSI) measurements taken using different shears. This improves the precision above that which can be achieved using a single shear. The improvement increases with the degree of modulation of the spectral intensity. In the extreme case of a spectrum with one or more islands of zero spectral intensity, or *spectral nulls*, the use of multiple shears can enable determination of the relative spectral phase across the null, a property to which a single shear reconstruction would be completely insensitive. A limitation of the method is that the null must be smaller than the largest nonzero spectral region. It is, however, a significant improvement over the single shear case which would require the null to be smaller than the spectral shear, a somewhat contradictory requirement since the spectral shear defines the resolution of the measurement. As such, the use of multiple shears offers partial resolution of the relative phase ambiguity of SSI.

The ambiguity in the relative phase of separate spectral components in SSI is well known, and the possibility of using multiple shears to recover it has been previously suggested [81, 186, 266] but not implemented.

This chapter develops the multiple shear spectral interferometry algorithm, which accepts as its input the extracted interferometric phase differences. This leads into chapter 4, which presents an arrangement for acquiring such data sequentially and provides a first proof-of-principle experimental demonstration of multiple SSI. Chapter 4 also presents an arrangement for acquiring multiple shears simultaneously on a two-dimensional detector, and demonstrates the resolution of the relative phase ambiguity. The work in this chapter formed part of the publications "High precision self-referenced phase retrieval of complex pulses with multiple-shearing spectral interferometry" [267] and "Resolution of the relative phase ambiguity in spectral shearing interferometry of ultrashort pulses" [267].

This chapter is organized as follows: section 3.1 presents some qualitative reasons for devel-

oping a multiple-shear algorithm, and section 3.2 presents a quantitative analysis of the preci-
sion of SSI. Section 3.3 discusses the nature of the raw data, and then section 3.4 addresses sam-
pling issues which immediately arise when multiple shears are being considered. Section 3.5 then
presents the main algorithim. Section 3.6 discusses the implications for alleviating the relative
phase ambiguity. Sections 3.7 and 3.8 together present a general argument that using multiple
shears is always beneficial for the precision; the former considers the cost in terms of signal-to-
noise ratio (SNR) of acquiring multiple shears, whilst the latter considers the benefits. Section 3.9
gives some numerical examples and section 3.10 presents a summary and outlook.

3.1   Motivation

This section presents several broad qualitative arguments for the use of multiple shearing spectral
interferometry.

All else being equal, the noise sensitivity of SSI — that is, the precision loss due to a given
detector noise amplitude — increases with the number of sampling points across the spectrum.
This introduces challenges of precision when characterising complex pulses, since in almost all
definitions complexity is quantified by the time-bandwidth product which is proportional to the
required number of sampling points. Practically speaking, complex pulses arise frequently in ul-
trafast optics, such as in coherent control [7], supercontinuum generation [268], filamentation
[269], telecommunications [270] and micromachining [271]. Two factors are at play in the loss of
precision: first, increasing the number of sampling points means increasing the effective resolu-
tion of the spectrometer, decreasing the amount of light available and inherently reducing the pre-
cision. In practice, this is often accomplished by increasing the passband of the Fourier-domain
filter used in the processing. Second, the concatenation introduces an error at each sampling
point; if the errors are independent (as is shown to be the usual case) then they accumulate in the
fashion of a random walk. It is plausible that a scheme involving multiple shears could bypass the
second of the aforementioned channels through which pulse complexity reduces precision; that
is, the random-walk like concatenation of errors. A small shear could resolve fine spectral details.
Errors in the small shear measurement could be prevented from accumulating by the incorpo-

ration of information from a larger shear, which links disparate frequencies directly without any concatenation.

The noise sensitivity of SSI also increases with the degree of spectral intensity modulation of the spectrum. In the extreme case of a spectral null, SSI develops an ambiguity in the relative phase across the null, as discussed in section 2.3.6.1. However, ultrashort pulses with highly modulated spectra arise frequently in ultrafast optics. Examples include the output of hollow-core fiber compression systems [60], soliton molecules [272], shaped two-color pulses for difference frequency generation [273], supercontinua generated in photonic crystal fibers [268] and shaped pulse sequences for spectroscopy and coherent control [274]. Furthermore, the relative phase between spectrally disjoint subpulses is usually physically significant. For example, the output of pulse compression systems based on self-phase modulation often consist of the superposition of temporally overlapped but spectrally disjoint subpulses, and the resulting temporal intensity is critically dependent on the relative phase of the subpulses. Also, the relative phase of spectrally disjoint subpulses can be important in pump-probe and coherent control experiments [275–277]. A multiple shear scheme could potentially alleviate the problems associated with spectral intensity modulation in the single shear case. A small shear would resolve spectral details, whilst a large shear could cross spectral nulls, precisely linking the spectral phase between frequencies of significant spectral intensity. The idea is depicted in Fig. 3.1.

Another motivation is the redundancy inherent in a multiple shear dataset, which should also offer an automatic self-consistency check, useful for validation of the measurement and identification of an incorrect setup. This is an advance upon the comparison of several single-shear reconstructions taken at different shears, an approach that is already available to a conscientious experimenter with a variable-shear apparatus.

Of course, in motivating multiple SSI one must justify the use of SSI in the first place, particularly over less experimentally complex spectrographic methods. Besides the features of SSI described in section 2.3.6, there are a number of specific reasons one may prefer SSI over spectography for complex and/or highly modulated pulses. The most complex ultrashort pulse measurements have been done with crosscorrelation FROG (XFROG) [175], which has elucidated the

3. PHASE RECONSTRUCTION ALGORITHM FOR MULTIPLE SPECTRAL SHEARING INTEROMETRY
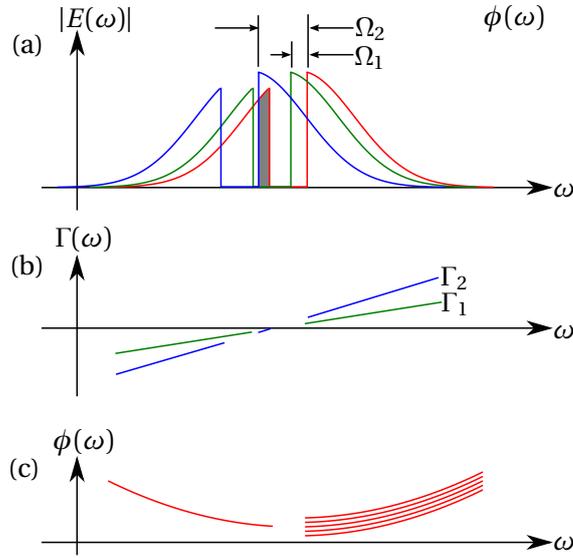
Figure 3.1: Using multiple shears to cross a spectral null; (a) original (red) and two replicas (green, blue) of a pulse with quadratic spectral phase and a spectral null; the shear Ω_1 of the green replica is less than the width of the null, whilst the shear Ω_2 of the blue replica is greater. The overlap of the blue replica and the original is shown in grey. The pulse has a quadratic spectral phase (not shown). (b) Phase difference between the replicas and original; the phase difference between the lobes is defined for the second replica. (c) Reconstructed phase; the relative phase of the two lobes is constrained by the blue-red overlap, allowing the correct reconstruction (strong red) to be chosen from amongst those consistent with the single shear data (faint red).

dynamics of supercontinuum generation in photonic crystal fibre [64]. However a limitation of these measurements is their multi-shot nature, necessitated by weak signal levels and the need to angle-dither the nonlinear crystal to obtain sufficient phase-matching bandwidth. Efforts to alleviate these disadvantages have limited sensitivity because of the need to use third order processes to obtain sufficient bandwidth [278]. All spectrographic methods require N well-resolved sampling points along both the time delay and frequency axes, and the system response must be well calibrated over both of these domains. This presents an experimental challenge, particularly to single-shot measurements. The algorithmic complexity of retrieving the pulse from the N^2 data points scales as $N^2 \log N$ leading to reconstruction times measured in hours.

Another concern with spectrographic methods is the possibility of ambiguities pertaining to, or caused by, the temporal structure of the pulse, which do not occur in SSI. For example, a spectrogram only encodes the relative phase of separate temporal components if the gate pulse

is longer than the separation. If the components are also separated spectrally, the spectral width of the gate must also span this separation, leading to a requirement for a nontransform-limited gate pulse [81].

3.2 Quantitative noise analysis for spectral shearing interferometry

Since the motivation for multiple shears is largely based on considerations of uncertainty and precision, an analysis of the noise sensitivity of SSI is warranted. This complements previous studies [279, 279], which have been largely numerical. The ability to describe and quantify uncertainty is an important feature of any measurement technique. For example, in spectrographic methods one may perform "boot-strapping" [280] — inferring the uncertainty from an ensemble of reconstructions performed using random subsets of the data. The direct connection between the raw data and the reconstruction in SSI means it is more amenable to a traditional error analysis [186].

A related issue is the representation of uncertainty in ultrashort pulse measurement. For exact ambiguities, such as the absolute and relative phase ambiguities and the direction-of-time ambiguity in second-harmonic generation FROG (SHG-FROG), one normally chooses the unknown parameter in some arbitrary way. For example, one usually sets the absolute phase and arrival time to zero. The ambiguity is not represented on the plot, and must be noted, explicitly or implicitly. However, when the ambiguity is approximate, such as the relative phase betweens regions of low spectral intensity, this is somewhat misleading, as the uncertainty may range from zero up to $\pm\pi$. Simply plotting the pulse corresponding to the mode or mean of the distribution and describing the uncertainty in the text would lead to excessively verbose descriptions as the pulse complexity increased. Instead, one may consider error bars. However, the errors produced by ambiguities are highly correlated. For example, in representing the relative phase ambiguity between two spectral lobes, by fixing the absolute phase of one spectral lobe to some arbitrary value, the uncertainty in the phase of the other spectral lobe is described by a single parameter.

For these reasons, uncertainty must be represented by a more expressive means than error bars. One proposal is to show an entire ensemble of reconstructions on a single plot [281]. Depending on the origin of the error in question, the ensemble may be generated using boot-strapping

3. PHASE RECONSTRUCTION ALGORITHM FOR MULTIPLE SPECTRAL SHEARING INTERFEROMETRY

[281], using different calibration or reconstruction parameters [282], or by simply taking many measurements.

Here I use an alternative representation which quantifies errors and their correlations caused by detector noise in SSI. It is based on the realisation that since SSI measures spectral phase differences, the uncertainty is well represented by an error matrix

$$Q(\omega_i, \omega_j) = \left\langle \left| \Delta \left[\phi(\omega_i) - \phi(\omega_j) \right] \right|^2 \right\rangle^{1/2} \tag{3.1}$$

where ω_i and ω_j are sampled frequencies, Δ denotes the fluctuation from the mean value and the angle brackets denote expectation over a large number of measurements. A low value of $Q(\omega_i, \omega_j)$ means that the phase difference between ω_i and ω_j is precisely constrained by the measurement. The error matrix is an intuitive depiction of uncertainty in SSI and some examples will be given in the experimental demonstrations in chapter 4. Although one may be ultimately concerned with some other measure of precision, such as the root-mean-square (RMS) field variation [283, 284] or the variation in the pulse duration, these quantities can be derived from knowledge of $Q(\omega_1, \omega_2)$ combined with knowledge of any uncertainty in the spectral intensity measurement.

For a single-shear reconstruction, the error matrix may be calculated using a Pythagorean summation of the errors. Concatenation gives the reconstructed phase difference between frequencies ω_j and ω_i as

$$\phi(\omega_j) - \phi(\omega_k) = \sum_{n=k}^{j-1} \Gamma(\omega_n). \tag{3.2}$$

The properties of the errors on $\Gamma(\omega)$ are derived in Appendix A. The main result is that the errors are independent and their standard deviation is given by $(A.12)$, repeated here:

$$\sigma_\Gamma^2(\omega) = \frac{\sigma_\xi^2}{2|\bar{D}(\omega)|^2} \tag{3.3}$$

where σ_ξ^2 is the variance of the noise in the filtered sideband, and $|\bar{D}(\omega)| = |E(\omega + \Omega)E^*(\omega)|$ is the amplitude of the filtered sideband in the absence of noise. The errors add in Pythagorean fashion

to give

$$Q^2(\omega_i, \omega_j) = \frac{\sigma_\xi^2}{2} \sum_{n=k}^{j-1} \frac{1}{|D(\omega_n)|^2}. \tag{3.4}$$

Equation (3.4) makes manifest the role of the amplitude of the overlap. If $E(\omega_n) = 0$ for some ω_n between ω_i and ω_j, then $Q^2(\omega_i, \omega_j)$ will diverge, producing the relative phase ambiguity.

Later in this chapter I will show that, in a multiple shear reconstruction, the error matrix may also be conveniently obtained.

3.3 Nature of the input data

This section discusses the features of the input data for the multiple shear algorithm. Regardless of the experimental implementation and details of the numerical phase extraction set, the input to the algorithm is a set of interferometric products $\{D_k(\omega)\}$, where

$$D_k(\omega) = E(\omega + \Omega_k) E^*(\omega) e^{i\eta_k} \tag{3.5}$$

are taken at shears $\{\Omega_k\}$[1] The phase factor $e^{i\eta_k}$ is the absolute interferogram phase. In a single-shear reconstruction, it determines the group delay of the reconstructed pulse. In implementations based on nonlinear optics, absolute interferogram phase is affected by the spectral phase of the ancillae used for the upconversion, which may have been derived from the unknown pulse (in most implementations) or from an external reference (in modified spectral phase interferometry for direct electric-field reconstruction (SPIDER) arrangements [81, 285, 286]). It is also affected by the optical path difference between the arms of the interferometer. This is normally unproblematic. In the multi-shear case, care must be taken to ensure the constants are consistent between shears i.e. the various shear measurements must imply a pulse located at the same position in the temporal window. To avoid burdening the algorithm with any dependency on the value of the absolute phase, I shall take it as unknown and random, and by design, ensure that its value does not affect the result.

I have deliberately written (3.5) in terms of complex numbers, rather than the implied differ-

---

[1] I use curly braces to refer to a set of quantities.

ence equation between the phases

$$\operatorname{Arg} D_k(\omega) = \Gamma_k(\omega) + \eta_k - 2\pi u_k(\omega), \tag{3.6}$$

where $\Gamma_k(\omega) = \phi(\omega + \Omega_k) - \phi(\omega)$ are the usual phase differences. Here, $u_k(\omega) \in \mathbb{Z}$ are the unwrapping integers required because the principal value $\operatorname{Arg} D_k(\omega)$[1] is only recoverable over the domain $[-\pi, \pi)$. In a single-shear reconstruction of a pulse whose duration is lower than the Nyquist-dictated upper bound, the phase differences are also within this interval $(-\pi, \pi]$, so that the principal value of the argument of $D(\omega)$ can be used for $\Gamma(\omega)$. In certain cases, correct retrievals can even be achieved for pulses which exceed the temporal window if $D(\omega)$ is unwrapped by removing 2π discontinuities. However, in a multiple shear reconstruction, neither approach will work in general: larger shears can sample phase differences of absolute value much greater than π, yet performing an unwrapping procedure on each shear individually may produce mutually inconsistent data. Like the absolute phase issue, this is another subtlety of multiple shear reconstructions which must be taken into account by the algorithm.

3.4   Sampling and uniqueness of solutions

I shall temporarily set aside the absolute phase and unwrapping issues, and develop some other general features of the algorithm.

I assume that the products $\{D_k(\omega)\}$ are well sampled by the spectrometer resolution, which means that they can be resampled via interpolation to any desired set of frequencies. Here I derive a linear least-squares method which is noniterative, efficient and a logical generalization of the concatenation algorithm. Like the concatenation algorithm the unknown phase is sampled at frequencies $\{\omega_n\}$, whose regular spacing Ω determines the temporal window of the unknown pulse. The strategy is to resample the $\{D_k(\omega)\}$ to these points, take the argument to yield the phase differences $\{\Gamma_k(\omega)\}$ and then write the difference equations (3.6) at each frequency and shear to produce a set of equations which are solved simultaneously. To express the difference equations

---

[1]I use Arg to denote the principal value, within $[-\pi, \pi)$, and arg to denote the local analytic continuation so that the small fluctuations caused by noise do not cause 2π phase jumps in any evaluation of the statistics.

simply in terms of the unknown phase at the sampled points, one requires that for every combination of n and k, $\omega_n + \Omega_k$ must be a sampled frequency ω_m for some m — that is, the frequencies connected by every equation must be sampled frequencies. Otherwise, expressing $\phi(\omega_n + \Omega_k)$ in (3.6) would require interpolation, complicating the expressions. The shears must therefore be an integer multiple of the sampling rate, so that one may write $\Omega_k = C_k \Omega$. Imposing this restriction on the shears permits the use of simple linear difference equations, of the form

$$\phi_{n+C_k} - \phi_n = \Gamma_{k,n} \tag{3.7}$$

for all n and k where $\phi_n = \phi(\omega_n)$ and $\Gamma_{k,n} = \Gamma_k(\omega_n)$. One also imposes $\phi_0 = 0$ as a boundary condition since the absolute phase is not defined.

I now consider the uniqueness of the solution to (3.7). Although not used in practice, a simple solution method is to start at the boundary condition ω_0, and to use superposition to "visit" other sampling points. One may determine which of the other sampling points can be visited by applying series of jumps, forwards or backwards, of the different shear sizes. If one applies a_1 jumps of shear Ω_1, a_2 jumps of shear Ω_2, and so on (the $\{a_k\}$ may be positive or negative) then one will have moved a total displacement of

$$\omega - \omega_0 = a_1 \Omega_1 + a_2 \Omega_2 + \dots \tag{3.8}$$

Therefore, one can reach frequency ω_n if and only if there exist integer solutions $a_1, a_2, \dots$ to

$$a_1 C_1 + a_2 C_2 + \dots = n. \tag{3.9}$$

By Bézout's Identity [287], this occurs if and only if n is an integer multiple of the greatest common divisor of $\{C_k\}$. If this condition is not satisfied, then there is no way that any linear combination of the difference equations can refer ω_n to ω_0, and hence ϕ_n is not determined by the experimental data. Therefore, a unique solution for all n requires that the $\{C_k\}$ be relatively prime. Examples of the two cases are illustrated in Fig. 3.2. I proceed assuming the shears are relatively prime, in which case one has N unknowns and approximately NM equations, where M is the number of
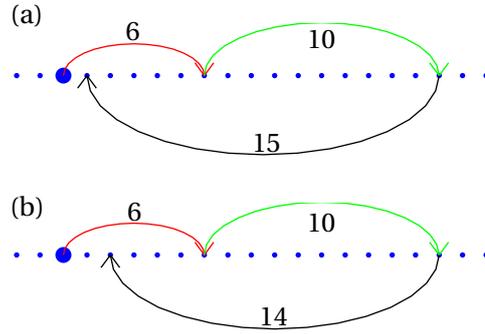
Figure 3.2: (Color online) Combining multiply-sized shears to achieve a unit spectral displacement along the sampling points (small blue dots). (a) Shears of size (relative to the sampling rate) 6, 10, and 15 can be combined to form a unit displacement from the starting point (large blue dot) because the sizes are relatively prime. (b) Shears of size 6, 10, and 14 have a common factor of 2, and no integer combination will produce a unit displacement.

shears.

3.5 Main algorithm

The problem consists of a set of linear equations in the discretized unknown variables $\{\phi_n\}$, $\{\eta_k\}$, and $\{u_{k,n}\}$. Therefore, standard methods for least-squares solution of an overdetermined set of linear equations may seem like an attractive proposition. However because the unwrapping numbers are integers, such approaches cannot be directly applied. Furthermore, since the unknown variable η_k appears in all equations involving the k-th shear, the problem is strongly coupled, so that the sparsity of the equations cannot be easily exploited. Instead, I have developed a heuristic method, based on the principle that the absolute phases and unwrapping numbers of large shears can be efficiently and accurately obtained from a reconstruction performed using smaller shears. This imposes one restriction: one of the shears, which for clarity I take as the first, must be equal to the sampling rate. Henceforth $\Omega_1 = \Omega$.

The shears are sorted in increasing order of absolute value and the algorithm loops through each shear in turn. For each shear, there are two steps: preprocessing, in which the absolute phase and unwrapping integers are determined, and generalized concatenation, in which the phase is reconstructed incorporating the information from all the shears encountered thus far. After the final shear has been treated, the generalized concatenation step has incorporated the information

from all the shears and the reconstruction is complete.

Throughout the algorithm, one must keep track not only of the reconstructed phase at each step, but also the *covariance* of the reconstructed phase. This summarises one's knowledge of the uncertainty in reconstructed phase samples and is essential for the preprocessing step. The uncertainty is assumed to arise from detector noise, which is represented as an additive contribution of variance σ_ξ^2 to the measured interferometric products $D_k(\omega)$. The value of σ_ξ^2 may be estimated from the raw data in a number of ways, depending on the encoding method used in the experimental implementation. Appendix A.2 presents one method, suitable for Fourier-transform interferometry. Using σ_ξ^2, the covariance is calculated through elementary, if somewhat involved, error propagation analysis.

3.5.1   Preprocessing

The preprocessing step for shear k computes the absolute phase η_k and the unwrapping integers $\{u_{k,n}\}$ for that shear. For $k = 1$, these are chosen in a similar manner to the single-shear case. For $k \geq 2$, the preprocessing step of shear k relies on a reconstruction performed using only the subset of shears $1, 2, \ldots, k-1$, which is denoted by $\phi_n^{(k-1)}$. It requires as its input the covariance of this reconstruction: that is

$$V_{n,m}^{(k-1)} = \langle \Delta \phi_n^{(k-1)} \Delta \phi_m^{(k-1)} \rangle \tag{3.10}$$

where Δ denotes fluctuations about the mean value.

Both of these inputs — the reconstructed phase and its covariance — are computed in the previous generalized concatenation step. I therefore take these as given in the description of this step.

3.5.1.1   First shear

The first shear is the sampling rate $\Omega_1 = \Omega$. Its absolute phase is arbitrary — the other absolute phases will be set accordingly. It is now convenient to set it to the intensity-weighted average

phase of the first interferogram

$$\eta_1 = \mathrm{Arg} \sum_n |D_1(\omega_n)|^2 \exp[i\,\mathrm{Arg}\,D_1(\omega_n)]. \tag{3.11}$$

Subtracting this absolute phase sets the group delay of the reconstructed pulse, weighted by the spectral intensity, to zero, and makes our procedure invariant against the absolute phase. The unwrapping integers of the first shear $\{u_{1,n}\}$ are chosen to remove 2π phase discontinuities using a standard unwrapping procedure. One can thus compute $\{\Gamma_{1,n}\}$ via (3.6).

3.5.1.2   Second and higher shears

From the previously reconstructed phase $\{\phi_n^{(k-1)}\}$, one obtains the *expected phase differences* $\{\widehat{\Gamma}_{k,n}\}$ for the present shear using

$$\widehat{\Gamma}_{k,n} = \phi_{n+C_k}^{(k-1)} - \phi_n^{(k-1)}. \tag{3.12}$$

One also computes the variance of the expected phase differences from the covariance

$$\sigma_{\widehat{\Gamma}_{k,n}}^2 = \left\langle \left| \Delta(\phi_{n+C_k}^{(k-1)} - \phi_n^{(k-1)}) \right|^2 \right\rangle \tag{3.13}$$

$$= V_{n,n}^{(k-1)} + V_{n+C_k,n+C_k}^{(k-1)} - 2V_{n,n+C_k}^{(k-1)}. \tag{3.14}$$

One now finds the absolute phase η_k that makes the measured $\{\Gamma_{k,n}\}$ most consistent with the expected $\{\widehat{\Gamma}_{k,n}\}$. Specifically, one solves

$$\widehat{\Gamma}_{k,n} - \mathrm{Arg}\,D_{k,n} = \eta_k \quad \mathrm{mod}\ 2\pi \tag{3.15}$$

for η_k. Since (3.15) represents a set of equations — one for each n — with only one unknown, η_k, it is massively overdetermined. Therefore one must perform some form of optimization; here I choose a least-squares method. One must also take into account the precision with which $\{\widehat{\Gamma}_{k,n}\}$ and $\{\mathrm{Arg}\,D_{k,n}\}$ are known. For example, $\widehat{\Gamma}_{k,n}$ will be essentially random if a region of low SNR — most commonly due to a spectral null — occurs between ω_n and $\omega_n + \Omega_k$, and has not yet been "bridged" by one of the smaller shears. This is simply the relative phase ambiguity at work. In this

situation, $\widehat{\Gamma}_{k,n}$ should play no role in determining η_k.

One can address this concern by scaling the equations represented by (3.15) according to the precision with which the left hand side is known. The least-squares optimization will thus ignore places where $\widehat{\Gamma}_{k,n}$ is unknown. The noise on the left-hand side of (3.15) is given by Pythagorean summation

$$\rho_{k,n}^2 = \sigma_{\widehat{\Gamma}_{k,n}}^2 + \frac{\sigma_\xi^2}{2|D_{k,n}|^2} \tag{3.16}$$

where I have used (3.3) for the noise on the argument of the interferometric product $D_{k,n}$.

To alleviate the difficulties of the modulo 2π in (3.15), I use complex exponentials and express the problem as a weighted least squares minimization

$$\min_{\eta_k} \sum_n \left| \exp(i\widehat{\Gamma}_{k,n} - i\operatorname{Arg} D_{k,n}) - \exp(i\eta_k) \right|^2 \rho_{k,n}^{-2} \tag{3.17}$$

which has solution

$$\eta_k = \operatorname{Arg} \sum_n \exp[i\operatorname{Arg} D_{k,n} - i\widehat{\Gamma}_{k,n}]\rho_{k,n}^{-2}. \tag{3.18}$$

I call this procedure the *registration* of the absolute phase. The success of this step depends on the existence of at least one frequency where the left-hand side of (3.15) is well-defined — that is, where $\rho_{k,n}$ is close to zero — and holds crucial implications for the presence of the relative phase ambiguity. This will be discussed in section 3.6 below.

The unwrapping can then be performed by choosing

$$u_{k,n} = \operatorname{round} \left[\frac{\widehat{\Gamma}_{k,n} - \operatorname{Arg} D_{k,n} + \eta_k}{2\pi} \right]. \tag{3.19}$$

This unwrapping step (3.19) will fail if the phase differences from the first shear are so noisy that the accumulated error over an interval equal to one of the larger shears exceeds π.

3.5.2 Generalized concatenation

The generalized concatenation step for shear k accepts as its input a set of measured phase differences for each shear up to and including k. These have been unwrapped and had their absolute

phases removed using (3.6), so that they may be directly related to the unknown phase via (3.7). This step outputs the reconstructed phase $\{\phi_n^{(k)}\}$ incorporating the information from all of these shears, as well as the associated covariance.

For $k = 1$, the phase reconstruction reduces to the standard single-shear concatenation procedure. Calculation of the covariance for this step is described in section 3.2. For $k > 1$, the concatenation procedure must be generalized for multiple shears. Of course, this generalization works for the single-shear case, and so for simplicity of programming it is unnecessary to include the single-shear procedure as a special case.

The equation set represented by (3.7) is linear and furthermore, the coupling of the unknowns is local — a pair of frequencies ω_i and ω_j are linked only if they differ by one of the shears. This suggests a sparse matrix representation leading to an efficient least-squares solution.

The unknown phase, sampled at ω_n, $n = 0, 1, \dots, N-1$, is written as a column vector $\bar{\phi}^{(k)} = \left[\phi_0^{(k)}, \phi_1^{(k)}, \dots, \phi_{N-1}^{(k)} \right]$ of length N. The post-processed phase differences, defined by (3.6), are written as a set of column vectors $\{ \bar{\Gamma}_j \}$, $j = 1, 2, \dots, k$, each of length $N - |C_j|$. These vectors are slightly shorter than $\bar{\phi}$ since the phase differences must not overrun the range of sampled frequencies. I introduce the shear matrices $\{ \mathbf{G}_j \}$, $j = 1, 2, \dots, k$ each of size $(N - |C_j|) \times N$, which, when right multiplied by $\bar{\phi}^{(k)}$, give the corresponding phase differences:

$$
\mathbf{G}_j = \begin{bmatrix}
\ddots & & & & \ddots & & & & & 0 \\
0 & -1 & 0 & \cdots & 0 & 1 & 0 & & & \\
& 0 & -1 & 0 & \cdots & 0 & 1 & 0 & & \\
& & 0 & -1 & 0 & \cdots & 0 & 1 & 0 & \\
0 & & & & \ddots & & & & \ddots &
\end{bmatrix}
\tag{3.20}
$$

so that each row consists of a -1 and a 1 separated by C_j columns. I also introduce the diagonal weights matrices $\{ \mathbf{W}_j \}$, each of size $(N - |C_j|) \times (N - |C_j|)$, with entries $(\mathbf{W}_j)_{n,n} = \sigma_{\arg D_{j,n}}^{-1}$, obtained by (3.3). Rewriting (3.7) in the present matrix formalism with the weights, one has

$$
\mathbf{W}_j \mathbf{G}_j \bar{\phi}^{(k)} = \mathbf{W}_j \bar{\Gamma}_j
\tag{3.21}
$$

for $j = 1, 2, \ldots, k$. I then vertically concatenate the equations for all the shears into a single set

$$\mathbf{B}^{(k)} \bar{\phi}^{(k)} = \bar{F}^{(k)} \tag{3.22}$$

where $\mathbf{B}^{(k)}$ and $\bar{F}^{(k)}$ are the vertical concatenations of $\{\mathbf{W}_j \mathbf{G}_j\}$ and $\{\mathbf{W}_j \bar{\Gamma}_j\}$ respectively, for all $j = 1, 2, \ldots, k$:

$$\mathbf{B}^{(k)} = \begin{bmatrix} \mathbf{W}_1 \mathbf{G}_1 \\ \mathbf{W}_2 \mathbf{G}_2 \\ \vdots \\ \mathbf{W}_k \mathbf{G}_k \end{bmatrix}, \bar{F}^{(k)} = \begin{bmatrix} \mathbf{W}_1 \bar{\Gamma}_1 \\ \mathbf{W}_2 \bar{\Gamma}_2 \\ \vdots \\ \mathbf{W}_k \bar{\Gamma}_k \end{bmatrix}. \tag{3.23}$$

The number of rows of $\mathbf{B}^{(k)}$ and $\bar{F}^{(k)}$ is equal to the total number of "observations"

$$P = KN - \sum_{j=1}^{k} |C_j|, \tag{3.24}$$

whilst $\mathbf{B}$ has N columns. Equation (3.22) can then be solved in the least-squares sense via a suitable linear algebra method. The normal equation matrix $(\mathbf{B}^{(k)})^T \mathbf{B}^{(k)}$ is sparse — an element nm is nonzero only if frequencies ω_n and ω_m are connected by one or more shears. It is therefore banded, with nonzero elements positioned above or below the main diagonal by no more than $\max\{|C_j|\}$. Equation (3.22) can therefore be solved by efficient LU factorisation and backsubstitution routines for banded matrices. For example, the LAPACK routines are DGBTRF and DGBTRS. For the numerical and experimental results presented below, I used MATLAB, achieving a computation time of 40 ms for a 1000-point inversion on a laptop computer.

Elementary linear algebra shows that the covariance, expressed as a matrix, is equal to the inverse of the normal equations matrix:

$$\mathbf{V}^{(k)} = \left[(\mathbf{B}^{(k)})^T \mathbf{B}^{(k)} \right]^{-1}. \tag{3.25}$$

The error matrix may be readily computed from the covariance matrix

$$Q^{(k)}(\omega_i, \omega_j) = V_{i,i}^{(k)} + V_{j,j}^{(k)} - 2V_{i,j}^{(k)}. \tag{3.26}$$

3.5.3 Pseudocode algorithm

For a precise summary, this section presents a pseudocode algorithm.

for $k = 1$ to M **do**

  if $k = 1$ **then**

    Compute η_1 using (3.11)

    Unwrap $\{D_1(\omega_n)\}$ to compute $\{u_{1,n}\}$ for $n = 0, 1, \ldots, N-1$

  else

    Compute $\{\widehat{\Gamma}_{k,n}\}$ for $n = 0, 1, \ldots, N-1$ using (3.12)

    Compute $\{\sigma^2_{\widehat{\Gamma}_{k,n}}\}$ $n = 0, 1, \ldots, N-1$ using (3.14)

    Compute $\rho^2_{k,n}$ using (3.16)

  end if

  Form $\mathbf{B}^{(k)}$ and $\bar{F}^{(k)}$ using (3.23)

  Compute $\mathbf{V}^{(k)}$ using (3.25)

  Compute $\bar{\phi}^{(k)} = \mathbf{V}^{(k)} \bar{F}^{(k)}$

end for

3.6 Implications for the relative phase ambiguity

The registration of the absolute phase using (3.18) for a given shear k requires the previous reconstruction $\phi^{(k-1)}$ to precisely determine the phase difference between at least one pair of frequencies ω_n and ω_{n+C_k} — that is, to imply a value of $\widehat{\Gamma}_{k,n}$ with which the measured $\Gamma_{k,n}$ can be directly compared so as to determine η_k. Mathematically, such a region will then have a low value of $\rho_{k,n}$ via (3.16), so that it will contribute significantly to the sum (3.18). If no such region exists, then the argument in (3.18) will give an essentially random value, the η_k will not be determined, and further reconstruction will be inaccurate.

This requirement for a precise registration of each shear carries implications for the use of multi-shear in determining the relative phase between disjoint spectral regions. The situation is most simply explained by considering the spectral intensity to be binary valued, so that the amplitude of the interferometric product $|D_k(\omega)| = |E(\omega + \Omega_k)E^*(\omega)|$ is either 1 (measurable signal) or 0 (signal swamped by noise). Consider a two-shear reconstruction of a spectrum with two separate spectral regions. To have any chance of measuring the relative phase of the spectral regions, the second shear must be larger than the separation. However, for the registration of the second shear to succeed, one of the spectral regions must also be wider than the separation. The two cases are depicted in Fig. 3.3. In Fig. 3.3(a), one of the lobes is broad enough to allow the larger shear to be registered, whereas in Fig. 3.3(b), this is not the case. Note that these conclusions are consistent with Keusters *et al.* in their previous discussion of this ambiguity. In their analysis of SPIDER they considered only the situation of Fig. 3.3(b) and concluded that resolving the ambiguity was impossible. For SPIDER with a well-characterized external reference (denoted X-SPIDER in [81]), they concluded that the ambiguity could be resolved. My analysis concurs; a well-characterized external reference would mean that the absolute phases of the ancillae are known and the registration procedure is unnecessary.

For more complicated situations, involving more than two separate spectral regions, the reasoning above can be applied inductively. For example, Fig. 3.4 shows how a spectrum with three spectral lobes could be characterized. Note that the gap between the second and third lobe, 3 units, is larger than the three lobes (2, 1 and 1 units respectively), so that if a two shear measurement were attempted then the shear of size 3 required to bridge the gap between the second and third lobe could not have its absolute phase registered. However, if three shears are used, of size 1, 2, and 3 units respectively, then the reconstruction succeeds. The second shear can be registered on the first spectral lobe (arrow R_2), and used to bridge the gap between the first and second lobes (arrow GC_2), as depicted in Fig. 3.4(a). The third shear can then be registered using the "compound lobe" formed by the first and second lobes, whose relative phase is now known (arrow R_3). It then determines the relative phase of the second and third lobes (arrow GC_3).
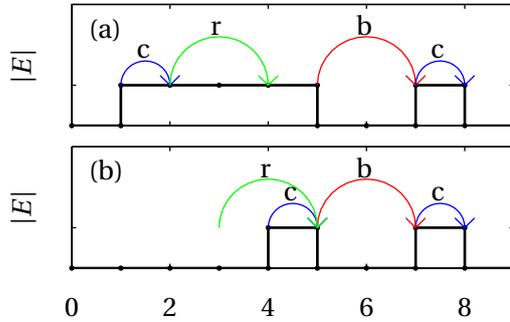
Figure 3.3: Schematic of a two-shear reconstruction procedure for a two-lobed spectrum where (a) one of the lobes is wider than the gap and (b) both of the lobes are narrower than the gap. The spectral amplitude (black, thick lines) is sampled (black dots) at frequencies separated by the small shear $\Omega_1 = 1$. The arrows represent connections between different frequencies given by the shears; only a subset of these are shown. The small shear allows concatenation between adjacent samples (blue arrows, labelled 'c' for concatenation). The large shear $\Omega_2 = 2$ bridges (red arrows, labelled 'b' for bridge) the gap between the lobes. In (a), the absolute phase of the large shear can be registered on one of the lobes (green arrows, labelled 'r' for register) but in (b), no sufficiently wide continuous spectral region is available.
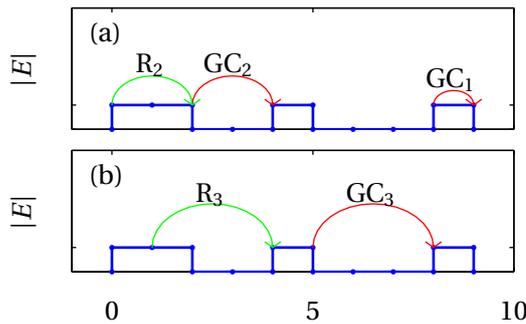
Figure 3.4: Reconstruction procedure for a three lobed spectrum using three shears of size $\Omega_k = 1, 2, 3$. (a) Incorporation of the second shear; the arrows represent the connection of frequencies by the registration R_2 and generalized concatenation GC_2 steps. One of the pairs of frequencies connected by the first shear is also indicated (GC_1). (b) Incorporation of the third shear, using the same notation as in (a).

3.7 Signal-to-noise ratio cost of acquiring multiple shears

Acquiring multiple shears will generally add complexity to the apparatus and will reduce the SNR of the acquired data for a fixed set of measurement resources, such as acquisition time and detector pixels — simply because more information must be acquired. For multi-shear to be beneficial, this must be more than offset by the gain in SNR of the reconstruction. In this section I consider the general properties of various experimental implementations of multi-shear and discuss the SNR reduction that they entail. The multiple shear measurements may be acquired simultaneously or sequentially and I shall consider each case in turn.

For simultaneous acquisition of the multiple shears, the loss of SNR in the raw data comes from two factors: the reduction in signal power caused by additional beam splitting to create the additional shears and the reduction in detector SNR caused by the multiplexing of the additional signals. I consider each in turn. One must distinguish between the detector being near saturation, with additional signal input power available, or being below saturation at full input power so that sensitivity is limiting the precision of the measurement.

Simultaneous acquisition is possible by adding additional beam splitting optics to generate more replicas. In conventional SPIDER, the unknown arm is split L ways and each replica mixed with a single chirped ancilla, whereas in ZAP- and SEA-SPIDER, the chirped ancillas are split. Either way, the energy of the signal is proportional to $1/L$, so that, for example, adding a third arm to provide three shears gives a 33% reduction in signal. If detector sensitivity, rather than saturation, is limiting the measurement, this translates into a corresponding reduction in SNR. In general, each pair of replicas can interfere, yielding $M = L(L - 1)/2$ shears assuming the upconversion frequencies are chosen to give no degeneracies.

Assuming that the detector is completely filled, the multiple interferograms must then be multiplexed. One potential method is to use a carrier-division multiplexing scheme, in which each interference term has a different carrier such as time delay or, for spatially encoded arrangements, beam convergence angle. The various terms, each corresponding to different shears, are disjoint in the Fourier domain and can be isolated by appropriate filtering. The detector resolution must

of course be sufficient to resolve the different sidebands. The coherent summation of L different replicas gives an L^2 enhancement of the peak signal, so that if detector saturation is limiting the measurement then the available SNR is reduced by L^2. If the detector is not completely filled, this can be mitigated somewhat by increasing the magnification then averaging over adjacent pixels. An alternative is to distribute the various shears over different parts of the detector, as is achieved by the inherently multi-shear chirped-arrangement for SPIDER (CAR-SPIDER) [200, 282].

For sequential acquisition, one must consider the reduction in integration/averaging time for each shear to preserve the total acquisition time. I assume that the time spent adjusting the apparatus is negligible. Then, the acquisition time for each of the M shears is inversely proportional to the number of shears. If the detector is near saturation, then the SNR is proportional to the square root of the acquisition time since the various shots must be averaged and the interferogram SNR is proportional to $M^{-1/2}$. Otherwise, the SNR is proportional to the acquisition time since the detector exposure can be adjusted i.e. proportional to M^{-1}.

In summary, for simultaneous acquisition, if the detector is not approaching saturation then the generation and detection of M interferograms using L replicas reduces the SNR by L, whereas if sufficient laser power is available and the detector is saturated then the SNR loss is a factor of L^2. For sequential acquisition, the two factors are M and $\sqrt{M}$, respectively.

3.8   Signal-to-noise ratio benefit of multi-shear

As described above, if the spectrum of the unknown pulse contains a null region between two spectral lobes then a multiple-shear measurement is absolutely necessary. Otherwise, taking additional shears is not a mathematical necessity in the sense that with a sufficient SNR, obtained via averaging or otherwise, a single-shear reconstruction will be accurate. In this section, I derive approximate formulae for the improvement in the precision of the reconstructed phase achieved by multi-shear. To do so, I shall derive explicit formulae for the reconstructed phase in the Fourier domain using the simplifying assumption that the spectral intensity and hence the SNR is uniform across the entire bandwidth. One can then ignore the weighting factors, permitting the use of a simple Fourier method.

The retrieval is formulated as the least squares minimization of the deviation between the unknown phase and the observed phase differences, summing over all shears:

$$R = \sum_{k=1}^{M} \sum_{n=0}^{N-1} |\phi_{n+C_k} - \phi_n - \Gamma_{k,n}|^2 \Omega. \tag{3.27}$$

Using Parseval's theorem and the discrete Fourier transform one has

$$R = \sum_{k=1}^{M} \sum_{m=0}^{N-1} |\tilde{\phi}_m (e^{i\Omega_k t_m} - 1) - \tilde{\Gamma}_{k,m}|^2 2\pi/B. \tag{3.28}$$

Here, the quasi-time points $t_m = m2\pi/B$, $m = -\lfloor \frac{N-1}{2} \rfloor, \ldots, \lfloor \frac{N}{2} \rfloor$ where $B = N\Omega$ is the measured bandwidth, $\{\tilde{\phi}_m\}$ and $\{\tilde{\Gamma}_{k,m}\}$ are the discrete Fourier transforms of $\{\phi_n\}$ and $\{\Gamma_{k,n}\}$ respectively, and M is the number of shears. Minimization of R with respect to $\tilde{\phi}_m$ then gives

$$\tilde{\phi}_m = \frac{\sum_{k=1}^{M} \tilde{\Gamma}_{k,m} \left(e^{-it_m\Omega_k} - 1 \right)}{2\sum_{k=1}^{M} 1 - \cos\Omega_k t_m}. \tag{3.29}$$

Equation (3.29) gives the discrete Fourier transform of the unknown phase in terms of the discrete Fourier transform of the observed phase differences. For a single shear, it reduces to

$$\tilde{\phi}_m = \frac{\tilde{\Gamma}_m}{(e^{it_m\Omega} - 1)} \tag{3.30}$$

which upon inverse Fourier transforming becomes concatenation in the frequency domain, the traditional single-shear algorithm. For multiple shears, (3.29) can be thought of as a weighted sum of the information from the various shears, with the contribution of each reflecting its SNR at each point in the quasi-time domain. The zero in the denominator at $t_m = 0$ corresponds to the absolute spectral phase, which cannot be retrieved with spectral shearing or any other self-referencing envelope-based technique.

The behaviour of the coefficient of $\tilde{\Gamma}_{k,m}$ in (3.29),

$$\beta_{k,m} = \frac{\left(e^{-it_m\Omega_k} - 1 \right)}{2\sum_{k=1}^{M} 1 - \cos\Omega_k t_m}, \tag{3.31}$$

3. PHASE RECONSTRUCTION ALGORITHM FOR MULTIPLE SPECTRAL SHEARING INTEROMETRY
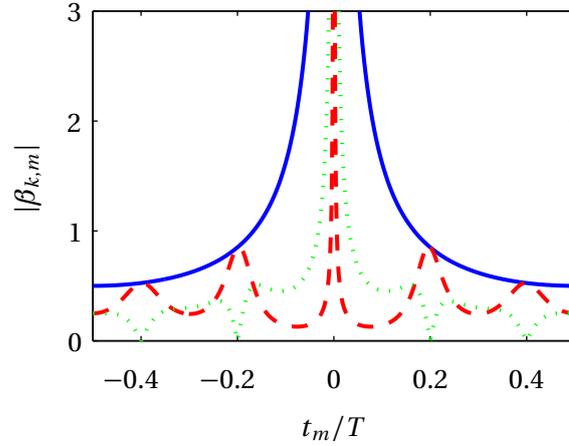
Figure 3.5: (Color online) Modulus of the reconstruction coefficient $|\beta_{k,m}|$ (eq. (3.31) in the text) for a single shear of $\Omega = 2\pi/T$ (solid blue) and for multi-shear retrieval with $\Omega_1 = \Omega$ (dashed, red) and $\Omega_2 = 5\Omega$ (dotted, green).

gives further insight into the algorithm. The modulus of this factor is shown in Fig. 3.5 for the two-shear case $\Omega_1 = \Omega$ and $\Omega_2 = 5\Omega$, and compared with the single shear case. The larger shear (dotted green) information is preferred at around $t = 0$ which corresponds to large scale spectral features. However, at $t = \pm T/5, \pm 2T/5$, the contribution from the larger shear drops to zero because it contains no information with this periodicity, and all information is obtained from the smaller shear, so that its coefficient equals the single-shear case at these points.

To evaluate the overall performance, one considers $\tilde{\Gamma}_{k,m}$ to be corrupted by uniform noise with unit variance. In the retrieval, the noise, along with the signal, is multiplied by (3.31), so that the total noise energy in $\{\tilde{\phi}_m\}$, and equivalently in $\{\phi_n\}$, can be calculated by taking the Pythagorean sum of $\beta_{k,m}$ over all shears and points in the quasi-time domain

$$A^2[\{\Omega_k\}] = \frac{1}{4T} \sum_{m=1}^{N-1} \frac{\sum_k \left| e^{-it_m\Omega_k} - 1 \right|^2}{\left[\sum_k (1 - \cos t_m\Omega_k) \right]^2} \tag{3.32}$$

One can compare the multiple shear case with simply taking the same number of measurements with the same shear to obtain a "noise suppression factor"

$$S = \frac{A[\Omega]}{A[\{\Omega_k = \Omega\}]} \frac{\sqrt{M}}{\alpha_{\text{SNR}}}. \tag{3.33}$$

Here, $A[\{\Omega_k = \Omega\}]$ is evaluated with all shears equal to Ω, in which case (3.32) scales with the inverse square root of the number of measurements, as expected from the averaging of statistically independent samples. The $\sqrt{M}$ in the numerator compensates for this. In the denominator, α_{SNR} is the SNR loss caused by the particular means used to acquire multiple interferograms, as described in Section 3.7. Although I have not explicitly done so, (3.33) can be rewritten to depend only on the normalized shears $\{C_k\}$ and the number of sampling points N. Numerical evaluation of (3.33) is then useful to investigate which shear combinations are likely to yield superior noise rejection for a particular problem.

Figure 3.6(a) shows S for a two-shear reconstruction, taking $\alpha_{\text{SNR}} = 2$, the worst case for sequential acquisition, Fig. 3.6(b) shows the same for a three-shear reconstruction, with $\Omega_3 = \Omega_1 + \Omega_2$ as is the case when three upconverted replicas are generated and interfere in a pairwise fashion. The SNR loss was $\alpha_{\text{SNR}} = (3/2)^2$, the worst case for simultaneous acquisition. It can be seen that $S > 1$ universally, indicating that even when the spectral intensity is uniform, a multiple shear reconstruction gives superior precision. Although I can give no formal proof, one intuitively expects that when the spectrum becomes more modulated, the benefit of multi-shear will gradually increase, up until the point where spectral nulls are present and multi-shear becomes essential for an ambiguity-free reconstruction. Note also that S increases with the number of samples and with the ratio $C_2 = \Omega_2/\Omega$, although the improvement saturates if only one of these parameters is increased. Since, for large C_2, the third shear differs only slightly from the second, a triple shear reconstruction carries only a marginal benefit over two shears and actually performs worse once the SNR cost is taken into consideration as Fig. 3.6(c) shows.

3.9 Numerical comparison of single- and multi-shear retrieval

I now use numerical Monte Carlo simulations to compare the precision of single- and multi-shear retrieval of two complex pulses.

The first pulse is that used in Figure 3.7. Originally, it was a Gaussian pulse with transform-limited duration (intensity FWHM) 30 fs and centre frequency 798 nm. Quadratic dispersion of $2000\,\text{fs}^2$ was applied, and a 20 mrad/fs (7 nm) wide notch cut out of center of the spectrum. A
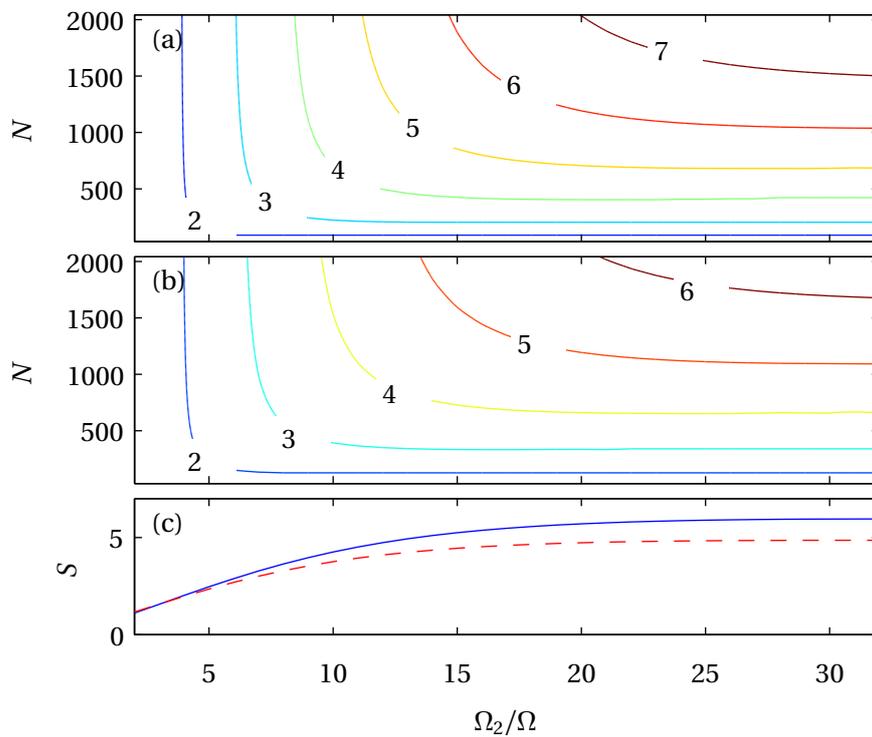
Figure 3.6: (Color online) Contour plots of noise suppression factor S vs number of sampling points and shear ratio for (a) two shears and (b) three shears, with the third shear chosen according to the text. (c) Noise suppression factor vs shear ratio for $N = 1024$ for two shears (blue) and three shears (red dashed).
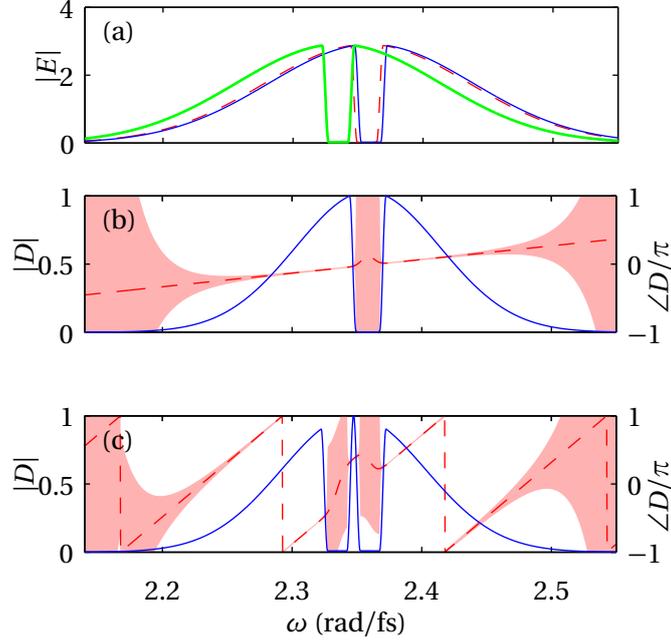
Figure 3.7: (Color online) (a) Spectrum of a two-lobed pulse (blue, solid) and spectrally sheared replicas with small (red, dashed) and large (green, thick line) shears. (b) and (c): Amplitude (blue) and phase (red, dashed) of interference product for the small (b) and the large (c) shear. The variation of the interference products with a peak SNR of 50 is shown by the shaded area.

phase retardation of $\pi/2$ was then applied to the higher frequency lobe. I compared a single shear reconstruction with $\Omega \approx 3.14$ mrad/fs with a two-shear reconstruction with $\Omega_1 = \Omega$ and $\Omega_2 = 8\Omega \approx 25.1$ mrad/fs. The pulse spectral amplitude and phase are shown in Fig. 3.8(a).

Each interferogram $D(\omega)$ was contaminated with additive white Gaussian noise with RMS amplitude equal to 2% of the peak signal amplitude for the single shear case and 4% in the double shear case to emulate the worst-case scenario for multi-shear with sequential acquisition. A random absolute phase, uniformly distributed over $[0, 2\pi)$, was also added to each interferogram. The algorithm of section 3.5 was used for the double shear whilst standard concatenation was used for the single shear case. Having determined the phase, I reconstructed the complex envelope using the spectral intensity of the unknown pulse. The entire process was repeated 1000 times to produce an ensemble of reconstructions so that the precision could be examined. Once this ensemble was computed, it was necessary to remove ambiguous zeroth and first order phase from each reconstruction so that only fluctuations of higher order would affect the statistics. To do this,

87
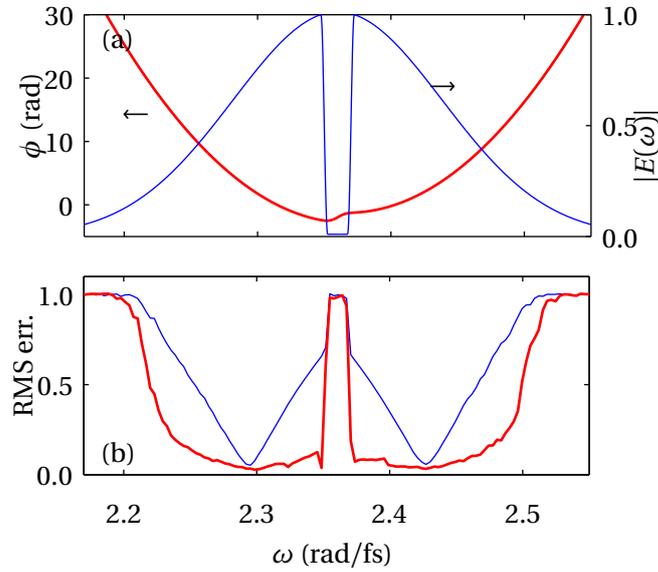
Figure 3.8: (Color online) (a) Spectral phase (red, thick line, left axis) and amplitude (blue, thin line, right axis) of the notched pulse. (b) RMS phase variation for single shear (blue, thin line) and double-shear (red, thick line) reconstructions.

I minimized the RMS field variation over the ensemble [283, 284] using the procedure described in section A.4. To verify the accuracy (difference between the reconstructed and original pulse) I then adjusted the zeroth and first order phase of the entire ensemble to minimize the RMS field error between the average reconstructed pulse and the unknown so that, once again, only errors of higher order would be observed.

To compute the precision of the reconstructed phase, I adopted the following definition of the RMS phase variation which takes into account its modulo-2π nature:

$$\sigma_\phi^2(\omega) = \left\langle \left| \exp[i\phi(\omega)] - \langle \exp[i\phi(\omega)]\rangle \right|^2 \right\rangle. \tag{3.34}$$

Again, angle brackets denote the expectation value over a large number of measurements. Equation (3.34) reduces to the standard RMS when the fluctuations are small, and saturates to $\sigma_\phi = 1$ when the fluctuations are so large that the phase is essentially random.

The results are that both the single- and double-shear reconstructions are accurate, in that their averages do both converge to the unknown pulse. However, the precision of the single-shear
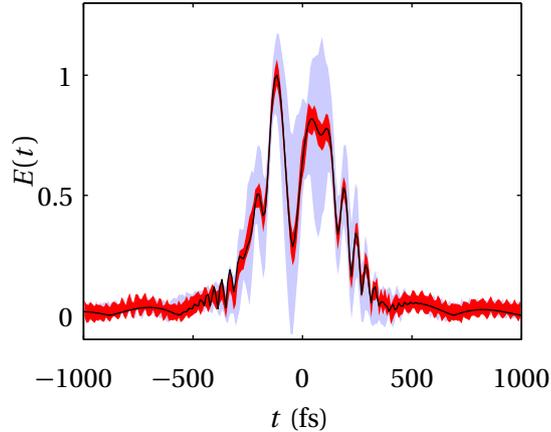
Figure 3.9: (Color online) Temporal amplitude of the notched pulse (black), and the single-shear (blue, light shade) and double-shear reconstructions (red, dark shade). The shaded areas represent the mean plus/minus one standard deviation over the whole ensemble.

reconstruction is very poor, with an RMS field variation of 0.33. This is reduced to 0.02 in the double-shear reconstruction. The typical RMS error between the reconstructed and unknown pulses is 0.26 and 0.006 for the two cases respectively. The precision of the phase of the single-shear reconstruction, shown in Fig. 3.8(b), is also very poor — it varies across the pulse but is greater than 0.5 for a significant fraction of the spectral energy. The double-shear reconstruction achieves a significant improvement, with a typical RMS phase variation of 0.03 rad in regions of significant spectral intensity. Because the relative phase of the two spectral lobes strongly affects the temporal profile, the double-shear measurement produces a dramatic improvement in the time domain as shown in Fig. 3.9. The amplitude fluctuations in the single-shear case are nearly 100%, whereas the double-shear gives a very precise reconstruction.

The preceding example demonstrated that multi-shear is successful bridging a large spectral gap where the signal falls well below the noise level. I now consider a situation where multi-shear is not strictly necessary but, as predicted in the previous section, will improve the precision. The unknown pulse was originally a Gaussian with transform-limited duration (intensity FWHM) 30 fs. Quadratic dispersion of $1000\,\text{fs}^2$ was applied, along with sinusoidal ripples in the spectral phase of period 12.6 mrad/fs and amplitude 0.5 rad. In the time domain, this produces a series of satellite pulses separated by 500 fs, the inverse of the period of the ripples. The pulse therefore has only a
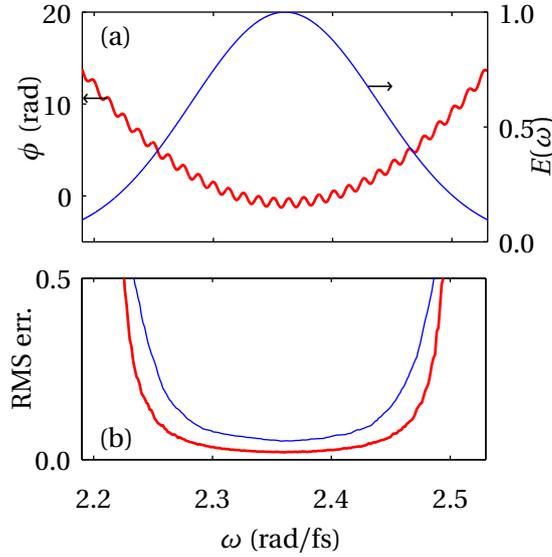
Figure 3.10: (Color online) (a) Spectral phase (red, thick line, left axis) and amplitude (blue, thin line, right axis) of the pulse with a spectral phase ripple. (b) RMS phase variation for single shear (blue, thin line) and double-shear (red, thick line) reconstructions.

smooth variation in its spectral intensity, but both fine and coarse structure in the spectral phase. The two shears used were $\Omega_1 \approx 0.698\,\mathrm{mrad/fs}$ and $\Omega_2 = 15\Omega_1 \approx 10.8\,\mathrm{mrad/fs}$.

The phase and amplitude of the pulse are shown in Fig. 3.10(a), and the RMS phase variation of the reconstructions shown in Fig. 3.10(b). The improvement in precision, whilst not as significant as in the case of the notched spectrum, is still approximately a factor of three across the entire spectrum, in rough agreement with the value of 2.8 predicted by (3.33) above. The agreement is expected since the spectral intensity does not contain any nulls apart from the endpoints. The RMS field precision is 0.052 and 0.028 for the two cases respectively.

3.10 Summary, critical evaluation, and outlook

In this chapter I argued that the use of multiple shears is beneficial to SSI because it improves the precision and resolves the relative phase ambiguity between separate spectral regions. However, it does introduce complexities relating to the absolute phase and unwrapping of the measured phase differences. I presented a reconstruction algorithm which handles these difficulties. It performs the absolute phase determination and unwrapping of a particular shear using a partial re-

construction done using smaller shears. I discussed the cost, in terms of measurement precision, of acquiring interferograms at several shears and then showed that even for a spectrum with no gaps, a multi-shear reconstruction yields superior precision. Finally I presented two numerical examples, demonstrating the robustness of the preprocessing step and the performance improvement achieved by adding a second shear.

There are nonetheless several limitations to the work presented in this chapter. The multiple-shear algorithm is considerably more complex than single-shear concatenation. Furthermore, it requires the shears to be exact integer multiples of one another. In a practical situation where the shears are acquired sequentially by adjusting the apparatus, this requires the parameter (such as a delay stage position) which controls the shear to be calibrated in advance. Finally, the relative phase ambiguity, a motivation for developing the algorithm, is only resolved for spectral separations which are smaller than the largest continuous spectral region. For larger spectral nulls, the ambiguity remains, in common with all known self-referenced method based on the second-order nonlinearity [81].

There are several immediate applications of the multi-shear algorithm. Many implementations of SSI, such as spatially encoded arrangement SPIDER (SEA-SPIDER) [196] or methods which prepare the ancillae using tunable spectral filters [288], allow the shear to be easily adjusted. Sequential acquisition of multiple shears could potentially be automated in implementations based on electro-optic modulation [289], in which the shear is proportional to the applied voltage. Finally, there are several implementations which acquire a range of shears simultaneously on a two-dimensional detector [200, 282]. In these setups, the use of multiple shears will allow more complicated pulses to be measured without losing precision.

A possible extension to this algorithm is the ability to handle a set of shears which are not in integer ratios. As discussed in section 3.4, expressing the phase differences (3.7) in terms of the sampled phase points $\{\phi_n\}$ requires interpolation when the shears are not integer multiples of the sampling rate. Adhering to the assumption of finite support, the "correct" basis for interpolation

3. PHASE RECONSTRUCTION ALGORITHM FOR MULTIPLE SPECTRAL SHEARING INTERFEROMETRY

is in the complex amplitude itself — the sampling theorem gives

$$E(\omega) = \sum_n E(\omega_n) \operatorname{sinc}(\frac{\omega - \omega_n}{\Omega}).$$

(3.35)

However this is nonlinear in the sampled phase, preventing the use of matrix methods. One option is to experiment with linear interpolants, which may perform acceptably under certain circumstances.

A related concept has been studied and implemented in lateral shearing interferometry in two spatial dimensions. Here, the unknown phase $\phi(x,y)$ is obtained using spatial shear measurements of the form $\phi(x+X, y+Y) - \phi(x,y)$. At least two linearly independent shears are required for complete wavefront characterization (up to the usual absolute phase ambiguity) [290]. Additional shears can be used to improve the accuracy [291]. There are connections between the one-dimensional multiple shear case considered here and two-dimensional reconstructions: both involve reconstructing the phase from an overdetermined set of wrapped phase differences. In the most common two-dimensional case, there are two orthogonal shears which define the sampling rate and the wavefront is assumed to not contain any *vortices* — singular points around which the closed path integrals of the phase gradient are nonzero. In this situation, the unwrapping problem which occurred in this chapter does not arise since all phase differences are assumed to be within $[-\pi, \pi]$. However, if larger shears are used, or optical vortices are present, reconstruction algorithms cannot assume the wrapped phase differences are the "true" values. Wavefront reconstruction in the presence of vortices, for example, require sophisticated algorithms [292] and are a topic of current research [293]. It is possible that the multiple-shear algorithm may shed light on such problems, or vice-versa.

In the next chapter, I present the first experimental demonstrations of multiple SSI.

4 Experimental implementations of multiple spectral shearing interferometry

This chapter presents the first experimental demonstrations of multiple spectral shearing interferometry (SSI). Two implementations of spectral phase interferometry for direct electric-field reconstruction (SPIDER) are used. For sequential acquisition of the shears, I used a spatially encoded arrangement SPIDER (SEA-SPIDER), the shear of which can be conveniently adjusted (without altering any other parameters) using a delay stage. For simultaneous acquisition of multiple shears, I used a spatially encoded arrangement CAR-SPIDER (SEA-CAR-SPIDER), a recently developed version of SPIDER that acquires a continuous range of shears mapped to the spatial dimension of an imaging spectrometer. In each case, I synthesized various test pulses and evaluated the precision of the reconstruction by examining the statistics over an ensemble of different acquisitions. The accuracy was tested either by comparison with a single shear reconstruction or by interfering the test pulse with a simple, well-characterized reference.

Tobias Witting provided assistance with the construction the apparatus used in this chapter, as well as with the data acquisition for section 4.1.

4.1 Sequential acquisition of shears

As a proof of principle, I compared the measurement of the spectral phase of a double pulse using standard concatenation and the multi-shear algorithm with two shears. Examination of the statistics of the reconstructions obtained from an ensemble of acquisitions reveals a significant improvement in the precision when using multiple shears. The work presented in this section formed part of the article "High precision self-referenced phase retrieval of complex pulses with multiple-shearing spectral interferometry" [267].

4.1.1 SEA-SPIDER

I used a SEA-SPIDER [196], modified to use an external ancillary pulse rather than splitting off a portion of the test pulse. This variation, when applied to a traditional temporally encoded SPIDER, has previously been described as an X-SPIDER [81], modifed-SPIDER [285] or crosscorrelation SPI-

4. EXPERIMENTAL IMPLEMENTATIONS OF MULTIPLE SPECTRAL SHEARING INTERFEROMETRY
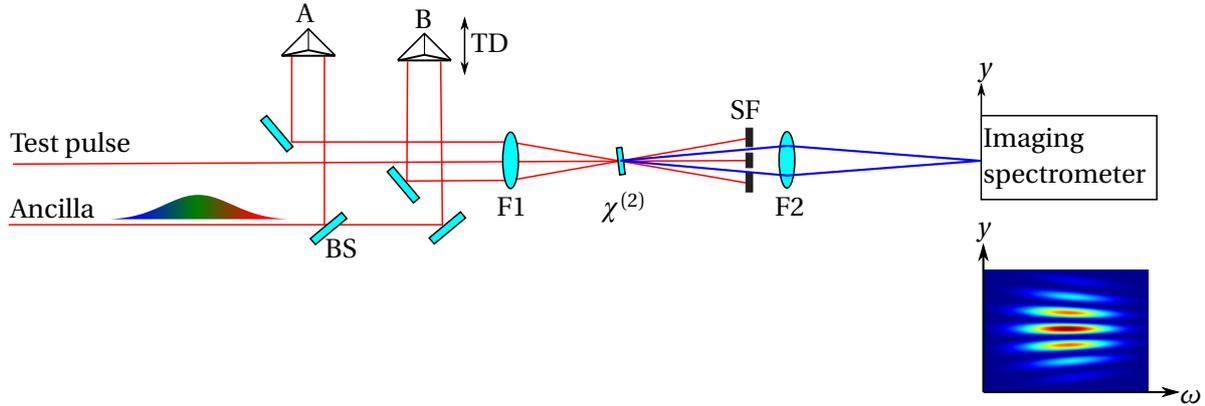
Figure 4.1: SEA-SPIDER setup; showing ancilla beamsplitter BS, time-delay stage TD used to adjust the shear, optic F1 focusing the beams into the crystal, $\chi^{(2)}$ crystal 30 μm BBO cut for Type II sum-frequency generation, spatial filter SF and focusing optic F2 for reimaging the upconverted beams onto the imaging spectrometer.

DER [286]. To obtain the large chirp necessary to ensure sufficiently monochromatic ancillae over the duration of the unknown pulse, a beamsplitter picked off 370 mW of the uncompressed output of our 2 kHz chirped pulse amplifier, which has a quadratic phase of $\phi_2^{\mathrm{anc}} = 1.82 \times 10^6 \, \mathrm{fs}^2$. The average power of the ancilla was reduced to prevent damage to the crystal by passing it through a spectral bandpass filter (centre wavelength 798 nm, FWHM of transmission band 3 nm). This does not reduce the power in the SEA-SPIDER signal because the test pulse only mixes with frequencies that lie within the passband of the filter.

The principle of the SEA-SPIDER was described in section 2.3.6.2. Figure 4.1 shows the details of the setup used here. The ancilla beam is split into two parts, denoted as the A and B arms in the figure. The B arm contains a retroreflector on a motorized delay stage TD. The test pulse and two ancillae are focused into a 30 μm BBO crystal cut for type II sum-frequency generation between the ancillae and the test pulse. The relative time delay τ of the B and A ancillae determines the spectral shear $\Omega = \tau / \phi_2^{\mathrm{anc}}$ between the two upconverted test pulse replicas. A spatial filter blocks all but the upconverted beams, which are reimaged onto the entrance slit of an astigmatism-compensated Czerny-Turner imaging spectrometer [294]. The spectrometer resolved the second order diffraction of the grating with a spectral resolution of 0.18 nm.

The spectral shear was calibrated by taking a series of spectra of the upconverted pulses using
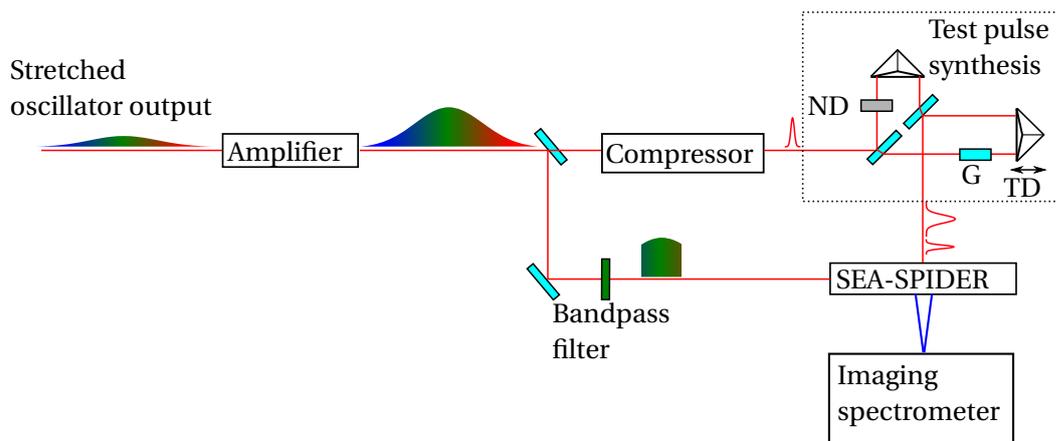
Figure 4.2: Experiment for testing multiple-shear the SEA-SPIDER. ND denotes neutral density filter, G denotes dispersive glass block, and TD denotes time-delay stage.

the B ancilla alone and then numerically crosscorrelating them with the spectra resulting from the A ancilla.

4.1.2 Experiment

The output of the chirped-pulse amplifier which remained after splitting off the ancilla beam was sent through the compressor to produce transform-limited pulses and then into a Mach-Zehnder interferometer which synthesized the test pulse, consisting of two subpulses. The time delay was 2 ps, and we placed a neutral density filter with 22% transmission in the shorter arm and a piece of LASF31A glass with a quadratic dispersion of $1.714 \times 10^3 \, \text{fs}^2$ in the longer arm. The pulse energy was 13.0 µJ, with 2.3 µJ in the leading weak pulse and 10.7 µJ in the trailing dispersed pulse.

To measure the unknown pulse, I used two shears of $\Omega_1 = 0.3672$ mrad/fs and $\Omega_2 = 6\Omega_1 = 2.2034$ mrad/fs, and a reference image with zero shear. To study the statistics of the entire acquisition and reconstruction process, I took an ensemble of 50 samples at each shear and processed each individually.

Figure 4.3(a) shows a typical raw interferogram. The grid-like pattern is a combination of temporal fringes, arising from the double-pulsed nature of the unknown pulse, and spatial fringes, used for the encoding. To isolate the interferometric term, I computed the two-dimensional discrete Fourier transform of each interferogram. A typical result is shown in Fig. 4.3(b). The signal

4. EXPERIMENTAL IMPLEMENTATIONS OF MULTIPLE SPECTRAL SHEARING INTEROMETRY
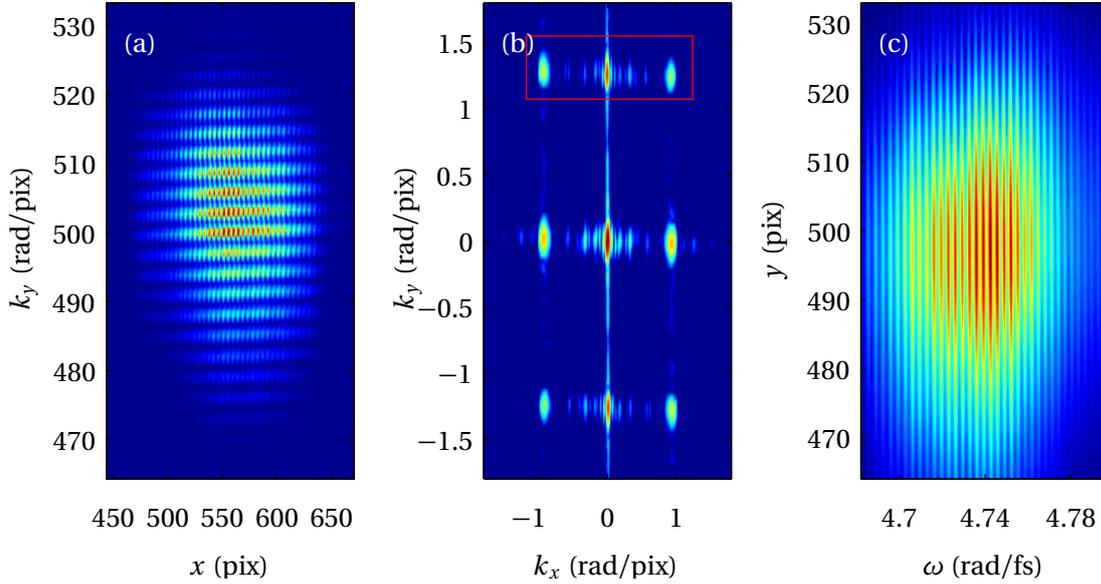
Figure 4.3: (Color online) (a) Raw SEA-SPIDER trace. (b) 2D discrete Fourier transform, showing filter passband (black rectangle). (c) Amplitude of the sideband after filtering and inverse Fourier transform.

is contained in the spatial sidebands, lying above and below the baseband term in the Fourier domain, and one of these is isolated by applying the filter indicated by the overlaid rectangle. The vertical stripe at $k_x = 0\,\mathrm{pix}^{-1}$ is an artefact of the "dead rows" of the detector, which introduce noise at all vertical spatial frequencies. The Fourier filtering substantially reduces the influence of these artefacts. After applying the filter and inverse Fourier transforming, the amplitude of the resulting interferogram is shown in Fig. 4.3(c). The SNR, defined here as the ratio of the maximum value of the signal to the RMS background noise level, was ≈ 50 before the filtering and ≈ 60 afterwards. For the noninterferometric images, obtained by blocking one arm and used for the shear calibration and obtaining the unknown pulse spectrum, I performed a similar filtering procedure to isolate the baseband components and reduce noise.

I performed all subsequent steps on lineouts taken at $x = 0\,\mu\mathrm{m}$, and subtracted the common upconversion frequency of 2.380 rad/fs to produce spectra centered around 800 nm. I obtained the spectral intensity, shown in Fig. 4.4(a), from the spectra taken with the B arm blocked. For the filtered interferometric phase, I subtracted the phase of the zero-shear interferogram to yield the complex valued $\{D_k(\omega)\}$, and then applied the algorithm of sections 3.5. The phase differences

$\{\Gamma_k(\omega)\}$ are shown in Fig. 4.4(b). The resulting retrieved phase, for the double-shear reconstruction, is shown in Fig. 4.4(c). Similarly to the numerical demonstration in section 3.9, I evaluated the RMS phase variation, shown in Fig. 4.4(d). For a fair comparison of the precision of single- and double-shear reconstructions, I averaged together pairs of single-shear reconstructions, thus increasing their SNR by a factor of $\sqrt{2}$ to create an ensemble of 25 before computing the statistics. The double-shear reconstruction is much more precise, particularly in the spectral wings. This manifests as a palpable improvement in the time domain, shown in Fig. 4.5. A global indication of the precision is given by the RMS field variation, eq. (A.13), which was 0.12 and 0.02 in the single- and double-shear cases respectively.
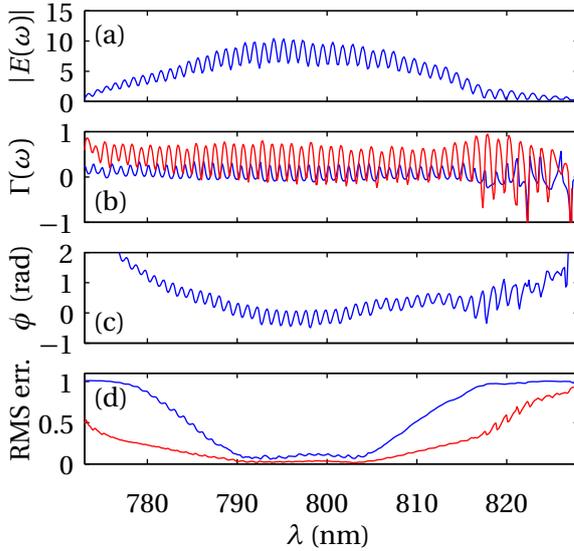
Figure 4.4: (Color online) (a) Spectral intensity of the unknown pulse. (b) Phase differences from small (blue) and large (red) shears, averaged over the ensemble. (c) Retrieved phase, averaged over the ensemble. (d) RMS variation of the retrieved phase over the ensemble for the single- (blue, thick line) and double- (red, thin line) shear reconstructions.
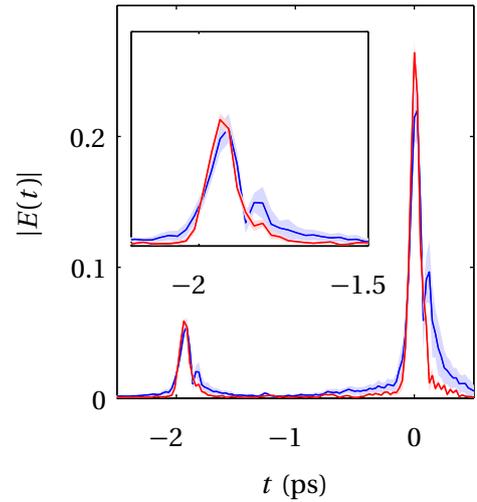
Figure 4.5: (Color online) Temporal amplitude of the single-shear (blue, light shade) and double-shear reconstructions (red, dark shade), averaged over the whole ensemble. The region represents the mean plus/minus one standard deviation. The inset is a zoomed-in view of the shorter subpulse.

I also checked the accuracy of the reconstruction by numerically isolating the two subpulses with a temporal filter and verified that their relative energies were in agreement with our mea-

surements and the difference in their group-delay dispersion equalled the known value for the LASF31A glass. Finally the RMS time-bandwidth product (defined using angular frequency ω) of the reconstructed pulse was 11.9, one of the largest ever measured using a spectral shearing method.

Having experimentally proven the principle of multiple-spectral shearing interferometry, I now describe a means of acquiring multiple shears simultaneously.

4.2   Simultaneous acquisition of shears

Sequential acquisition prohibits rapid updates and requires the shears to be chosen in advance. This motivates the simultaneous acquisition of a range of shears in a two-dimensional data set. One potential way of acquiring multiple shears simultaneously is by adding more arms to the interferometer in a SPIDER. This would incur extra complexity and would require some form of multiplexing for each shear to be detected. Another approach is a chirped-arrangement for SPIDER (CAR-SPIDER), introduced in section 2.3.6.2. Here, I used a spatially encoded arrangement for CAR-SPIDER, a new implementation developed during the course of this dissertation [282].

The work presented in this section formed part of the article "Resolution of the relative phase ambiguity in spectral shearing interferometry of ultrashort pulses" [267].

4.2.1   Spatially encoded arrangement chirped-arrangement SPIDER

The operating principle of the SEA-CAR-SPIDER is depicted in Fig. 4.6. The test pulse (TP) $E(\omega)$, propagating along the z-axis, is sum-frequency mixed with two ancilla beams A & B which are spatially chirped along the x-axis and crossing at an angle θ in the xz-plane. The spatial chirps are oppositely aligned, such that the local frequencies of the two ancillae are $\omega_{\text{A}}(x) = \omega_{\text{up}} + \alpha x$ and $\omega_{\text{B}}(x) = \omega_{\text{up}} - \alpha x$ where ω_{up} is their common frequency at $x = 0$ and α the degree of spatial chirp. Two noncollinear replicas of the test pulse are produced, upconverted by the local ancilla frequency. These are reimaged onto the entrance slit of an imaging spectrometer, producing spa-
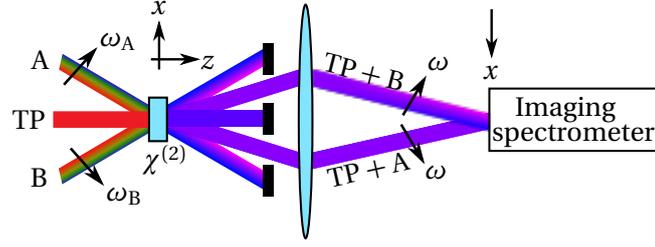
Figure 4.6: SEA-CAR-SPIDER concept: test pulse (TP), spatially chirped ancillae A and B with local frequencies ω_A and ω_B increasing along the arrows, and sum-frequency beams TP+A and TP+B. The coordinate system used in the text is shown. The entrance slit of the imaging spectrometer is parallel to the x-axis; the x-axis is reversed here by the reimaging of the crystal plane.

tial fringes caused by the beam crossing angle. The resulting signal is

$$S(\omega, x) = |E(\omega - \omega_A)|^2 + |E(\omega - \omega_B)|^2 +$$

$$2|E(\omega - \omega_A)||E(\omega - \omega_B)| \cdot \cos\left[\phi(\omega - \omega_B) - \phi(\omega - \omega_A) + \frac{\theta\,\omega_{up}\,x}{c}\right]. \tag{4.1}$$

For each x, the measured data is an SSI trace with shear $\Omega(x) = \omega_A - \omega_B = 2\alpha x$.

The experimental implementation of SEA-CAR-SPIDER is depicted in Fig. 4.7. The test pulse TP is split with beamsplitter BS1 and focussed into the nonlinear crystal $\chi^{(2)}$ with a cylindrical lens CL of focal length $f_{cl} = 500\,\text{mm}$. The delay-line DL1 is only for coarse adjustment of the delay between ancillae and test pulse and only requires initial adjustment. The ancillae are prepared as follows: the reflected beam from BS1 is dispersed by a grating G of pitch $\Lambda = 300\,\text{mm}^{-1}$ and angle of incidence $\gamma = 10°$. The spherical lens L1 is positioned one focal length $f_1 = 500\,\text{mm}$ away from the grating and the crystal is mounted f_1 further downstream. Between the lens and the crystal the ancilla beam is split at 50/50 beamsplitter BS2 and one arm is spatially inverted by applying an odd number of reflections. The coarse delay line DL2 is used to achieve temporal overlap between the ancillae and needs no subsequent adjustment. The $\chi^{(2)}$ crystal (250 μm BBO) is placed in the back Fourier plane of L1, where the spatial chirp of the ancillae is $\alpha = \cos[\beta(\omega_{up})]\omega_{up}^2/(2\pi c f_1 \Lambda)$ where $\beta(\omega_{up})$ is the diffracted angle of the beam from the grating as depicted. A half-wave retarder $\lambda/2$ is introduced into the ancilla arm to ensure type-II upconversion with the ancilla polarization aligned to the e-axis and the TP polarization to the o-axis of the crystal, respectively. After spatial
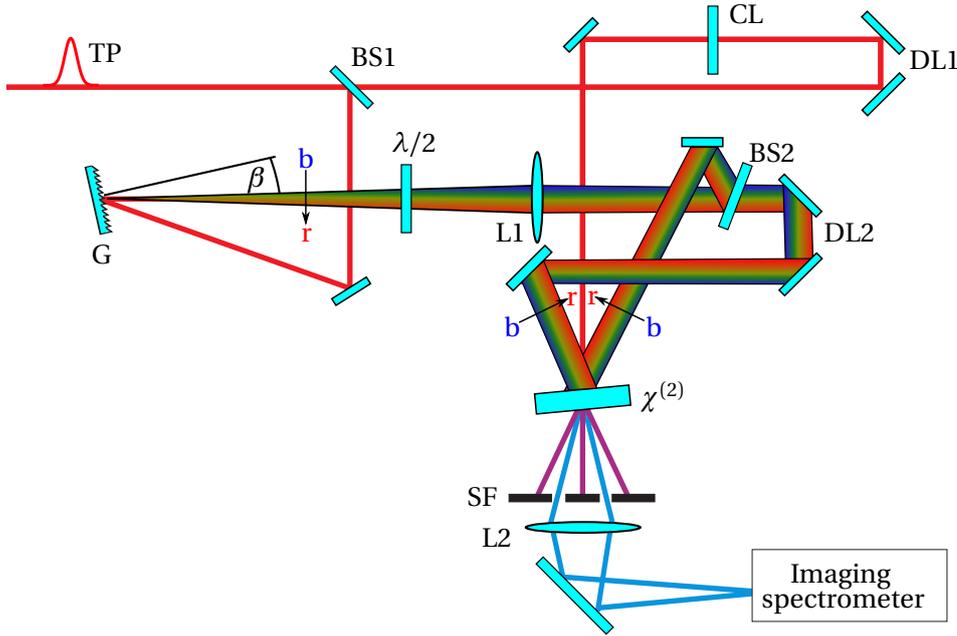
Figure 4.7: SEA-CAR-SPIDER setup. The symbols are defined in the text.

filter SF the sum-frequency beams are re-imaged onto an imaging spectrometer by lens L2 with focal length $f_2 = 300\,\text{mm}$. The spectrometer's entrance slit is oriented in the dispersion plane of the ancillae and the signal is internally dispersed perpendicular to this plane. The signal is recorded with a 1280×1024 pixel CMOS detector with 8 bit analog-to-digital conversion. Single-shot data is recorded at a rate of 30 Hz limited by memory and data-transfer capacities.

4.2.2 Experiment

Figure 4.8 shows the experimental setup. To demonstrate recovery of the spectral phase across a null, I characterized a bichromatic double pulse, synthesized using a $4f$ pulse shaper with a 19.4 mrad/fs wide stop ($\approx 25\%$ of the full-width at half-maximum (FWHM) bandwidth) in the Fourier plane. To one side of the stop, a glass slide delayed frequencies above 2.378 rad/fs by 380 fs. The temporal intensity, spectrum and spectrogram of the resulting pulse is shown in Fig. 4.9. Figure 4.11(c) shows the measured SEA-CAR-SPIDER interferogram, consisting of the coherent sum of the two single-arm spectra shown in Figures 4.11(a) and (b). The large fringe-covered areas at lower and higher frequencies in Fig. 4.11(c) represent self-overlap of the near transform-limited
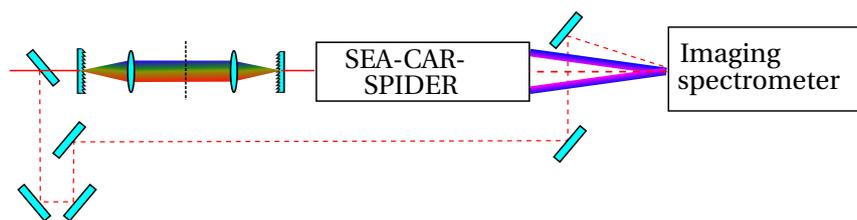
Figure 4.8: Experiment for demonstrating and verifying measurement of shaped pulses with SEA-CAR-SPIDER. The dashed lines represent the beams used for spectral interferometry.
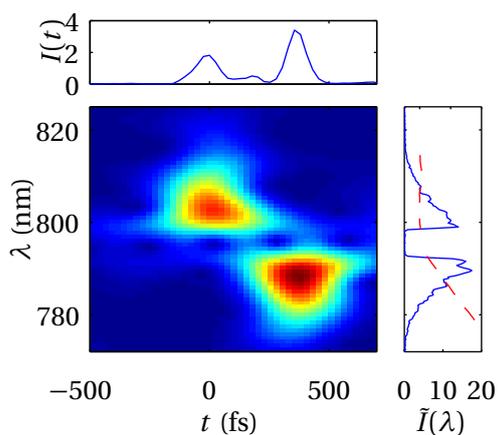
Figure 4.9: Temporal intensity, spectrum and spectrogram of the bichromatic double pulse, as reconstructed using SEA-CAR-SPIDER.
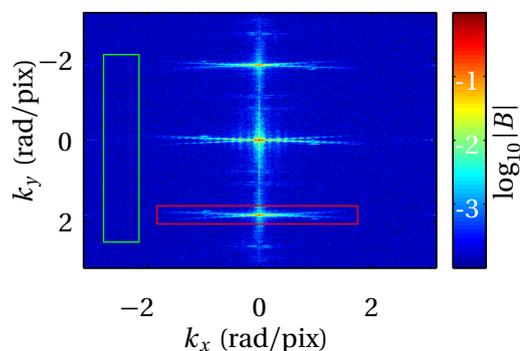
Figure 4.10: Amplitude of Fourier transform of AB interferogram, $\log_{10}$ color scale. The boxes show the filters used for the interferogram (red) and noise estimation (green).

spectral lobes. The fringes are therefore largely flat, with small modulations caused by high-order phase distortions present in our laser system. The triangular fringe-covered area at the bottom of Fig. 4.11(c) represents cross-term interference between the upper and lower lobes. The fringe tilt is due to the group-delay difference between the lobes, whilst their phase offset (with respect to the self-overlap fringes) encodes the relative phase of the lobes.

I extracted the fringe phase using Fourier filtering and applied the multiple-shear algorithm using the 6 shears indicated by the dashed lines in Fig. 4.11(c). I verified the accuracy of the measurement using spectral interferometry (SI), shown by the dashed lines in Fig. 4.8. I first characterized an unshaped pulse, produced with no mask in the pulse shaper, using a single shear with the SEA-CAR-SPIDER. I then obtained the phase of the reference arm by interferometry with the characterized unshaped pulse. The phase of any shaped pulses could then be simply and accurately
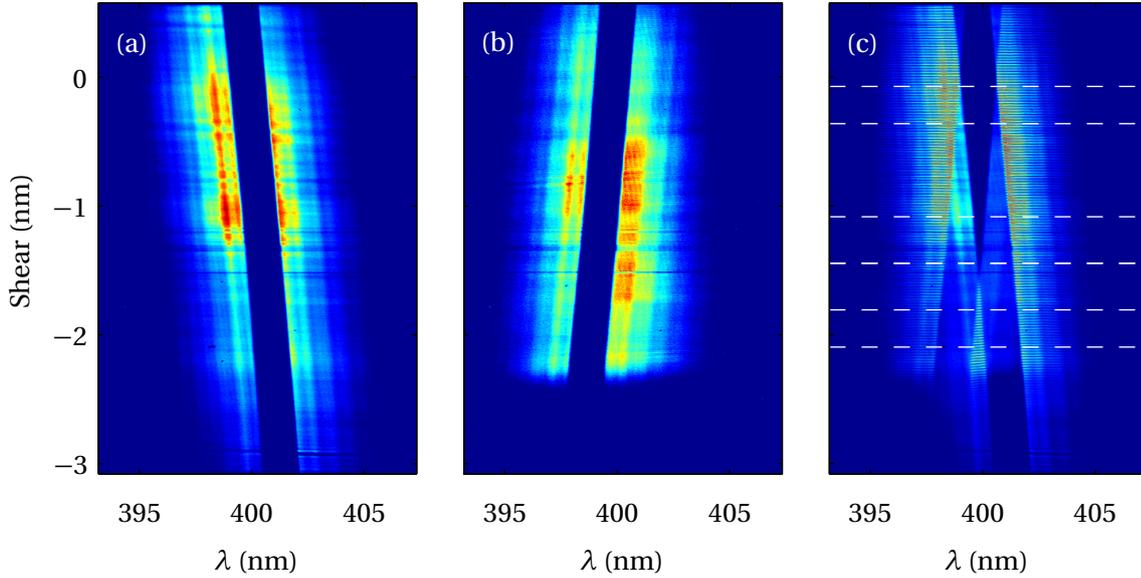
Figure 4.11: Individual spectra of upconverted replicas in A arm (a) and B arm (b); (c) SEA-CAR-SPIDER trace.

compared to the reference. Figure 4.12 shows the phase of the two-lobed pulse reconstructed using multiple shears, along with that obtained using SI. The disagreement (black dotted line) does not exceed 0.25 rad over the full-width at 10% maximum intensity.

4.2.3 Error analysis

In this section I demonstrate the use of the error matrix, introduced in section 3.2, to describe and quantify uncertainty in the spectral phase reconstruction. Figure 4.13 shows several experimentally obtained $\mathbf{E}$ for single- and double-lobed spectra, with blue indicating precise determination of the relative phase. For a single-lobed spectrum with single-shear reconstruction (Fig. 4.13(a)), $\mathbf{E} < 0.2$ rad between all pairs of frequencies down to the 10% intensity level. For a two-lobed spectrum with a single-shear reconstruction (Fig. 4.13(b)), the error between frequencies of the same lobe is less than 0.4 rad but is essentially infinite between the lobes — the off-diagonal terms of $\mathbf{E}$. Using multiple shears reduces $\mathbf{E}$ to less than 0.03 rad (Fig. 4.13(c)) between all frequencies down to the 10% intensity level.
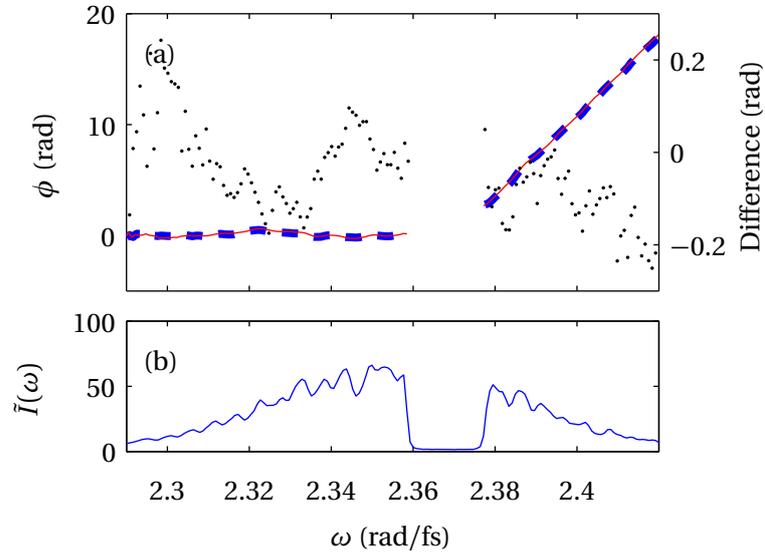
Figure 4.12: (a) Measured spectral phase (left axis) using multiple-shear SEA-CAR-SPIDER (solid red) and SI (dashed blue); disagreement between the two (dotted black, right axis). (b) Spectral intensity.
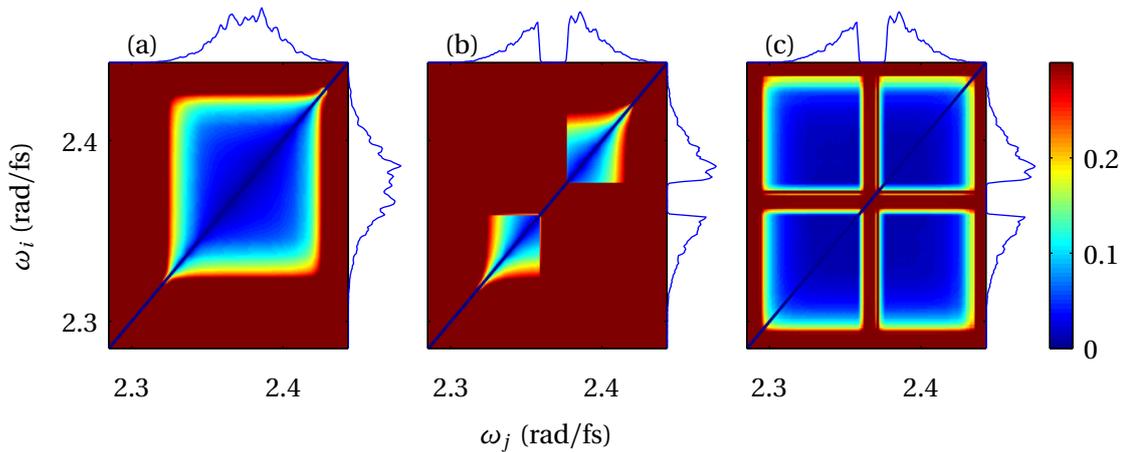
Figure 4.13: Error matrix **E** for (a) single-lobed spectrum, single-shear reconstruction; (b) double-lobed spectrum, single-shear reconstruction; (c) double-lobed spectrum, multiple-shear reconstruction. Spectra are shown for reference.

4.2.4 Discussion

One limitation of the present setup arises because the ancillae are derived from the test pulse. Upconversion therefore cannot be achieved with these missing frequencies. This can be somewhat inconvenient — the central upconversion frequency ω_{up} must be carefully chosen to lie in the centre of a spectral lobe so as to provide the widest possible range of upconversion frequencies either side. Fortunately, ω_{up} can be modified without altering any of the other parameters by rotating the grating used to disperse the ancillae. This issue can also be alleviated by deriving the ancillae from an external source, as was done for the SEA-SPIDER in section 4.1. However for general use a self-contained device is preferable.

4.3 Summary and outlook

This chapter described multiple shearing spectral interferometry using both sequential and simultaneous acquisition of the shears. For the measurement of a complex pulse without spectral nulls, the addition of a second shear significantly improved the precision. For a pulse with a spectral null, the use of multiple shears gave a precise and accurate phase reconstruction which would have been impossible using only a single shear.

These results prove the principle of multiple spectral shearing interferometry. An interesting application would be measurement of the output of hollow fibre compressors for few-cycle pulse generation [60], which can produce highly structured spectra, especially when pushed to their maximum bandwidth. Single-shear measurements using SEA-SPIDER have already been presented [197]. The combination of high spectral resolution and precision offered by multiple shear reconstructions would enable high dynamic range measurements with a longer time window than is currently achieveable with a single shear.

5 Compact Space-time SPIDER

In this chapter I describe a new version of space-time SPIDER (ST-SPIDER), a self-referenced spatio-temporal pulse characterization technique demonstrated by Dorrer *et al.* [58, 59]. The new version is more compact, involving a single interferometer rather than two separate spatial and spectral interferometers as used in the original. Furthermore, it addresses several subtle issues which arise from an in-depth analysis of the previous design.

Section 5.1 discusses some motivation for the use of spectral and spatial shearing interferometry for spatio-temporal pulse characterisation. All such measurements have an implied object plane which must be considered when performing reconstructions and this is discussed in section 5.2. An analysis of the current ST-SPIDER design is given in section 5.3 and several issues are identified which are considered in the new design. The compact ST-SPIDER is presented in section 5.4 and an experimental example using angular dispersion given in section 5.5. The sensitivity of the device to instability is analysed in section 5.6. Section 5.7 offers a summary and conclusions.

5.1 Motivation for ST-SPIDER

Most of the spatio-temporal characterization techniques presented to date consist of a spatially resolved wavefront measurement combined with a one-dimensional temporal characterization of a single point in the beam. The wavefront measurement returns $\phi(\omega, \mathbf{x}_T) + f(\omega)$, where $\mathbf{x}_T$ is the transverse position in the beam and $f(\omega)$ an unknown function of frequency. I write $\mathbf{x}_T$ as a 2-vector for generality; however many techniques resolve only one spatial dimension. The temporal characterization returns the phase $\phi(\omega, \mathbf{x}_{T,0})$ at point $\mathbf{x}_{T,0}$. If all frequencies of the pulse are present at $\mathbf{x}_{T,0}$, then $f(\omega)$ and hence $\phi(\omega, \mathbf{x}_T)$ can be determined.

Techniques of this class have one general limitation. If the field at $\mathbf{x}_{T,0}$ does not contain all the frequencies present in the beam, then the unknown function $f(\omega)$ remains unknown at these frequencies. The situation is depicted in Fig. 5.1(a) for a single spatial dimension y. The sampling points in (ω, y) space are black dots and one of these is greyed out to represent a missing frequency at the point in the beam where the spectral measurement is performed. The relative phase at this frequency of all the other points remains unknown. Whilst this situation may seem
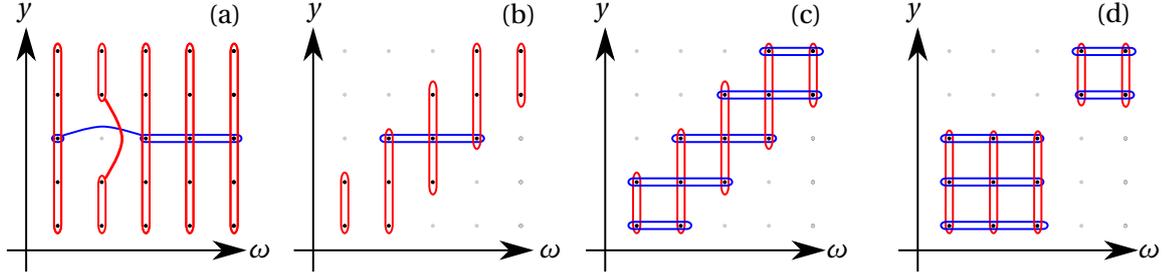
Figure 5.1: Spatio-spectral phase reconstruction with a spatially resolved wavefront measurement (red) combined with a one-dimensional spectral phase measurement (subfigures (a) and (b)) and a spatially resolved spectral phase measurement (subfigures (c) and (d)). The greyed out point represents a frequency which is missing from the pulse. In (a), a single frequency-point is missing. In (b) and (c) the beam is spatially chirped. In (d), the beam consists of two spectrally and spatially disjoint regions.

somewhat artificial and one could simply choose a different point to perform the spectral phase measurement, there are realistic situations where this is not possible. Figure 5.1(b) depicts a pulse with some spatial chirp. Regardless of where the spectral phase is measured, not all frequencies will be present.

One way of avoiding this limitation is to perform a spatially resolved spectral phase measurement, which returns $\phi(\omega, \mathbf{x}_\text{T}) + g(x)$. Combining this with the spatially resolved wavefront measurement, the phase at missing frequencies can be obtained, as depicted in Fig. 5.1(c). Note that this still does not guarantee an ambiguity-free reconstruction — if several spectrally and spatially disjoint regions are present, as depicted in Fig. 5.1(d), then their relative phase will not be returned. However, a much larger class of pulses can be measured. In particular, all pulses with a simply connected spatio-spectral intensity profile can be measured with this approach. Another potential advantage is the introduction of redundant information to the retrieval. In general, this adds noise robustness and flags measurement errors, turning the reconstruction algorithm into an optimization problem of an overdetermined system. This raises the possibility of using algorithms which weight the signal according to its local signal-to-noise ratio (SNR). Such approaches have found success in the related problem of two-dimensional phase unwrapping [295], particularly around signal dropouts which occur, for example, near optical vortices.

Spectral shearing interferometry is well suited to the acquisition of a spatially resolved spec-

tral phase measurement because the data is one-dimensional, leaving another dimension free for spatial resolution if an imaging spectrometer is used. In principle, it could be combined with any of the spatial phase methods described in Section 2.4.4. To date, however, the only demonstration has been that of Dorrer *et al.*[58, 59, 265] who combined spectral phase interferometry for direct electric-field reconstruction (SPIDER) with lateral shearing interferometry (LSI). This chapter develops this combination.

Before presenting the new apparatus, it is necessary to discuss a subtlety which arises when spatial and temporal measurements are performed using separate devices.

5.2   Role of the measurement plane

In spatio-temporal metrology, it is important to consider the *object plane* in which the measurement is being performed. This is because beams undergo diffraction as they propagate in free space. For example, if a spectrally and spatially transform-limited pulse is passed through a thin diffractive element, then just after the element, the pulse possesses angular dispersion. Upon subsequent propagation, the frequencies travel in different directions and arrive in different positions; the pulse has acquired spatial chirp. A space-time characterization occurs in some tranverse plane, and this plane must be specified, or at least kept in mind.

This consideration also applies to purely temporal measurement because of the dispersion of air. However dispersion is simply an multiplicative phase factor $\exp[i\phi(\omega)]$, whereas diffraction is a convolution with the spatial Fourier transform of the paraxial propagator $\exp(-ik_T^2 z/k)$, which is frequency-dependent. In both cases, if the plane of measurement differs from the plane of interest, then numerical propagation can relate the two. However with diffraction the relation is more complex, contravening the practical and philosophical principle that the mapping between experimental data and the physical quantities should be as simple as possible.

When the spatial and spectral phase profiles are being measured separately, differences between the two planes of measurement warrant particular consideration. Consider the two measurements $\phi(\omega, \mathbf{x}_T, z_T) + f(\omega)$ and $\phi(\omega, \mathbf{x}_T, z_S) + g(\mathbf{x}_T)$ at different planes z_T (subscript T for transverse spatial) and z_S. To combine these measurements, they must be first specified at a common
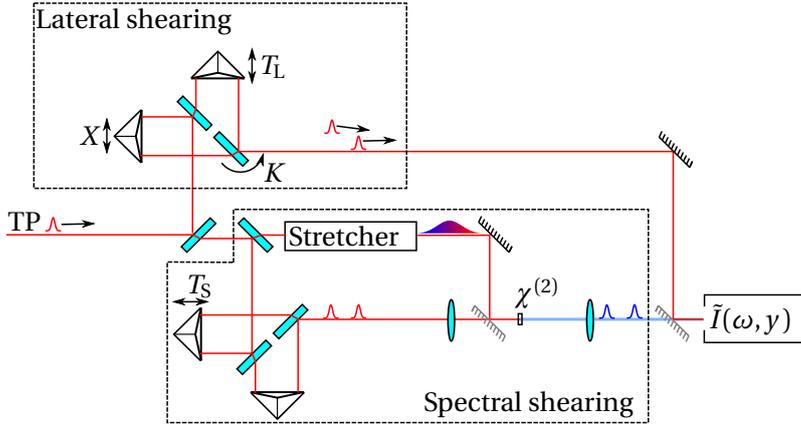
Figure 5.2: ST-SPIDER implemented by Dorrer *et al.* [58, 59, 265]. The test pulse (TP) is split and simultaneously passed through lateral and spectral shearing interferometers. The lateral shearing interferometer applies time delay T_L, tilt K and a spatial shear which depends on X as well as the displacement induced by propagation to the detector if $K \neq 0$. The spectral shearing interferometer is a standard SPIDER design with time delay T_S. Mirrors drawn in gray are at a lower height allowing the beam incident from the rear to pass overhead.

plane. In principle this is possible by propagating the partially reconstructed field at z_T to z_S, since the unknown spectral phase $f(\omega)$ plays no role in linear paraxial propagation. However, doing so complicates the reconstruction and creates an additional calibration parameter. Furthermore, it places at risk one of the proposed advantages of acquiring redundant data, which is the potential to apply reconstruction algorithms which consider the local SNR. Since propagation is spatial convolution, errors which are local in $\mathbf{x}_T$ become global, complicating the implementation of weighted algorithms. Therefore, in this study I stipulate that the two planes of measurement must be equal. The implications of this for ST-SPIDER are discussed below.

5.3 Analysis of ST-SPIDER

The original ST-SPIDER of Dorrer *et al.* [58, 59, 265] is shown in Fig. 5.2. It consists of separate lateral and spectral shearing interferometers, the latter being a standard SPIDER design.

The first motivation for this chapter is reducing the complexity of the device, which involves two Mach-Zehnder interferometers (MZIs). The second motivation is some subtle issues concerning the device's accuracy, which are the topic of this section. I emphasize that none of these issues are catastrophic, nor do they invalidate previous results. They do however increase the complexity

of design and operation of the device and of reconstruction.

5.3.1 Role of the telescope in ST-SPIDER

The double Fourier transform configuration introduces several subtleties. Because frequency conversion occurs in the Fourier plane, there is a frequency-dependent reduction (in the sense of inverse magnification) of the upconverted beams. To see this, recall that in the Fourier-transform property of a lens, the mapping of the spatial coordinate x in one focal plane to the transverse wavenumber k_x in the other is $x/f = k_x/k$. For the purposes of this argument, I take the ancilla as a plane wave. Upconversion therefore involves no modification to the transverse wavenumber. However, the longitudinal wavenumber changes from ω/c to $(\omega + \omega_{\mathrm{up}})/c$. Therefore, in the output Fourier plane of the second lens, the upconverted beam is magnified by $\omega/(\omega + \omega_{\mathrm{up}})$. (There is of course an inversion which is not important in this discussion).

The seriousness of this distortion depends on the bandwidth of the pulse and the shear used. In the narrowband and small shear limit, the magnification varies by $B/(\omega_0 + \omega_{\mathrm{up}})$ across the spectrum, and therefore the effect can be approximately compensated by spatially scaling the acquired image by the inverse of this factor. As the spectrum broadens, the scaling becomes frequency dependent and must be taken into account. Additionally, the difference between the two upconversion frequencies causes the two replicas to have different magnifications. To first order in the shear, the difference between the magnifications is $\omega\Omega/(\omega + \omega_{\mathrm{up}})^2$. For $\omega \approx \omega_0$, and assuming $\omega_{\mathrm{up}} \approx \omega_0$, this is equal to $\Omega/(4\omega_0)$. When the difference in the magnifications of the two replicas becomes significant, the two replicas will acquire a radial shear, severely complicating the reconstruction.

The telescope also affects the object plane of the device. The amplitude and phase of the field at the input focal plane is reproduced at the output focal plane, a distance $4f$ downstream. In terms of diffraction, it is as if the intervening space did not exist. In the light of the comments in section 5.2, this must be considered in laying out the device.

5.3.2 Role of ancilla spatial profile

In the ST-SPIDER design of Fig. 5.2, the ancilla are not focused into the crystal. The reason for this design decision, which reduces signal intensity, is to do with the role played by the spatial profile of the ancilla. I shall now explain this. Sum-frequency generation, which occurs at the focus, is multiplicative in the spatial domain. The spatial profile of the upconverted beam is therefore the product of the spatial profile of the unknown with that of the ancilla. In the collimated beam, which is detected, the two spatial profiles are therefore convolved. Such an operation will, in general, scramble the phase. It is therefore necessary to ensure that the ancilla is close to a plane wave propagating along the z-axis at the two frequencies ω_{up} and $\omega_{up}-\Omega$. This can be achieved by spatially filtering the ancilla [296], or as in Fig. 5.2, making the ancilla much larger than the unknown pulse focus.

5.4 A common-path re-imaging ST-SPIDER

I now present the new design, shown in Fig. 5.3, and explain how it addresses the issues of section 5.3. A common path is used for the spatial and spectral shearing; only a single interferometer is needed, which performs both the spatial shear and the time delay that is necessary for the spectral shear. In the diagram, the spatial characterization is in the vertical plane i.e. out of the page in the top view of Fig. 5.3. However, this only requires minor modification to change. In layout, the design is very similar to the SPIDER in Fig. 5.2. A crucial difference is that the lenses L_1 and L_2 are not in a telescopic configuration; their separation is larger than the sum of their focal lengths. The crystal C is located slightly downstream of the focus of $L_1$the first lens. It therefore has conjugate plane O some distance upstream. Lens L_2 re-images the crystal plane onto the entrance slit of the imaging spectrometer D.

The re-imaging configuration prevents the ancilla spatial profile from scrambling the spatial phase of the upconverted beams, as discussed in section 5.3.2. The spatial profile acts multiplicatively, but is frequency-independent since only one frequency of the ancilla participates in the upconversion. The ancilla can therefore be focused with the beams, although it is useful to enlarge it
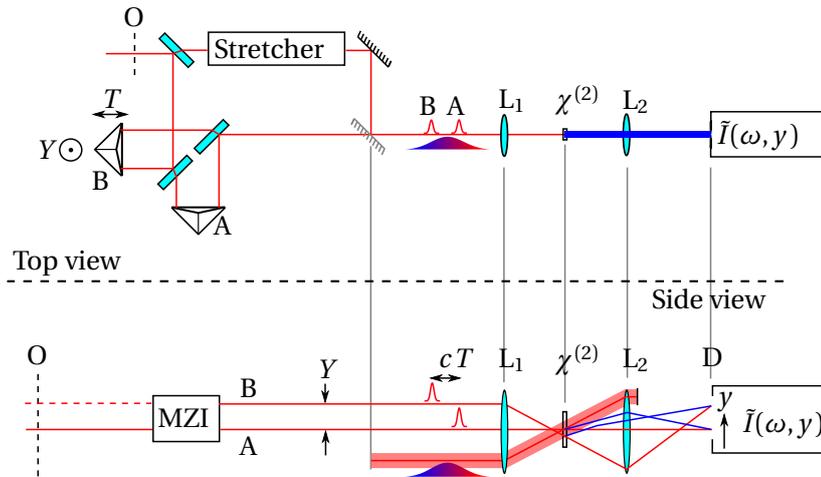
Figure 5.3: Common-path re-imaging ST-SPIDER. Interferometer arm A is stationary whilst the corner-cube in arm B is displaced to apply time delay T and vertical spatial shear Y. In the side view, the MZI has been reduced to a "black box", so that beam A passes straight through whilst beam B receives the vertical displacement and time delay. The mirror drawn in gray is at a lower height allowing the beam incident from the rear to pass overhead.

slightly with a telescope (not shown) to ensure a strong upconverted signal in the wings.

At the detector, the fundamental and upconverted beams arrive at the same positions, since they are re-imaged from their positions in the crystal. Either one may be selected by color filtering; in fact both may be detected simultaneously for single-shot operation. The fundamental beams have a spatial shear, whilst the upconverted beams have both spectral and a spatial shear. The latter is a side-effect of the common-path configuration but, as shown below, is straightforward to accommodate in the reconstruction. Both pairs of beams have the same temporal delay T. A tilt, perhaps applied inadvertedly, is to be avoided; although in principle it does not cause a problem, it complicates the use of a convenient calibration procedure, explained in section 5.4.2 below.

The re-imaging system after the interferometer introduces a subtlety into the interpretation of the device as a LSI.

5.4.1 Lateral shearing interferometry using a Mach-Zehnder followed by a lens

In the original ST-SPIDER, the interferometers and the detector lie in spaces, that, when viewed as an imaging system, are geometrically similar. This means that angles are preserved when transforming between the two domains. Less abstractly, it means that when a lateral shear with zero

tilt is applied in the interferometer, a lateral shear with zero tilt is what results on the detector. In the new design, this is not the case, because an imaging system separates the two. A lateral shear causes both a shear and a tilt in the crystal plane, which is re-imaged onto the detector. Although this hardly represents a drastic complication, I nonetheless found no analysis of the situation in the literature. This section provides an intuitive understanding of LSI in such a configuration.

The situation is depicted in Fig. 5.4(a). The MZI outputs two beams with a relative time delay, lateral shear and tilt. These pass through a lens and their interference pattern is detected by an imaging spectrometer. (In fact, as with the new ST-SPIDER design, the plane may D may be re-imaged, but this only results in a magnification of the interference pattern and so it suffices to analyse the single-lens case.) In "normal" LSI, the measurement procedure is to set the device up so that the shear on the detector is zero, and thereby obtain the carrier phase which results from the tilt only. Then, the interferometer shear Y is adjusted to the desired setting. However, because of the lens, the tilt and shear of the beams in the detector plane are some linear combination of the shear and tilt leaving of the interferometer. Adjusting Y changes the tilt at the detector, as shown by the green and blue lines in Fig. 5.4(a), adding a linear spatial phase to the interferogram. The reconstruction therefore acquires an additional quadratic term.

Instead of calculating the amplitude of the quadratic term, insight is gained by observing that the lens is re-imaging some object plane O onto D. The position of D is given by the imaging equation. Of course, the interferometer may stand between O and D, as depicted in Fig. 5.4(a). However, by imagining that the interferometer were not there, one can find the intersections of virtual rays cast backwards (as in Fig. 5.4(a)) or forwards from the lens with plane O. One can also backpropagate the corresponding virtual fields of the two beams. The fields at D are identical (in amplitude and phase) to the virtual fields at O, except for a magnification and a quadratic spatial phase produced by the imaging. In particular, the sum of the virtual fields of the two beams at O is itself an interference pattern. The object plane and the interferometer lie in geometrically similar spaces — there is no lens in between. This means that the interference pattern at O may be interpreted as a standard lateral shearing interferogram — adjustments to the interferometer shear setting do not affect the tilt. For convenience, Y is defined as the separation of the beams in
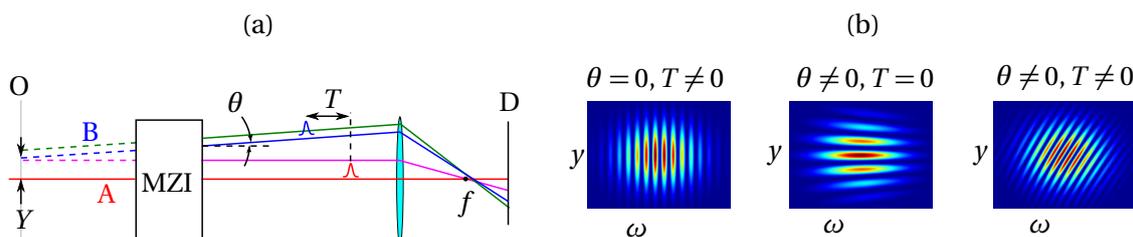
Figure 5.4: (a) Re-imaging LSI. The reflections caused by the mirrors of the A arm of the MZI have been applied, so that the beam traversing the A arm (red) travels in a straight line. The lens re-images object plane O onto the spatially and spatially resolving detector D. A delay T, tilt θ and shear Y is applied to the B arm output (blue). Some B arm outputs for different settings of shear and tilt (purple and green) are drawn for explanatory purposes in the text. (b) Sample fringe patterns for different interferometer settings.

the object plane. The spatio-spectral intensity profile of this *virtual interferogram* is re-imaged to D. This leads to the general result that LSI with a lens between the interferometer and the detector returns the phase of the field at the plane conjugate to the detector, with a spatial magnification factor. This may be generalized to an arbitrary imaging system. The spatial resolution is limited by the numerical aperture as seen from the object plane.

Figure 5.4(b)–(d) shows the nature of the interferograms formed for different delays and tilts. In particular, note Fig. 5.4(b) in which no tilt is applied and the fringes are purely spectral. This is somewhat counterintuitive since, as the purple line in Fig. 5.4(a) shows, even with no tilt in the interferometer the beams cross with a tilt in the detector. The paradox is resolved by considering the effect of the quadratic phase resulting from imaging O on D. This quadratic phase causes a linear term in the lateral shearing interferogram which exactly cancels the linear phase resulting from the tilt.

5.4.2 Calibration phase

As is common in SPIDER, the reference phase is recorded using the fundamental interferogram. The calibration phase should purely reflect the geometry and be independent of the unknown pulse. It therefore must be taken at zero spatial shear. If the spatial shear can be adjusted without changing the time delay, then this is straightforward. A more precise method is to double-pass the interferometer, a technique I shall now describe.

The arrangement is shown in Fig. 5.5. A plane mirror M_1 is placed in the unused output port. The mirror is perpendicular to the nominal beam direction. In fact, this mirror defines the plane of measurement. From the other output port, four double-passed beams emerge, labelled AA, AB, BA and BB by the order of the arms they traverse. For example, the path of beam AB is: reflection off beamsplitter BS_A, reflection off corner-cube CC_A, reflection off beamsplitter BS_B, reflection off mirror M_1, tranmission through BS_B, reflection off CC_B, and finally reflection off BS_A. Two mirrors M_2 and M_3 bring the double-passed beams back onto the main line for convenient measurement.

For the calibration, one must relate the time delay, tilt, and shear of the double-passed beams to those of the single pass. It is fairly intuitive that, with respect to the AA beam, AB and BA are delayed by T since they take one pass through the B arm, and BB is delayed by $2T$ since it takes two passes through the B arm. The shear and tilt are more challenging to visualize (for this author at least) but can be understood using the diagrams of Fig. 5.6. The beams are "straightened-out" and viewed side-on i.e. up on the page is always up in the lab, whilst right on the page is always "forward" for the beam under consideration. The corner-cubes (retro-reflectors) flip the beam in the vertical plane about their vertex, whilst the reflections apply a tilt. The effects of shear and tilt are additive and may be considered separately for simplicity. The incident beam is drawn straight and level with the vertex of corner-cube CC_A. However, since transformations of the incident beam apply equally to all the outputs, the implications of the diagrams are general. Corner cube CC_B is raised by an amount $Y/2$, whilst beam splitter B is rotated in the vertical plane by $\theta/2$.

To introduce the notation used in the diagrams, Fig. 5.6(a) shows the effects of adjusting Y and θ on the single-pass output B.

Figure 5.6(b) shows the effects of Y on the double-passed beams. The AB and BA outputs are sheared by Y in opposite directions, whilst adjustments to Y have no effect on the AA and BB outputs.

Figure 5.6(c) shows that adjusting θ does not affect BB, whilst AB and BA are deflected in opposite directions by θ. AA is deflected downwards by 2θ, which causes a z-dependent shear. However, as the construction lines show, the shear is zero in the plane of the M_1.

This is the basis of the calibration method — interference of the AA and BB outputs produces
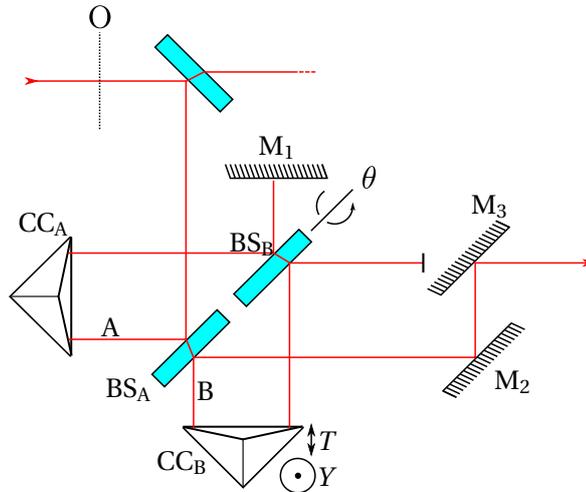
Figure 5.5: Double-pass arrangement for calibrating the tilt and time delay. Mirror M_1 sends the beams back through the interferometer. The double-passed beams exit from beamsplitter BS_A and are put back on the main beamline by mirrors M_2 and M_3, the latter being on a flip mount for convenient insertion and removal.

an interferogram with carrier $2\omega T - 2\theta k y$. If $\theta = 0$ or M_1 coincides with the object plane, then the interferogram has zero lateral shear, and is therefore a suitable calibration trace from which T and θ may be obtained. In principle, it is possible to position mirrors M_2 and M_3 so that the distance from the focusing lens to M_1 in the double-pass configuration is equal to the distance from the lens to the object plane in the single-pass configuration. Then, during calibration, the lens will re-image M_1 onto the detector, so AA and BB beams will always have zero shear.

5.4.3   Device operation and reconstruction

The details of the measurement protocol depend on the spectrometer. The present setup uses a Czerny-Turner imaging spectrometer with a diffraction grating. The fundamental and upconverted beams are measured in the first and second diffracted orders respectively and hence fall on the same pixels of the detector. One simple option is to obtain them sequentially. Alternatively, because the second order is twice as dispersive, the spectral fringes of the upconverted beams are twice as closely spaced as those of the fundamental. Therefore, provided the time delay is large enough, the fundamental and upconverted fringes may be simultaneously recorded and separated
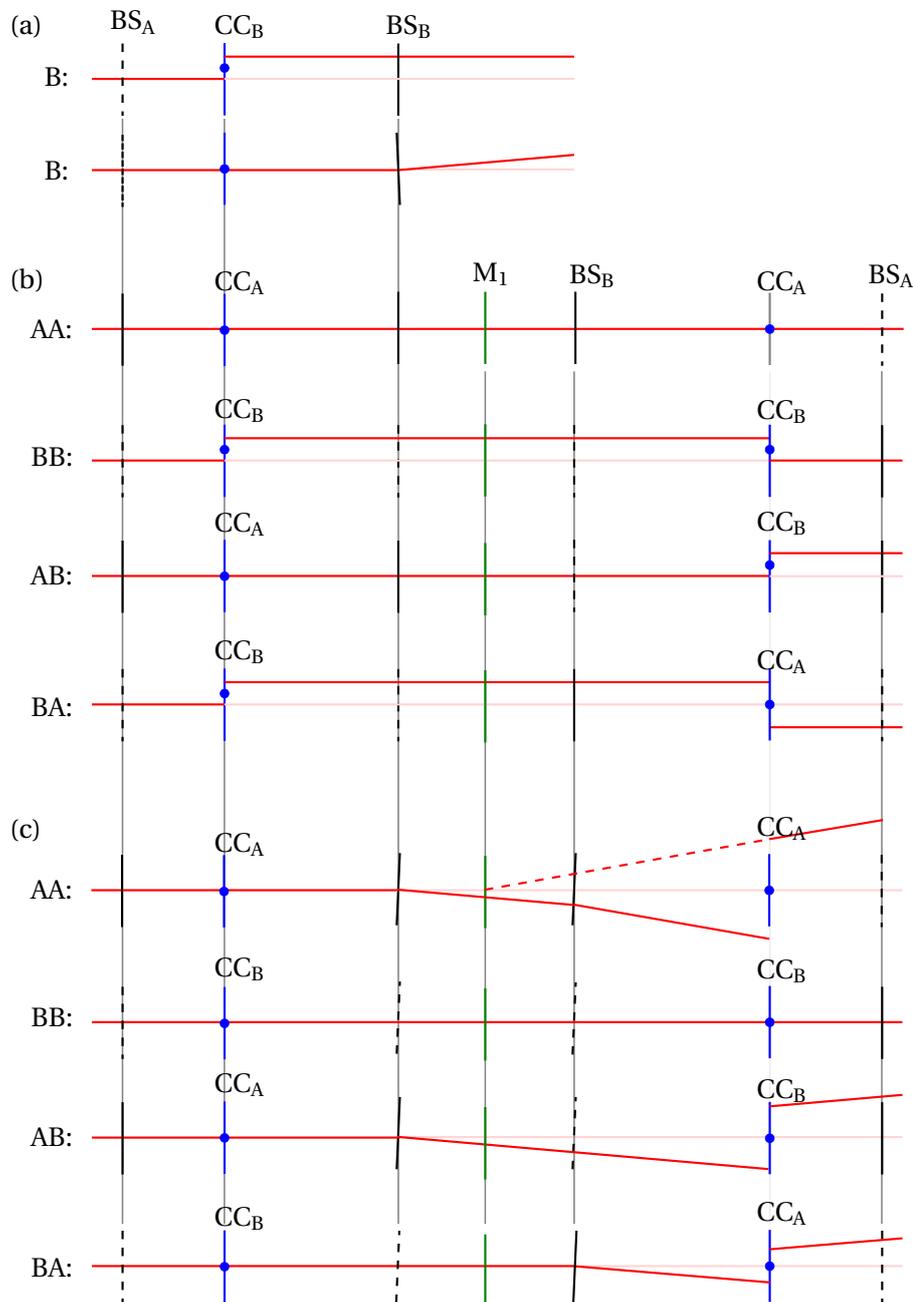
Figure 5.6: Analysis of vertical shear and tilt in double-passed MZI of Fig. 5.5. Beam-splitters are shown in reflection (solid black) and transmission (dashed black). The vertex of corner-cubes (blue) are shown with a dot. In each case, the A or AA beam, which does not experience any transformations, is shown faintly for reference. (a) Single-pass, with shear (upper) and tilt (lower). (b) Double-pass with shear. (c) Double-pass with tilt. For the AA beam, the output ray is projected backwards to the main axis.

in the Fourier filtering step.

After the usual Fourier filtering to isolate the sidebands and subtraction of the carrier, the phase of the fundamental interferogram is

$$\Lambda_{\text{f}}(\omega, y) = \phi(\omega, y - Y) - \phi(\omega, y) \tag{5.1}$$

whilst the phase of the upconverted interferogram is

$$\Lambda_{\text{up}}(\omega + \omega_{\text{up}}, y) = \phi(\omega - \Omega, y - Y) - \phi(\omega, y) + \eta(y) \tag{5.2}$$

Here, $\eta(y)$ is the spatially dependent phase of the upconverted interferogram, which depends on the ancilla phases. To avoid making assumptions about the uniformity of the ancilla spatial profile, one must treat $\eta(y)$ as unknown and incorporate this into the algorithm.

In developing a reconstruction algorithm, I note that much research has been done into two-dimensional wavefront reconstruction [295, 297–300] for Shack-Hartmann and shearing interferometers. Expressed in the present notation i.e. in the spatio-spectral domain, these algorithms act upon discretely sampled orthogonal phase differences

$$\Gamma_{\omega}(\omega_m, y_n) = \phi(\omega_{m+1}, y_n) - \phi(\omega_m, y_n) \tag{5.3}$$

and

$$\Gamma_y(\omega_m, y_n) = \phi(\omega_m, y_{n+1}) - \phi(\omega_m, y_n) \tag{5.4}$$

where $\omega_m = \omega_0 + m|\Omega|$ and $y_n = y_0 + n|Y|$ are the sampling points. To deploy these algorithms, one must convert the measured interferogram phases (5.1) and (5.2) into the forms of (5.3) and (5.4).

The spatial shear component of the upconverted interferogram is removed by subtracting the

fundamental phase:

$$\Lambda'_{\text{up}}(\omega, y) = \Lambda_{\text{up}}(\omega + \omega_{\text{up}}, y) - \Lambda_{\text{f}}(\omega, y) \tag{5.5}$$

$$= \phi(\omega - \Omega, y - Y) - \phi(\omega, y - Y) + \eta(y). \tag{5.6}$$

The shears are made positive using the following substitutions:

$$\Gamma_y(\omega, y) = \begin{cases} \Lambda_{\text{f}}(\omega, y) & Y < 0 \\ -\Lambda_{\text{f}}(\omega, y + Y) & Y > 0 \end{cases} \tag{5.7}$$

$$= \phi(\omega, y + |Y|) - \phi(\omega, y) \tag{5.8}$$

$$\Lambda''_{\text{up}}(\omega, y) = \begin{cases} \Lambda'_{\text{up}}(\omega, y + Y) & \Omega < 0 \\ -\Lambda'_{\text{up}}(\omega + \Omega, y + Y) & \Omega > 0 \end{cases} \tag{5.9}$$

$$= \phi(\omega + |\Omega|, y) - \phi(\omega, y) + \eta'(y). \tag{5.10}$$

Note that, although I am expressing this process in terms of the phase, in the numerical implementation it is generally better to work with the complex numbers, because i) interpolation and resampling routines are more robust in this form, and ii) the use of SNR-weighted reconstruction routines requires that one keep track of the signal intensity, as well as its phase.

After discretizing, the only remaining obstacle is the presence of the spatial-dependent phase $\eta'(y)$. This can be found by comparing the cross-derivatives, which, except for the absolute phase, are identical:

$$\left[\Lambda''_{\text{up}}(\omega, y + |Y|) - \Lambda''_{\text{up}}(\omega, y)\right] - \left[\Gamma_{\text{f}}(\omega + \Omega, y) - \Gamma_{\text{f}}(\omega, y)\right] = \eta'(y + Y) - \eta'(y) \tag{5.11}$$

Thus $\eta'(y)$ is chosen, by concatenation, to satisfy (5.11), up to a constant value which corresponds to the arrival time ambiguity of the pulse. Subtraction of $\eta'(y)$ from $\Lambda''_{\text{up}}(\omega, y)$ yields $\Gamma_\omega(\omega, y)$. The data are then ready for the application of wavefront reconstruction algorithms. Here, I use the weighted, pre-conditioned conjugate-gradient phase-unwrapping algorithm (3) of Ghiglia and

Romero [295].

5.5   Experimental example

I generated an angularly dispersed test pulse using an LAK21 prism, inducing an expected angular dispersion of $126 \pm 1\,\text{fs/mm}$, the uncertainty arising from the orientation of the prism. The apex of the prism was placed in the object plane O of the ST-SPIDER so that the spatial chirp introduced subsequent propagation would not be detected. The measurement protocol was:

1. Set spatial and spectral shear/time delay using translation stage.

2. Setup the ST-SPIDER to acquire the fundamental. This involved blocking the ancilla, removing the spatial and spectral filters and placing a neutral density filter in the beam.

3. Double-pass the interferometer and acquire a zero-shear interferogram at the fundamental wavelength.

4. Return to the single-pass configuration and acquire a fundamental interferogram.

5. Block the B arm and acquire the spectral intensity.

6. Setup the ST-SPIDER to acquire the upconverted pulses. This involves unblocking the ancilla, inserting the spatial and spectral filters, and removing the neutral density filter.

7. Acquire the upconverted interferogram.

Figures 5.7(a) and (b) show typical raw data, whilst Fig. 5.7(c) and (d) show the spatial and spectral phase difference after Fourier filtering and processing according to the procedure in section 5.4.3. The spatial phase differences, which approximate the local wavefront tilt, clearly show the expected linear dependence on frequency. Correspondingly, the spectral phase differences, which approximate the group delay, also show a linear dependence on position. They also show a linear dependence on frequency which is due to a small amount of dispersion present in the pulse.
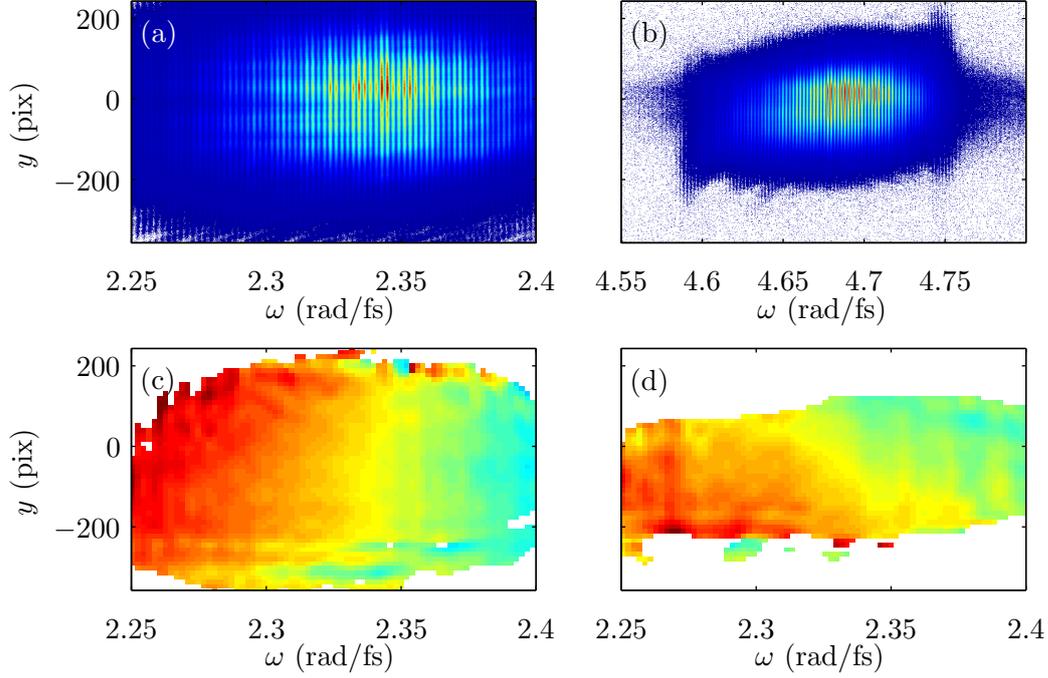
Figure 5.7: ST-SPIDER raw data and phase differences for an angularly dispersed pulse. (a) Raw spatially sheared fundamental. (b) Raw spatially and spectrally sheared upconverted pulse. (c) Spatial phase differences Γ_y. (d) Spectral phase differences Γ_ω.

The reconstruction is shown in Fig. 5.8. Fig. 5.8(a) shows the spatio-spectral phase. The hyperbolic contours are due to the angular dispersion which manifests in the spatio-spectral domain as a phase $\gamma(\omega - \omega_0)y$. Fitting the angular dispersion gives $127 \pm 1\,\mathrm{fs/mm}$, with the uncertainty taken by examining the shot-to-shot fluctuations over an ensemble of 25 measurements. Within experimental error this agrees with the expected value. The inferred spatio-temporal intensity is shown in Fig 5.8(b), whilst the angular spectrum is shown in Fig. 5.8(c). Also note that no spatial chirp was observed, verifying that the imaging system was correctly relaying the plane of the apex of the prism.

5.6 Effect of miscalibration

Interferometer instability and miscalibration can cause errors in the spectral and spatial carriers — the time delay and tilt between the beams. In spectral shearing interferometry (SSI), a time-
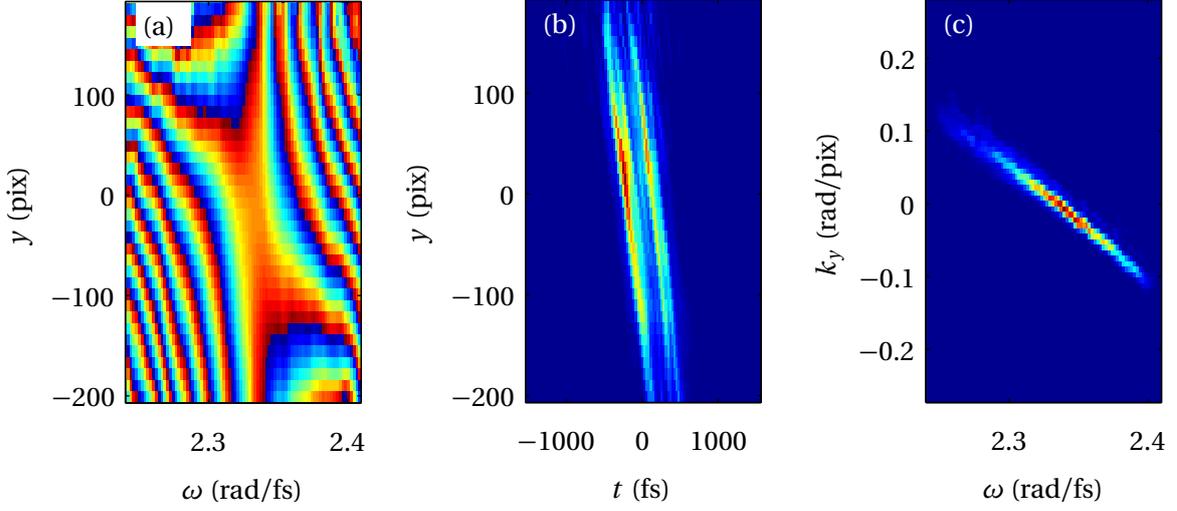
Figure 5.8: Angularly dispersed pulse reconstructed using ST-SPIDER. (a) Spatio-spectral phase, colour scale wrapped to $[-\pi, \pi]$. (b) Spatio-temporal intensity, linear colour scale. (c) Wavenumber-spectral intensity, linear colour scale.

delay error τ produces an erroneous quadratic phase. Analogously, in spatial shearing interferometry, an incorrectly calibrated tilt δ produces an erroneous quadratic phase. In the compact ST-SPIDER, one must distinguish between carrier errors that are common to the fundamental and upconverted interferogram, and errors that are different.

If a time-delay error is common to the fundamental and upconverted interferograms, as is always the case when the two are acquired simultaneously, then the subtraction step (5.5) removes the time-delay error from the upconverted interferogram. This mode of time-delay error therefore does not produce an erroneous quadratic phase. However it does affect the spatial phase reconstruction in a subtle way. The introduced linear spectral phase is $\Gamma_y(\omega, y) = \tau\omega + n2\pi$, where an unknown integer multiple of 2π is included because the phase is only obtained modulo 2π. Spatial concatenation of this erroneous term gives

$$\phi(\omega, y) = \frac{\tau\omega y}{Y} + \frac{n2\pi y}{Y}. \tag{5.12}$$

The first term corresponds to a beam tilt — purely a rotation of the co-ordinate system. This may or may not be important but is not normally considered a distortion. In particular it does

not correspond to pulse-front tilt, because the group front tilt $c\partial^2\phi/\partial y\partial\omega$ remains equal to the phase front tilt $c/\omega\partial\phi/\partial y$. However, the second term corresponds to a fictional pulse-front tilt of $cn2\pi/(\omega Y)$. Now n will be nonzero only if $|\tau| > \pi/\omega_0$. Therefore if the time-delay error is kept below half an optical period then it causes no distortions to the reconstructed pulse.

Any tilt error common to the fundamental and upconverted interferograms is also subtracted from the upconverted interferogram. It affects the spatial concatenation in the usual way, causing an erroneous quadratic spatial phase $\delta\omega y^2/(2cY)$.

Any time-delay error present in the upconverted interferogram, but not the fundamental interferogram, affects the spectral phase in the usual way, adding erroneous quadratic phase. Tilt error of a similar nature results in an inconsistency between the spatial and spectral shears which, if significant, will be observed in residuals of the reconstruction procedure. In the worst case, where one assumes that the incorrect upconverted interferogram has higher intensity and therefore outweighs the fundamental in the reconstruction algorithm, an erroneous pulse-front tilt of $\delta(\omega - \omega_0)y/(cY)$ results.

5.7   Summary and outlook

Spatio-temporal pulse measurement benefits from multidimensional spectral and spatial phase measurements, as opposed to combining a multidimensional spatial phase measurement with a spectral phase measurement of a single point in the beam. I presented a shearing interferometer capable of simultaneously obtaining the spectral and one spatial phase derivative of an ultrashort pulse from which the spatio-spectral phase may be reconstructed. The design is a geometrical simplification of ST-SPIDER, and in fact requires no substantial modifications to a conventional one-dimensional SPIDER except the use of an imaging spectrometer. The use of a single MZI ensures that the spectral and spatial gradients are measured at a common object plane with identical magnifications. The two derivatives may be obtained in a single-shot by carrier-multiplexing.

The realisation that simultaneous spectral and spatial shearing can be performed using a single interferometer makes spatio-temporal measurements a more practical proposition. If an imaging spectrometer is available, then wherever standard SPIDER setup is currently used, one can

now perform spatio-spectral measurements. The only modifications are extra optics for the double pass and supplying one arm of the interferometer with transverse motion. This will be useful for potentially space-time coupled sources such as filaments or strongly driven hollow-fibre compressors, to which one-dimensional SPIDER has already been applied.

Moving to few-cycle pulses should present little difficulty so long as a suitable imaging spectrometer is available. All dispersion up to the crystal must be calibrated and subtracted. The current design already uses mirrors rather than lenses for focusing, which also avoids chromatic aberrations. It would be necessary to calibrate the dispersion of the beam splitters used in the interferometer. A convenient way to do this would be to flip the orientation of one of them so that the interferometer becomes unbalanced and take a zero-shear interferogram at the fundamental wavelength.

This concludes the first part of this thesis. The second part deals with high-harmonic generation (HHG). The next chapter provides the necessary background on this topic.

6 High-harmonic generation

When the electric field of a laser becomes sufficiently strong, it may ionize an atom or molecule in a single nonresonant step. The laser field initially accelerates the electron away from the ion, before switching direction and driving the electron back towards the ion. The resulting collision has a number of effects — the electron may be scattered, the ion may be left in an excited state or doubly ionized and radiation is generated from the recombination of the electron and the ion. High-harmonic generation (HHG) is the third of these processes, and has been the subject of several decades worth of extensive research for several reasons. The emitted radiation is in the extended ultraviolet (XUV) and soft x-ray region and under certain conditions has subfemtosecond duration. High-harmonic generation is therefore the source of choice for attosecond science, and holds potential as a compact XUV source for imaging [301]. The emitted radiation also encodes certain properties of the target atom or molecule with subfemtosecond temporal resolution and sub-Ångström spatial resolution. Therefore, HHG serves as a probe of ultrafast dynamics of atomic and molecular structure.

In this dissertation, I present two new diagnostics for HHG. The first achieves *spectrally resolved wavefront sensing* for HHG, enabling the spatial amplitude and phase profiles to be determined at many frequencies simultaneously. I present a comprehensive set of measurements using this technique, and demonstrate good agreement with theory. The second diagnostic involves altering the generation process itself so that the contributions from different *quantum trajectories* — each corresponding to a particular electron ionization and return time — can be observed independently. One purpose of this chapter is to explain enough of the physics of HHG to justify why such measurements are useful, and also to interpret the results which I present. This chapter also outlines the current state of HHG metrology, in order to justify the development of new techniques.

The structure of this chapter is as follows: section 6.1 traces the historical development of HHG up to its present role as a useful source and diagnostic. Section 6.2 introduces the physics of HHG, starting with response of a single atom or molecule to the applied laser field. Section 6.3 discusses how the single-atom responses sum coherently to produce the observed field. Section 6.4 outlines

current methods of measuring the temporal and/or spectral profile of HHG, and section 6.5 does the same for the spatial profile. Section 6.6 then places the new results in context.

6.1 Development and applications

Advances in laser technology throughout the 1990s provided focused intensities above $10^{13}\,\mathrm{Wcm}^{-2}$. At these levels, irradiation of an atomic gas produces high-order harmonics in the XUV and soft x-ray regions [302–310]. Early results include production of the 17th harmonic of a 248 nm KrF laser in Neon [311], the 33rd harmonic of a 1064 nm laser in Ar [312], and the 135th harmonic of a 1053 nm laser in Neon [313]. High-harmonic generation was immediately recognized as a possible compact short-wavelength source, notwithstanding its relative weakness, and was used as a source for atomic core level spectroscopy [314], as a probe for plasma dynamics [315] and in time resolved x-ray fluorescence studies [316].

6.1.1 Attosecond pulse generation

The high harmonic generation process consists of the ionization and acceleration of electron wavepackets before their eventual radiative recombination with the atom. This process happens once in every half-cycle of the driving field, and the coherent superposition of these repetitions shapes the spectrum into odd harmonics. It was quickly realized that the individual half-cycle emissions potentially had subfemtosecond [317, 318] duration i.e. high harmonic emission actually consisted of a train of attosecond-duration pulses separated by half the laser period. Thus HHG seemed a promising embodiment of a scheme proposed by Hänsch [319] to produce pulses shorter than an optical period through Fourier synthesis. Trains of pulses with durations as short as 130 as [320, 321] have been observed.

Whilst the attosecond pulse train produced by a many-cycle drive field is a useful tool for studying ultrafast processes, a single attosecond pulse is required for general time-resolved measurements. An initial proposal was polarization gating, exploiting the fact that the electron only recollides in a linearly polarized laser field [322–325]. Initial experiments [326–328] succeeded in gating the harmonic emission to a duration of several optical cycles. More recent techniques

[329–333] have enabled the production of broadband XUV continua, indicative of emission from a single optical half-cycle.

The experiments mentioned thus far used many-cycle laser pulses. It was realized that the generation efficiency could be improved by reducing the pulse duration [334, 335]. Subsequent experimental efforts with 25 fs pulses [336], and then few-cycle pulses [337, 338], produced harmonics which reached the window of water transparency (2.3–4.4 nm) for the first time. It was at this time that the possibility of producing isolated attosecond pulses from a few-cycle drive field was raised theoretically [339]. In a few-cycle pulse, one half-cycle is significantly more intense than all the others, and hence one recollision event is significantly more energetic. The highest frequencies of the emission are therefore produced in a single event of attosecond duration and can be selected through high-pass spectral filtering. Using Ti:Sapphire laser systems with hollow fibre compressors, pulses of durations as short has 100 as [340–343] have been produced.

Attosecond pulses are finding ever-wider application because the importance of subfemtosecond dynamics is appreciated in a rapidly growing range of contexts. Examples include nanoplasmonics [344], charge migration in biomolecules [345], atomic inner-shell processes [346], and condensed matter physics [347]. These presage the emergence of *attoscience*, for which high-harmonic generation is an underpinning technology.

Besides aiming for ever shorter pulses, current research aims to increase the energy of the attosecond pulses, discover simpler and more robust means of generating them, and to shorten their wavelengths, particularly into the water transparency window (284–532 eV).

6.1.2 Structural and dynamical imaging

The radiation produced by HHG can also be used to infer properties of the atoms or molecules of the generating medium itself. The founding principle is the fact that the radiation is proportional to the dipole overlap of the returning electron wavepacket and the bound state from which the electron was ionized. The spatial resolution depends on the de Broglie wavelength of the returning electron, which is much smaller than the wavelength of the emitted photon. Furthermore, the electron wavepacket has significant momentum spread, thus sampling many spatial Fourier

components of the bound state. Temporal information can be inferred because of the brevity and chirp of the returning wavepacket, and the resolution may be significantly below 1 fs.

However, there are many challenges associated with extracting meaningful information from high-harmonic radiation. The properties of the returning electron wavepacket affect the radiation and must be known or accounted for in some way. A strong laser field is inherent in the process and its influence upon the sample is also a consideration. Finally, the structural and dynamical information in the measured harmonic spectra is often encoded in a subtle and complex manner. This section reviews various approaches.

In molecules, the ground state has many centres of higher electron density. The recombination dipole element is the coherent superposition of the dipole element resulting from each of these centres. Both the location of the centres [348–350] and the alignment of the molecule with respect to the laser polarisation [351–354] affect the interference of the emission from the centres.

Instead of looking for specific correspondences between certain features in the harmonic spectrum and the molecular orbitals, tomographic reconstruction [355] was used to image the orbitals by measuring a series of harmonic spectra at different alignment angles. The process is similar to computer-assisted tomography used in medicine. The interpretation of the molecular orbital created controversy [356], chiefly because of the single-electron approximation used in the reconstruction. Theoretical analysis revealed that tomographic reconstructions represent a mixture of molecular orbitals [357, 358] but for the nitrogen molecule used in the initial demonstration, the reconstructed orbital was a good qualitative picture of the highest-occupied molecular orbital in the Hartree-Fock picture. Multi-electron behaviour is expected to become more important as molecular size increases.

Concerning dynamics, one strand of research has focused on the evolution of the core (comprising the other electrons and the nuclei) during the excursion of the ionized electron. The state of the core is then sampled by the electron at the time of collision. This motivated examination of HHG spectra for signs of sub-fs dynamics; extending previous studies [359–361] of the kinetic energy of fragments produced by double-ionization caused by the returning electron. Another example is the sensitivity of HHG to the vibrational motion of the H_2^+ ion created when the process

is performed in hydrogen gas [362]. Because of the differing equilibrium internuclear separations of the molecule and the ion, the latter is left in an excited vibrational state after the ionization step of HHG. During the electron's excursion in the continuum, the nuclei begin to separate. Because the harmonics are modulated by the overlap of the time-evolved molecular ion wavepacket with the ground state molecular wavepacket, this separation manifests in a suppression of the emission. Experimental comparison of hydrogen and deuterium, the latter exhibiting slower nuclear motion due to its increased atomic mass, verified this effect [363, 364]. The intrinsic chirp of the returning electron wavepacket creates a mapping between excursion time and frequency in the harmonic spectra which simplifies the analysis of such experiments. This technique is known as probing attosecond dynamics by chirp encoded recollision (PACER).

Another proposal is to perform HHG in a system where rapid electronic motion has been previously excited [365]. The bound-state motion is reflected in the harmonic spectrum through its interplay with the chirp of the returning electron wavepacket. Such experiments offer the flexibility to probe the wave-packet motion of arbitrary excited states.

6.2   Theory of the single-atom response

This section commences the theoretical description of HHG. The process is understood on two conceptual levels. The physics of the *single-atom response* is the topic of this section. However it should be kept in mind that what is observed is the *macroscopic response* — the coherent sum over all single-atom responses within the interaction region. This aspect is discussed in section 6.3.

The simplest description of HHG is the three step model [366], illustrated in Fig. 6.1. One considers the most weakly bound electron of the target atom or molecule. The target is exposed to a laser field of sufficient intensity to significantly distort the electronic binding potential, causing *optical field ionization* within a single laser cycle (step one). The field accelerates the electron away from its ionized parent atom or molecule (step two). Upon reversal of the field direction after approximately one quarter-cycle, the electron is driven back towards the atom/molecule. As it passes the ion, it may recombine, releasing its acquired kinetic energy in the form of radiation (step three), which for typical laser intensities is in the extreme ultraviolet or soft x-ray region.
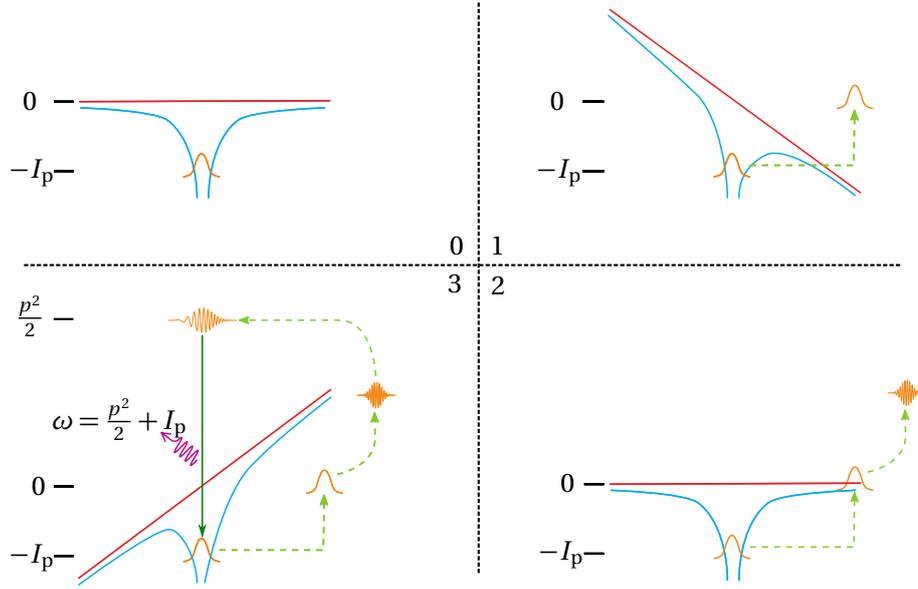
Figure 6.1: The three-step model of high-harmonic generation.

When discussing the single-atom response I use atomic units, in which the mass and charge of an electron, the reduced Planck constant $\hbar$ and Coulomb's constant are all set to unity.

6.2.1 Classical trajectory analysis

Many important features of the single-atom response may be predicted by treating the electron as a classical charged particle and following its motion in the time-varying electric field. One treats the ionized electron as if it were "born" with zero velocity at the nucleus, taken as the origin. After ionization, the effect of the Coloumb potential on the electron's motion is ignored. For electric field strengths below 10^{18} Wcm^{-2}, the electron velocity remains nonrelativistic and hence the optical magnetic field may also be ignored. The electron thus experiences a force $-\mathscr{E}(t)$. Integration of the equations of motion for an electron born with zero velocity at time $t_{\rm b}$ gives velocity $\mathbf{v}(t, t_{\rm b}) = \mathbf{A}(t) - \mathbf{A}(t_{\rm b})$ and position

$$\mathbf{x}(t, t_{\rm b}) = \int_{t_{\rm b}}^{t} \mathbf{A}(t')\, \mathrm{d}t' - \mathbf{A}(t_{\rm b})(t - t_{\rm b}) \tag{6.1}$$

where $\mathbf{A}(t) = -\int \mathscr{E}(t)\, \mathrm{d}t$ is the field vector potential. The recombination times are found by solving $\mathbf{x}(t, t_{\rm b}) = 0$ for $t > t_{\rm b}$. Depending on the birth time and the electric field profile, particularly its

polarization, the electron may return once, several times, or not at all. The return time is denoted $t_r(t_b)$. Each recombination produces a photon with energy equal to the sum of the kinetic energy of the electron and the ionization potential I_p of the bound state:

$$\omega_r(t_b) = I_p + v_r^2(t_b)/2 \tag{6.2}$$

where $\mathbf{v}_r(t_b) = \mathbf{v}(t_r(t_b), t_b)$ is the recombination velocity.

The existence, timing and recollision energy of the solutions give an intuitive picture of the generation process and the nature of the emitted radiation. The classical solutions are also quantitatively accurate in certain respects; for example, as shall be demonstrated, the locus of points $(t_r(t_b), \omega_r(t_b))$ does accurately reflect the time-frequency distribution of the harmonics. For this reason I shall discuss in detail the classical solutions in the paragraphs below.

For single colour fields[1], recollision will only occur if the field is linearly polarized. Although it is possible to engineer recollision in two-colour fields with complex polarizations [367], this is not relevant to this dissertation and so I restrict the discussion to linearly polarized fields.

I first consider the simple case of a monochromatic field and then discuss how the situation changes for a short pulse. Some sample trajectories in a monochromatic linearly polarized laser field are shown in Fig. 6.2(b); the field itself is shown in Fig. 6.2(a). The entire process repeats itself, with a change of sign, every half cycle of the laser, and so for clarity I have only drawn trajectories starting in the first half-cycle. The odd-numbered harmonic spectrum follows from this repetition using Fourier analysis. Following the trajectories in Fig. 6.2(b) it can be seen that only these electrons born in the first quarter-cycle return to the core. Of these some return more than once. (Up to three recollisions are shown in Fig. 6.2.) For reasons going beyond the classical description, these *higher order* trajectories generally play a minor role. Electrons born at the start of the cycle return 2π radians later; later birth times give earlier return times.

The return kinetic energy is plotted against the return time in Fig. 6.2(c). The *cutoff energies*

---

[1]This definition requires clarification for pulses. Generally, a pulse is said to be n-colour if it contains n spectral regions each separated, as a minimum, by the inverse of the transform-limited duration of the pulse. For example, a field consisting of the sum of two 30 fs pulses at 800 nm and 400 nm is two colour. However a few-cycle pulse whose spectrum spans 400 nm and 800 nm is single colour.
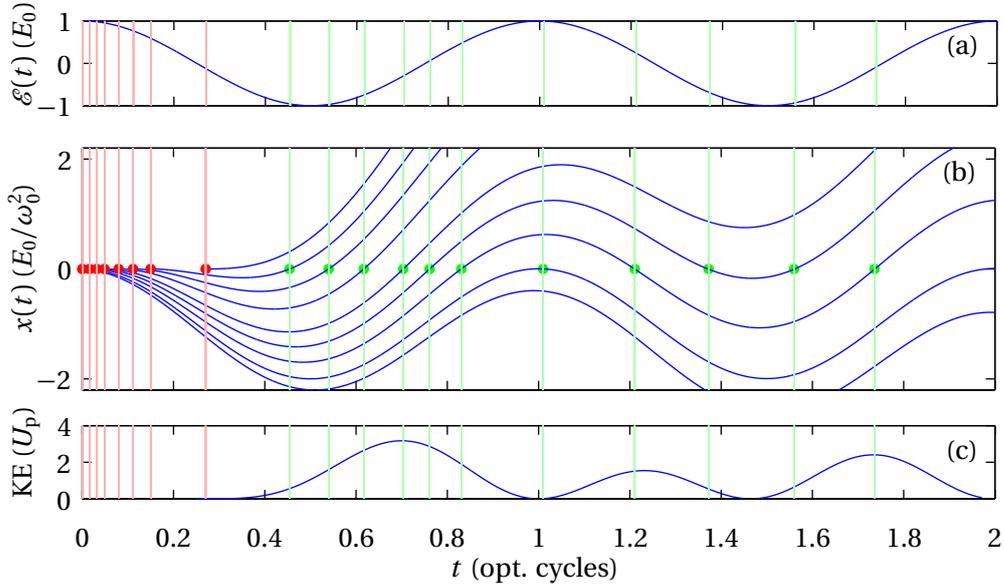
Figure 6.2: Classical picture of high-harmonic generation. All the axes are in normalized units. (a) Drive laser field. (b) Selected electron trajectories born during the first half cycle (-0.25–0.25 optical cycles). Births times (red dots) and recollision times (green dots) are shown. Faint lines guide the eye to the other plots. (c) Return kinetic energy normalized to the pondermotive energy versus recollision time.

$\omega_{c,max}$ are local maxima of the photon energy with respect to the birth or recollision time. There is one cutoff for each half-cycle of the laser field and for each recollision. For a sinusoidal electric field of frequency ω_0 and amplitude E_0, the cutoff energy for the first recollision is $I_p + 3.17U_p$ [368], where U_p is the *ponderomotive energy* $U_p = E_0^2/4\omega_0^2$. The corresponding birth and recollision phases are $\omega_0 t_b = 18°$ and $\omega_0 t_r = 252°$. This result also applies to slowly varying envelopes where the intensity does not change greatly during a single half-cycle. For any photon energy below the cutoff, there are two pairs of birth and recollision times, known as the short and the long trajectories.

For pulses, every half-cycle has a different intensity. For many-cycle pulses, each half-cycle may be treated as approximately monochromatic. In this *adiabatic* picture, the harmonic emission can be considered as the sum over many monochromatic half-cycles, so that, for example, the cut-off law applies to each. The electron trajectories are identical except for a scaling factor between each half-cycle. For few-cycle pulses however, the envelope influences the trajectories,
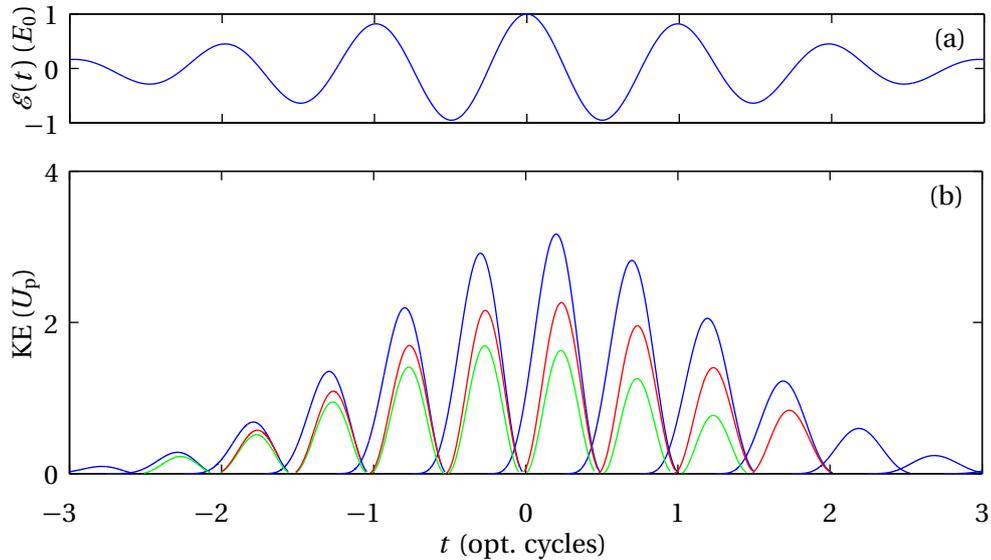
Figure 6.3: Classical trajectory analysis for 2.62 optical cycle pulse (equivalent to 7.0 fs pulse at 800 nm). (a) Electric field; (b) Recollision kinetic energies for first (blue), second (green) and third (red) returns.

and for an accurate picture one must treat each half-cycle individually.

Figure 6.3 shows a classical trajectory analysis for a 7.0 fs pulse at 800 nm, with a carrier-envelope phase offset of zero. The greatest recollision energy is attained by electrons born one half-cycle before the peak of the envelope. For photon energies below the maximum, several trajectories contribute.

The classical description is physically intuitive and gives a quantitative prediction of the structure of the emitted radiation in the time-frequency domain (t, ω). That is, it predicts significant energy at time-frequency co-ordinates $(t_r(t_b), \omega_r(t_b))$ for all birth times that result in a recollision. However it does not attach an amplitude or a phase to the emitted radiation — this requires the incorporation of wave-like properties of the electron [366]. The quantum-mechanical evolution of the electron is most accurately described by the time-dependent Schrödinger equation, which I shall introduce in the next section. However it should be noted that many features of the classical model are recovered by applying a realistic set of assumptions to the time-dependent Schrödinger equation (TDSE). These assumptions constitute the strong field approximation (SFA) and the

stationary-phase approximation (SPA), and shall be described in a later section.

6.2.2   Time-dependent Schrödinger equation

The TDSE in the length gauge and with the dipole approximation applied is

$$i\frac{\partial|\psi\rangle}{\partial t} = \left[\frac{\hat{\mathbf{p}}^2}{2} + V(\mathbf{x}) + \mathscr{E}(t)\cdot\mathbf{x}\right]|\psi\rangle \tag{6.3}$$

where $\hat{\mathbf{p}} = -i\nabla_{\mathbf{x}}$ is the momentum operator, $V(\mathbf{x})$ is the potential, $\mathscr{E}(t)$ the external field and $|\psi\rangle$ is the electron wavefunction. Several simplifications lie behind (6.3). A single active electron is considered [369]. This gives accurate results for atoms but generally needs modification for molecules. The potential here is static, a reasonable approximation for atoms. In molecules the time-evolution of the potential, caused, for example, by nuclear motion, can play a significant role. The electric field is homogeneous; this follows from the dipole approximation, valid as long as the electron excursion is signifcantly smaller than the wavelength. As mentioned in section 6.2.1, the magnetic field is ignored since the electron motion is nonrelativistic.

The result of a TDSE simulation of HHG is shown in Fig. 6.4. The parameters of this simulation are used in several examples throughout this chapter. At peaks of the electric field ionization is clearly visible: a significant component of the wavefunction detaches from the bound state which is correspondingly depleted. Recombination events are rather subtle, manifesting in a slight 'ripple' of the ground state where the returning electron trajectories overlap. This ripple leads to a high frequency oscillation of the dipole moment, which is the source of the harmonics. This illustrates the weakness of the high harmonic generation process, with typical conversion efficiencies being 10^{-6}.

The radiation emitted by a single atom is proportional to the acceleration of the dipole moment $\mathscr{D}(t) = \langle\psi(t)|\mathbf{x}|\psi(t)\rangle$ of the electron. Although $\mathscr{D}(t)$ can be obtained directly from the calculation results for $|\psi(t)\rangle$ and twice differentiated, it is instructive to use Ehrenfest's theorem to relate the dipole acceleration to the expectation value of the force

$$\frac{\partial^2\mathscr{D}(t)}{\partial t^2} = \langle\psi(t)|-\nabla_{\mathbf{x}}V(x)|\psi(t)\rangle + \mathscr{E}(t). \tag{6.4}$$
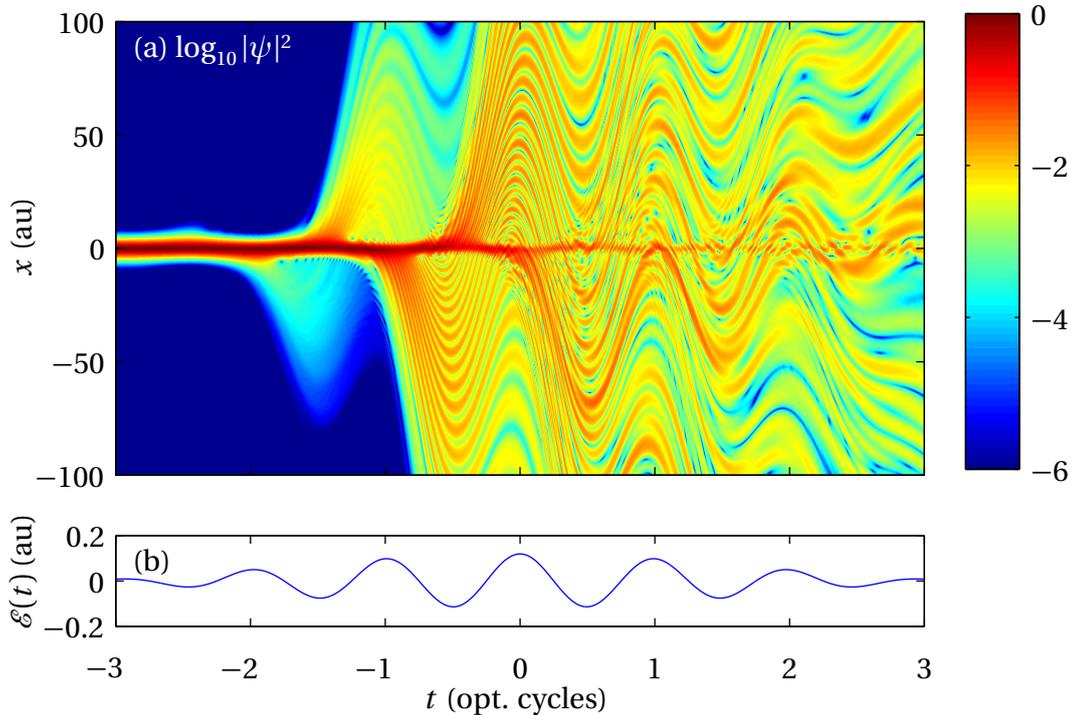
Figure 6.4: (a) Amplitude-squared of electron wavefunction in HHG computed using the 1D TDSE for a 7 fs pulse at 800 nm with peak intensity 5×10^{14} Wcm^{-2} in Argon. The "soft-core" potential $V(x) = -1/\sqrt{x^2 + a^2}$ is used, where $a = 1.1892$ is chosen such that the binding energy of the ground state equals the ionization potential of Argon. (b) Drive laser field.

The contribution $\mathcal{E}(t)$ is at the fundamental frequency; therefore the harmonics are solely produced by acceleration of the electron caused by the potential.

The complexity of the emitted radiation is such that plots of the temporal and spectral intensity are difficult to interpret. Time-frequency distributions are generally more meaningful. Figure 6.5 shows the spectrogram (defined in section 2.2.3.1) of the dipole acceleration for the simulation of Fig. 6.4.

6.2.3 Strong field approximation

Although the TDSE is the most quantitative model of HHG, greater physical insight and reduced computational demands are possible using the formalism of the SFA, also referred to as the Lewenstein model in this context [368]. The SFA incorporates one of the assumptions of the classical model, namely that the laser field is so strong that for ionized electrons the potential of the core
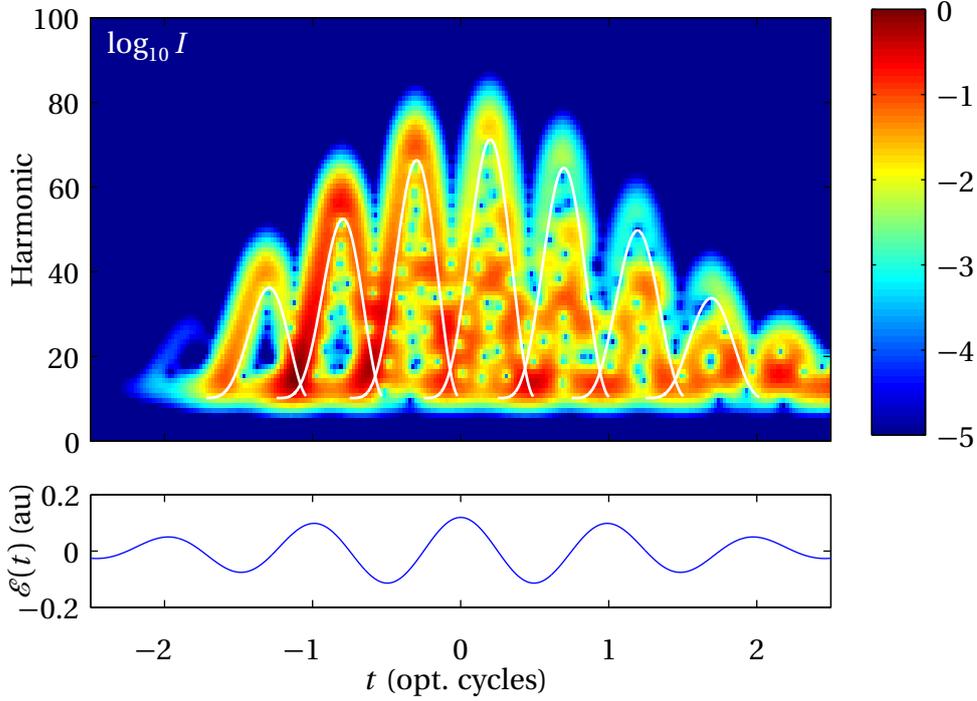
Figure 6.5: Spectrogram of high-harmonic generation using a 7 fs pulse at 800 nm with peak intensity 5×10^{14} Wcm^{-2} in Argon. The photon energy from a classical calculation (solid white lines) and a quantum trajectories calculation (dashed white lines) is shown.

can be ignored. One introduces the ansatz

$$|\psi\rangle = e^{iI_{\mathrm{p}}t}\left[g(t)|g\rangle + \int \mathrm{d}^3\mathbf{v}f(\mathbf{v},t)|\mathbf{v}\rangle\right] \tag{6.5}$$

which comprises a ground state $|g\rangle$ with amplitude $g(t)$ and a continuum of plane waves $|\mathbf{v}\rangle$ with amplitude $f(\mathbf{v},t)$. The oscillations of the ground state are factored out. The key idea behind the SFA is that the laser field transfers population from the ground state into the continuum states; the ground state is otherwise unperturbed. The continuum states are driven only by the laser field; on them the effect of the ionic potential is ignored.

Substituting (6.5) into (6.3) and applying the bra $\langle\mathbf{v}|$ yields

$$i\frac{\partial f(\mathbf{v},t)}{\partial t} = g(t)\mathscr{E}(t)\cdot\mathbf{d}(\mathbf{v}) + \left(\frac{v^2}{2} + I_{\mathrm{p}}\right)f(\mathbf{v},t) + i\mathscr{E}(t)\cdot\nabla_{\mathbf{v}}f(\mathbf{v},t). \tag{6.6}$$

where $\mathbf{d}(\mathbf{v}) = \langle \mathbf{v}|\mathbf{x}|g \rangle$ is the atomic dipole matrix element for a bound-free transition. I have assumed that the orbital is symmetric so that $\langle g|\mathbf{x}|g \rangle = 0$ and also that the relevant continuum states are highly energetic so that $\langle \mathbf{v}|g \rangle \approx 0$. At this stage I assume that a solution for $g(t)$ can be found or approximated by other means, such as the Ammosov, Delone, and Kraïnov (ADK) formula for tunnel ionization [370]. Equation (6.6) can be transformed into an ordinary differential equation by writing $f(\mathbf{v}, t) = b(\mathbf{p}, t)$, where $\mathbf{p} = \mathbf{v} - \mathbf{A}(t)$ is the *canonical momentum*. Upon performing this substitution, (6.6) is solved to yield

$$b(\mathbf{p}, t) = -i \int_{-\infty}^{t} dt_{\mathrm{b}} \mathscr{A}_{\mathrm{ion}}(t_{\mathrm{b}}, \mathbf{p}) e^{iS(t, t_{\mathrm{b}}, \mathbf{p})} \tag{6.7}$$

where

$$S(t, t_{\mathrm{b}}, \mathbf{p}) = -\int_{t_{\mathrm{b}}}^{t} dt' \left\{ \frac{[\mathbf{p} + \mathbf{A}(t')]^2}{2} + I_{\mathrm{p}} \right\} \tag{6.8}$$

and

$$\mathscr{A}_{\mathrm{ion}}(t_{\mathrm{b}}, \mathbf{p}) = g(t_{\mathrm{b}}) \mathscr{E}(t_{\mathrm{b}}) \cdot \mathbf{d}(\mathbf{p} + \mathbf{A}(t_{\mathrm{b}})). \tag{6.9}$$

Equation (6.7) can be interpreted as follows: electrons of canonical momentum $\mathbf{p}$ are ionized into the continuum at time t_{b} with amplitude $\mathscr{A}_{\mathrm{ion}}(t_{\mathrm{b}}, \mathbf{p})$, which is the product of the ground state amplitude, the electric field strength and the amplitude of the corresponding dipole transition. Because the canonical momentum $\mathbf{p}$ is an invariant of the motion of a free charged particle in an time-varying electric field, the *Volkov states* $|\mathbf{p}\rangle$ are eigenstates of the field Hamiltonian. They have time-varying eigenvalues and their evolution in the field is represented by a phase $S(t, t_{\mathrm{b}}, \mathbf{p})$. The ionization potential enters (6.8) because the oscillations of the ground state are factored out in (6.5).

The harmonic emission results from rapid oscillations of the atomic dipole $\mathscr{D}(t)$, computed using the ansatz (6.5) and the solution (6.7). This gives

$$\mathscr{D}(t) = -i \int d^3\mathbf{p} \int_{-\infty}^{t} dt_{\mathrm{b}} \mathscr{A}_{\mathrm{ion}}(t_{\mathrm{b}}, \mathbf{p}) e^{iS(t, t_{\mathrm{b}}, \mathbf{p})} \mathscr{A}_{\mathrm{rec}}(t, \mathbf{p}) + \mathrm{c.c.} \tag{6.10}$$

where

$$\mathscr{A}_{\text{rec}}(t, \mathbf{p}) = g^*(t)\mathbf{d}^*(\mathbf{p} + \mathbf{A}(t)). \tag{6.11}$$

Here, I have ignored dipole transitions between the continuum states because they generally do not contribute significantly to the harmonic emission, although this currently a topic of theoretical investigation [371]. Equation (6.10) may be interpreted in a similar way to (6.7): electrons are ionized at t_{b}, acquire a phase $S(t, t_{\text{b}}, \mathbf{p})$ during their excursion in the continuum, and then recombine with the core at time t with rate $\mathscr{A}_{\text{rec}}(t, \mathbf{p})$ given by the product of the ground state amplitude and the dipole transition. All possible combinations of birth time and momenta must be considered and are therefore integrated over; their contributions add coherently.

The phase $S(t, t_{\text{b}}, \mathbf{p})$ plays a fundamental role in HHG, and it may be interpreted in several ways. One may view it as the accumulated phase difference between the continuum and ground state wavefunctions during the electron's sojourn in the continuum. The contribution of the kinetic energy integral is the continuum contribution; it is the integral of the time-dependent eigenvalue corresponding to the Volkov eigenstate of the field Hamiltonian. Specifically, the solution to the potential-free TDSE is $|\mathbf{p}\rangle e^{S_{\mathbf{p}}(t)}$ where

$$S_{\mathbf{p}}(t) = -\int_{\infty}^{t} \mathrm{d}t' \frac{[\mathbf{p} + A(t')]^2}{2}. \tag{6.12}$$

The contribution to $S(t, t_{\text{b}}, \mathbf{p})$ proportional to the ionization potential simply arises from the free oscillation of the ground state during the electron's motion in the continuum. A second point of view is to interpret the phase as a semiclassical action [368]. The SFA then folds into the Feynman path integral picture of quantum mechanics, in which all possible system evolutions from the unperturbed ground state, through the continuum, and recombining back to the ground state, are summed together with a phase factor given by the action.

The spectrum of the radiation is found by taking the Fourier transform of the dipole expectation:

$$\mathscr{D}(\omega) = \frac{-i}{\sqrt{2\pi}} \int_{-\infty}^{\infty} \mathrm{d}t_{\text{r}} \int \mathrm{d}^3\mathbf{p} \int_{-\infty}^{t} \mathrm{d}t_{\text{b}} \mathscr{A}_{\text{ion}}(t_{\text{b}}, \mathbf{p}) e^{iS_{\omega}(t_{\text{r}}, t_{\text{b}}, \mathbf{p})} \mathscr{A}_{\text{rec}}(t_{\text{r}}, \mathbf{p}) + \text{c.c.} \tag{6.13}$$

where

$$S_\omega(t_r, t_b, \mathbf{p}) = -\int_{t_b}^{t_r} dt \left\{ \frac{[\mathbf{p} + \mathbf{A}(t)]^2}{2} + I_p \right\} + \omega t_r \tag{6.14}$$

is simply the action modified to include the phase ωt_r of the emitted photon. Equation (6.13) expresses the radation as an integral over all birth times, momenta and return times.

6.2.4 Quantum trajectories

One aspect of the semi-classical description which is *not* recovered by (6.13) is exactly which birth times, return times, and momenta contribute significantly to the harmonic emission. However, in (6.13), the action has a much greater dependence on the integration variables t_b, t_r, and $\mathbf{p}$ than the transition amplitudes. Since the cycle-averaged integral of a rapidly varying phase is approximately zero, significant contributions to the integral only occur for those values of t_b, t_r, and $\mathbf{p}$ for which the action is stationary. The stationary-point approximation [372] provides a way of formalising this intuitive notion. Specifically, the dipole response integral (6.13) is approximated by a finite sum over each of the stationary points. Besides reducing the computational complexity, the stationary-point approximation provides physical insight and makes explicit the connection between the quantum and classical pictures of high-harmonic generation.

Requiring the birth time, return time, and momentum to be stationary leads to three conditions, each with a physical interpretation that justifies the assumptions made in the classical three-step picture.

- Saddle points of the momentum integral occur when

$$\nabla_{\mathbf{p}} S_\omega(t_r, t_b, \mathbf{p}) = \mathbf{x}(t_r) - \mathbf{x}(t_b) = 0 \tag{6.15}$$

which shows that strong contributions occur only when the electron returns to the position at which it was born. One thus recovers one of the assumptions of the classical description. This equation can be solved directly to give

$$\mathbf{p} = \frac{-1}{t_r - t_b} \int_{t_b}^{t_r} dt \mathbf{A}(t). \tag{6.16}$$

- Saddle points of the birth time integral yield

$$\frac{\partial S_\omega(t_r, t_b, \mathbf{p})}{\partial t_b} = \frac{v^2(t_b)}{2} + I_p = 0 \tag{6.17}$$

  which under a classical interpretation suggests that electrons are born with negative kinetic energy. Clearly, (6.17) has no real solutions. However, by analytically continuing the domain of all the functions to the entire complex plane, complex-valued solutions can be found. The classically nonsensical state of possessing negative kinetic energy is interpreted as tunneling of the evanescent quantum wavefunction under a classically disallowed barrier. The imaginary part of the birth time is interpreted as a tunneling time [368, 370]. The assumption of zero kinetic energy at birth in the classical picture is "as close as possible" to this intrinsically nonclassical process.

- Saddle points of the recombination time integral give

$$\frac{\partial S_\omega(t_r, t_b, \mathbf{p})}{\partial t_r} = \frac{v^2(t_r)}{2} + I_p - \omega = 0 \tag{6.18}$$

  which shows that the photon energy equals the total energy relinquished by the electron upon recombination.

The equivalence of the physical interpretations of the stationary-point conditions and the assumptions of the classical model lends a comforting mutual consistency between these two approximate pictures of high-harmonic generation.

In general, for a given ω, the stationary-point equations will have multiple solutions, which I label with superscripts so that the solution j is $\{t_b^{(j)}(\omega), t_r^{(j)}(\omega), \mathbf{p}^{(j)}(\omega)\}$. For notational simplicity I define $S_\omega^{(j)}$ as the corresponding action, and $\mathscr{A}_{\mathrm{ion}}^{(j)}(\omega)$, $\mathscr{A}_{\mathrm{rec}}^{(j)}(\omega)$ the corresponding ionization and recombination amplitudes. The stationary points may be systematically classified according to their birth and return times. I shall use the superscript label $j = \alpha\beta m$, similar to references [373–375]. In this labelling system, $\alpha = \mathrm{S}$ or L denotes whether the orbit is a short or a long trajectory. Next, $\beta = 1, 2\ldots$ indicates the number of returns the electron has made to the core, and hence
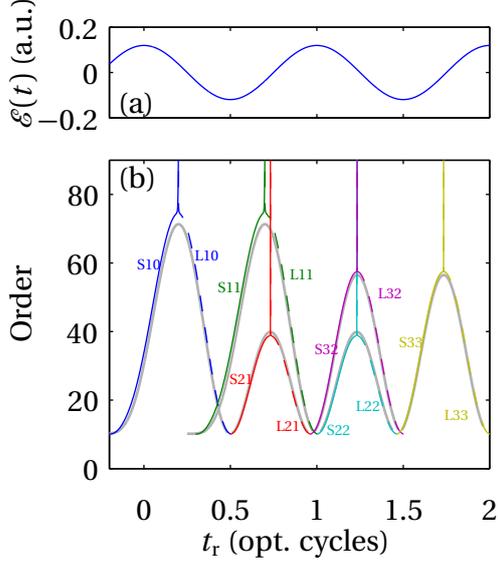
Figure 6.6: (a) Drive laser field. (b) Harmonic order versus recombination time of the quantum orbits of a CW drive field of $5 \times 10^{14}\,\mathrm{W/cm^2}$ in Argon. The orbits are labelled by type (short/long), excursion length and recombination event $\alpha\beta m$. Short trajectories are solid, long trajectories are dashed. The classical solutions are also shown (grey).
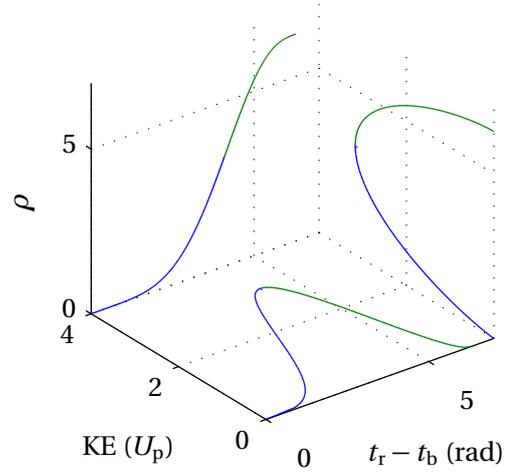
Figure 6.7: Intensity-derivative of the action for a monochromatic field, for the short (blue) and long (green) trajectories.

the trip time, with $\beta = 1$ being the shortest. Finally, $m = \dots, -1, 0, 1, \dots$ labels the recombination times, with $m = 0$ being the first recombination after $t = 0$. Another useful label is the birth event $\gamma = m - \beta$.

Figure 6.6 shows the emission frequency versus the real part of the return time for a CW field, and also illustrates the trajectory labelling system. Additionally, the classical trajectories are shown, and agree closely with the stationary-point solutions.

In the stationary-phase approximation, the contribution from an individual stationary point is proportional to the value of the integrand at that point, but modified by a prefactor which depends inversely on the local second derivative of the integration variables. Details of the derivation are

given in references [373, 375, 376]. The contribution from the jth saddle point is

$$\mathscr{D}^{(j)}(\omega) = \sqrt{\frac{(2\pi)^4}{i\left[t_{\mathrm{r}}^{(j)}(\omega) - t_{\mathrm{b}}^{(j)}(\omega)\right]^3 \det S_{\mathrm{s}}''\left[t_{\mathrm{r}}^{(j)}(\omega), t_{\mathrm{b}}^{(j)}(\omega)\right]}} \mathscr{A}_{\mathrm{ion}}^{(j)}(\omega) \mathscr{A}_{\mathrm{rec}}^{(j)}(\omega) e^{iS_\omega^{(j)}} \qquad (6.19)$$

$$= \mathscr{B}^{(j)}(\omega) e^{iS_\omega^{(j)}}$$

where I have introduced the symbol $\mathscr{B}^{(j)}(\omega)$ for later purposes. Equation (6.19) permits a similar interpretation to (6.13), with the additional prefactor accounting for spreading of the wavepacket due to diffraction.

One complication of the quantum orbit model is that not all of the stationary point solutions are valid. In consequence, the net dipole response is *not* simply the sum of the individual contributions, each given by (6.19). Examining Figures 6.5 and 6.6, one observes that for frequencies between the ionization potential and the classical cut-off, the long and short trajectories are distinct, both in terms of the exact TDSE solution as represented in the spectrogram and the quantum trajectories. Approaching the cutoff, the two recombination events shown in the spectrogram coalesce into one and this is also reflected in the real part of the recombination times of the quantum trajectories. However, the trajectories do not actually merge into one — there are still two stationary points, even above the cutoff. One of these, the *recessive* trajectory, is unphysical — its contribution to the amplitude diverges rapidly. Mathematically, this solution is no longer reached by the steepest-descent contour integral used in the stationary-point approximation. The other, *dominant* solution, possesses the appropriate rapid decay and accurately approximates the integral above the cutoff.

The preceeding discussion shows that the stationary-point approximation to the net dipole response is the sum over all trajectories well below the cutoff and the sum over all dominant trajectories well above the cutoff. Around the cutoff, the *uniform approximation*, which deals with coalescing stationary-points, is necessary for an accurate result. The uniform approximation to

the combined contribution from coalescing stationary-points $S\beta m$ and $L\beta m$ is [373, 376]

$$\mathscr{D}^{(\beta m)}(\omega) = \sqrt{6\pi S_-}\, e^{i(S_+ + \pi/4)} \left[\frac{\mathscr{B}_-}{z}\, \mathrm{Ai}'(-z) + i\, \frac{\mathscr{B}_+}{\sqrt{z}}\, \mathrm{Ai}(-z) \right] \tag{6.20}$$

where

$$S_\pm = \frac{1}{2} \left[S_\omega^{(S\beta m)} \pm S_\omega^{(L\beta m)} \right], \; \mathscr{B}_\pm = \frac{1}{2} \left[\mathscr{B}^{(S\beta m)}(\omega) \pm i\eta\, \mathscr{B}^{(L\beta m)}(\omega) \right],$$

$$z = (\frac{3S_-}{2})^{2/3} \cdot \begin{cases} 1 & \omega < \omega_{AS}^{(\beta m)} \\ e^{i2\eta\pi/3} & \omega >= \omega_{AS}^{(\beta m)} \end{cases}.$$

Here, Ai and Ai$'$ are the Airy function and its derivative respectively, and $\eta = 1$ if the short trajectory is dominant; otherwise $\eta = -1$. The cutoff frequency $\omega_{AS}^{(\beta m)}$ is formally defined as the frequency of the *anti-Stokes transition*, at which $\mathrm{Im}[S^{(S\beta m)}] = \mathrm{Im}[S^{(L\beta m)}]$.

6.2.5 Intensity-dependence of the action

One important and very general property of the quantum paths is the dependence of their phase on the intensity of the laser. From the definition of the phase (6.14), one can show that the intensity derivative at a stationary-point is:

$$\frac{\partial S_\omega^{(j)}}{\partial I} = -\int_{t_b^{(j)}(\omega)}^{t_r^{(j)}(\omega)} \frac{\left[\bar{\mathbf{p}}^{(j)}(\omega) + \bar{\mathbf{A}}(t') \right]^2}{2}\, \mathrm{d}t' \tag{6.21}$$

where $\bar{\mathbf{p}}$ and $\bar{\mathbf{A}}$ are the values of the momentum and vector potential normalized to unity intensity i.e. $I = 1$. This shows that the intensity-derivative of the phase of the single-atom response (at a given frequency and for a given trajectory) depends on the frequency solely through the birth and return times. Of course, these vary with intensity, but this dependence is quite weak except around the cutoff. For example, above the cutoff the birth and return phases remain fixed at 18° and 252°. For low emission frequencies the short trajectory birth and return phases both tend towards 90°, and those of the long trajectory tend towards 0 and 360°. Therefore the phase often exhibits a near-linear dependence on the intensity. In fact, using appropriate scaling and assuming that the harmonic frequency is much greater than the ionization energy, one can show that the intensity

derivative can be completely summarized by a normalized coefficient $\rho^{(j)}$, defined by

$$\frac{\partial S_\omega^{(j)}}{\partial I} = -\frac{\rho^{(j)}\left[\frac{\omega - I_p}{U_p}\right]}{4\omega_L^3}.\tag{6.22}$$

I have used this definition for consistency with a recent review [377]. The intensity coefficient for a monochromatic field is plotted against excursion time and return kinetic energy in Fig. 6.7. One can see that ρ rises monotonically with the excursion time, and is therefore greater for the long trajectories. For an 800 nm field at cutoff, the intensity-derivative takes the value of $12.8 \times 10^{-14}\,\text{W}^{-1}\text{cm}^2$ in SI units[1]. The intensity-dependent dipole phase plays an important role in the macroscopic effects which shape the single-atom response into what is actually observed. These are the subject of the next section.

6.3   Theory of the macroscopic response

The field produced HHG is the coherent superposition of the single-atom response from every particle in the target, and is strongly influenced by the nature of this superposition. In understanding this one must consider the variation of the laser field throughout the interaction region, how this in turn leads to a variation in the amplitude and phase of the single-atom response, and how these varying single-atom responses combine to give the observed field.

In this discussion it is useful to distinguish between purely transverse macroscopic effects, arising from variations in the single-atom response in a single transverse plane, and longitudinal macroscopic effects, arising from the superposition of the response from different transverse planes. One can think of the transverse macroscopic response as that which would be observed if the gas target were very thin. An important transverse macroscopic effect arises from the combination of the radial decay of the laser intensity with the differing intensity-dependent phases of the long and short trajectories [378]. This is illustrated in Fig. 6.8. Both trajectories have a divergent wavefront (Fig. 6.8(a)), but for the long trajectory the effect is greater. At the focus, this

---

[1]In SI units, (6.22) is

$$\frac{\partial S_\omega^{(j)}}{\partial I} = -\frac{\rho^{(j)}}{\omega_L^3}\frac{Z_0}{2E_{au}^2\,t_{au}^3}\tag{6.23}$$

where Z_0 is the vacuum impedance and E_{au} and t_{au} are the atomic units of electric field strength and time respectively.
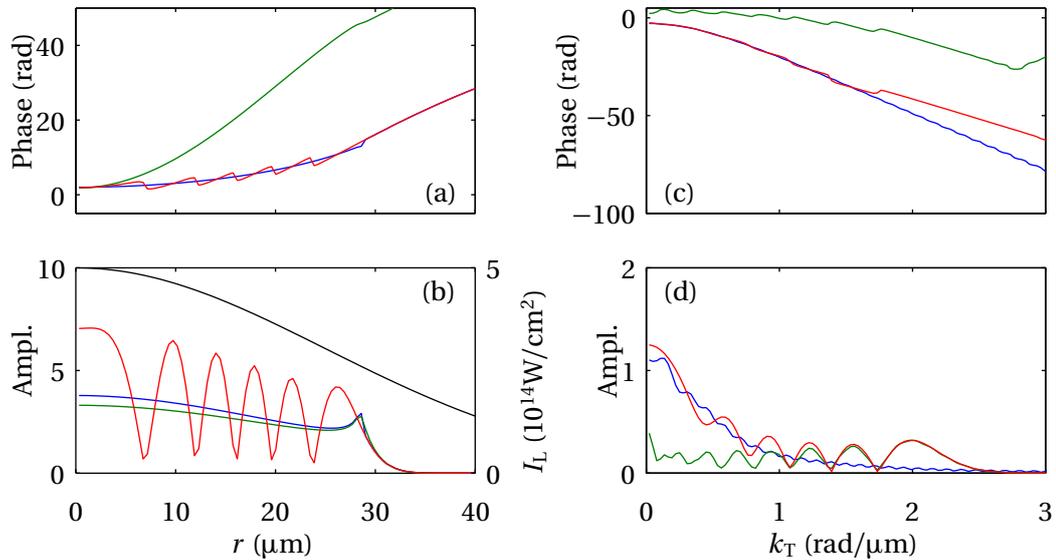
Figure 6.8: Transverse profiles at the beam waist of harmonic 43 produced by one half-cycle of a $5 \times 10^{14}\,\text{W/cm}^2$ 800 nm pulse with 50 µm beam waist in Argon. The short trajectory (blue), long trajectory (green) and their uniform-approximation sum (red) are shown. (a) Near-field harmonic phase. (b) Near-field harmonic amplitude (left vertical axis), and laser intensity (black, right vertical axis). (c) Far-field phase. (d) Far-field amplitude. The irregular shapes of the short and long trajectory far-field profiles are largely caused by their cusps in the near field profile — these are a numerical artifact that disappears when the trajectories are summed using the uniform approximation.

causes radial interference fringes in the sum of the long and short trajectories (Fig. 6.8(b)) until the intensity drops below the cutoff and the trajectories coalesce. After propagation to the far-field, (Fig. 6.8(c) and (d)), the short trajectory emission remains on-axis whilst the long trajectory emission has an annular profile. At intermediate divergence angles, with significant contributions from both trajectories, interference may also be observed [379].

The most important longitudinal macroscopic effect is phase matching, which often plays a critical role in HHG. A crucial difference between HHG and perturbative processes such as second-harmonic generation (SHG) is the intensity dependence of the phase in the former. This allows the intensity profile of the laser field to affect the phase matching. This in turn leads to different phase-matching conditions for the different trajectories. Therefore, macroscopic effects can select certain quantum paths and suppress others, leading to a "cleaning" of the harmonic spectrum. Under certain conditions this selection can be controlled and exploited.

Phase matching is governed by the difference between the wave-vector of the spatially varying single-atom response $\mathbf{k}_S = \nabla\phi_S(x, y, z, \omega)$ and the wavenumber of freely propagating harmonics $\beta = n(\omega)\omega/c$. The generation is most efficient when $k_S = |\mathbf{k}_S| = \beta$. The direction of the generated radiation is $\mathbf{k}_S$. The *coherence length* $L_{\mathrm{coh}} = \pi/|k_S - \beta|$ is the distance over which the single-atom response becomes π out of phase with propagating harmonics. If, as is often the case, the coherence length is shorter than other limiting distances such as the gas jet thickness or the absorption length, then the generation is phase-mismatched and the coherence length sets the effective generation length.

Much insight can be gained by considering the various contributions to the single-atom response wave-vector [380, 381]. Generally, there are several significant contributions. Spatial variations in the laser intensity leads to variation in the phase of the single-atom response, as discussed in section 6.2.5. This gives a contribution $-\rho\nabla I_L$ to $\mathbf{k}_S$. The phase ϕ_L of the focused laser, which includes the converging/diverging wavefront, gives a contribution $q\nabla\phi_L$, where q is the harmonic order. The dispersion of the gas at both the fundamental and harmonic wavelengths may also play a role. Summing these effects enables one to determine the location and direction of efficient harmonic generation within the laser focus.

Figure 6.9 shows the various contributions to the axial phase mismatch Δk_z for the same parameters as Fig. 6.8. The Gouy phase gives a negative contribution over the entire focus, whilst the intensity-dependent phase changes sign at the focus as the beam goes from converging to diverging. For the chosen pressure of 0.01 atm, the dispersive contributions are small (they can easily become significant at higher but experimentally realistic pressures). Under these conditions, on-axis phase matching may occur downstream of the focus, where the intensity-dependent and Gouy phase contributions can cancel. The higher intensity-dependence of the long trajectory means this cancellation occurs closer to the focus than for the short trajectory. Moving towards the cutoff (higher harmonics), the trajectories become more similar in character. However, the Gouy phase contribution becomes larger. The interplay of these effects means that the conditions for optimal phase matching depend on both frequency and the desired trajectory.

Additional phase-matching possibilities arise away from the axis, where the radial variation of
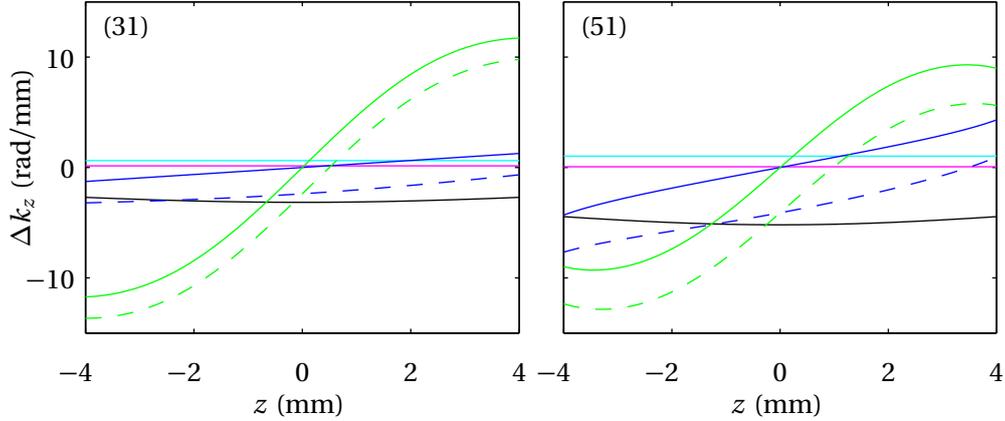
Figure 6.9: Axial phase-mismatch for harmonics 31 (left) and 51 (right) for a 5×10^{14} W/cm^2 800 nm field of waist 50 μm in 0.01 atm Argon. The contributions are: dispersion at laser wavelength (cyan), dispersion at harmonic wavelength (magenta), Gouy phase (black), and the intensity-dependence of the short (blue) and long (green) trajectories. The total phase-mismatch (dashed) is shown for the short (blue dashed) and long (green dashed) trajectories.

the laser intensity makes a contribution. Figure 6.10 shows the coherence length and the direction of the single-atom response wave-vector $\mathbf{k}_S$ using the same parameters as in Fig. 6.9. Away from the axis, phase matching can occur behind the focus, with the emitted radiation being divergent. This is particularly true for the long trajectories at lower harmonics, where the intensity-dependent phase is greatest.

As the laser pulse duration decreases to the single-cycle regime, and also as the laser intensity intensity rises to produce a significant fraction of ionised atoms, more complicated longitudinal macroscopic efforts occur. For example, even in the absence of dispersion and nonlinearities, the carrier-envelope phase of a few-cycle pulses changes by π as the pulse goes through a focus due to the Guoy phase [382]. For few-cycle pulses it is inaccurate to treat each half-cycle as if it were part of a monochromatic pulse — instead, the electric field itself, which depends on the envelope and the carrier-envelope phase, must be taken into consideration. Also, since HHG is intrinsically linked to ionization, it results in a free-electron plasma which affects the refractive index [383] in a time-dependent fashion. The free-electron contribution to the refractive index is much greater for the drive field than the harmonics, and can significantly alter phase matching [384, 385] in a spatially nonuniform manner. Furthermore, since ionization occurs during the pulse, the refrac-
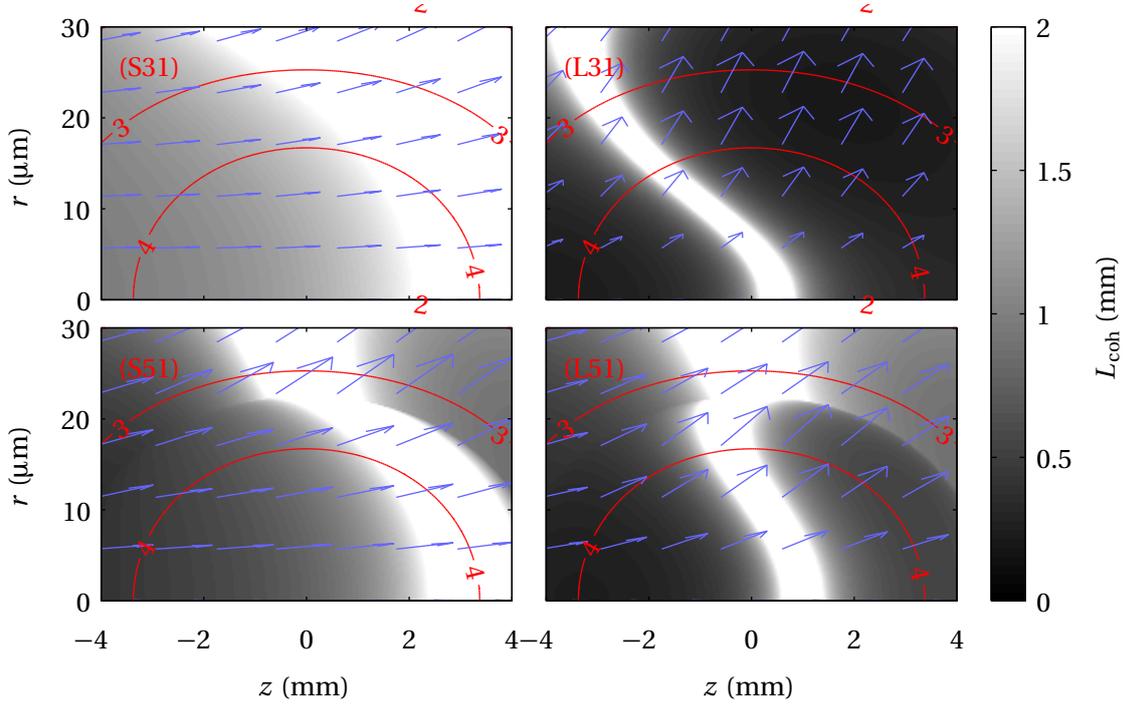
Figure 6.10: Phase-matching diagram for short trajectories (left column) and long trajectories (right column) and harmonics 31 (top row) and 51 (bottom row) for a $5 \times 10^{14}\,\mathrm{W/cm^2}$ 800 nm field of waist $50\,\mu\mathrm{m}$ in 0.01 atm Argon. The purple arrows show the single-atom response wave-vector (the arrows are scaled by the same factors as the axes of the plots, so that the off-axis component is greatly exaggerated.) The contour lines show the laser intensity in units of $10^{14}\,\mathrm{W/cm^2}$.

tive index changes dynamically, providing an ultrafast temporal phase modulation. This causes a blue-shift, which is most pronounced at the most intense regions of the pulse where ionization occurs rapidly [386–388]. For few-cycle pulses this temporal phase modulation actually assists with phase matching for high photon energies [384, 389]. Ionization also causes defocusing since the center of the beam is more intense, leading to greater ionization and a lower refractive index. Dispersion and absorption from neutral atoms play a significant role as the gaseous region becomes longer than a few millimeters [390–392]. Self-phase modulation may also play a role, particularly with molecular targets [393].

To include all of this behaviour, a numerical code is needed. Any such code consists of two parts. A propagation code for the drive field, which incorporates diffraction, dispersion and non-linear effects such as ionization, outputs the temporal profile of the drive field at all points in the interaction region. This is fed into a single-atom response code, which computes the single-atom

response at all points in the interaction region. The individual single-atom responses are summed together to calculate the final result.

6.4 Temporal metrology of high-harmonic generation

The principles of femtosecond metrology in the optical and IR also apply to HHG, but there are substantially greater challenges due to the weakness of the radiation and the lack of suitable optical components such as mirrors, beamsplitters and nonlinear crystals.

Since a well-characterized reference is usually unavailable, HHG temporal characterization methods must be self-referencing, and therefore require a time-nonstationary or nonlinear operation. Although the usual second and third order frequency mixing processes used for optical and IR characterization are not currently available for HHG, there are several alternatives. The most widely used is two-colour ionization using both the XUV pulse and the laser. In this method, the two fields are focused into an atomic gas where the XUV produces photoelectrons due to single-photon absorption. The amplitude and phase of the XUV field are mapped directly onto the wave-function of photoelectrons. The photoelectrons are then *streaked* — shifted in momentum by the electric field of a laser pulse — and the photoelectron spectrum is monitored as a function of the XUV-laser delay using, for example, a time of flight (TOF) electron spectrometer. A typical setup for photoelectron streaking is shown in Fig. 6.11.

Using the framework of the SFA, one can show that the streaking acts as a temporal phase modulation. The derivation mirrors that leading to the continuum amplitude (6.7), except that optical field ionization by the laser is replaced by photoionization by the XUV. Recombination is ignored, and one is interested in the momentum distribution after the fields have passed i.e. $t \to \infty$. The result is

$$b(\mathbf{p}, \infty) = -i \int_{-\infty}^{\infty} \mathbf{d}(\mathbf{p} + \mathbf{A}(t_{\mathrm{b}})) \cdot \mathbf{E}_{\mathrm{X}}(t_{\mathrm{b}}) \exp\left[i\phi_{\mathrm{s}}(t_{\mathrm{b}}) + i(W + I_{\mathrm{p}})t_{\mathrm{b}}\right] \, \mathrm{d}t_{\mathrm{b}} \qquad (6.24)$$

where $W = \frac{\mathbf{p}^2}{2}$ is the final energy of the photoelectrons. The phase modulation is

$$\phi_{\mathrm{s}}(t) = -\int_t^\infty \mathbf{p} \cdot \mathbf{A}(t') + \frac{\mathbf{A}(t')^2}{2} \, \mathrm{d}t'. \tag{6.25}$$

The photoelectron spectrometer measures projections of $|b(\mathbf{p}, \infty)|^2$. In the absence of the laser field, $\phi_{\mathrm{s}}(t) = 0$ and (6.24) is simply the Fourier transform of the product of the XUV pulse and the dipole transition element, evaluated at frequency $\omega = W + I_{\mathrm{p}}$. Generally, the dipole transition element is taken as given and may be calibrated out. The laser field applies a temporal phase modulation $\phi_{\mathrm{s}}(t)$ to the XUV field. One immediately infers that methods for the optical and IR which use a temporal phase modulation, such as spectral shearing interferometry (SSI) and tomography, may be implemented for HHG using two-colour photoionization. Whilst to a certain extent this has indeed occurred, unique properties of HHG mean that several new approaches have been invented.

I will now discuss the main elements of HHG temporal metrology.

6.4.1 Linear time-stationary measurements

XUV spectrometry. In the extended ultraviolet and soft x-ray region, the need to keep reflections to a minimum and use grazing incidence results in modifications to standard spectrometer designs in the optical and infrared region. The roles of focusing and dispersive optics are often combined. One common design is the grazing incidence flat-field spectrometer [394, 395], shown in Fig. 6.12. The grating has a cylindrical shape and smoothly varying groove spacing, which together focus the different frequencies onto a plane. The detector may be a charge-coupled device (CCD) or a micro-channel plate (MCP).

Photoelectron spectrometry. When an XUV photon ionizes an atom, its amplitude and phase is mapped directly onto the wavefunction of the resulting photoelectron. The energy spectrum of the photoelectron can be then be measured. Most photoelectron spectrometers for HHG use the TOF method [396, 397], in which the momentum of a photoelectron is inferred from its arrival time. Electrostatic lens systems collect and collimate the photoelectrons from a gas target, and
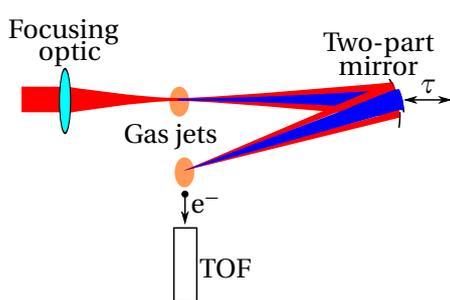
Figure 6.11: Typical setup for two-colour photoelectron spectroscopy.
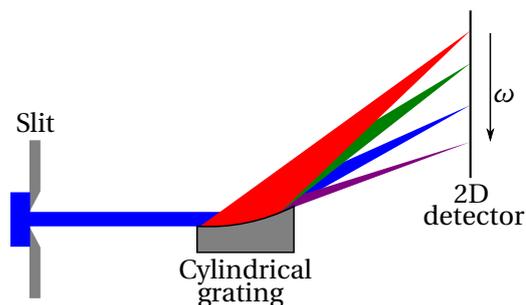
Figure 6.12: Flat-field XUV spectrometer.

can be used to retard or accelerate the electrons for a desired spectral range and resolution.

6.4.2 Nonlinear autocorrelations

Intensity autocorrelations have been implemented using two-photon absorption [398, 399]. Two separate harmonic sources are produced, and they are focused into an atomic gas. The spectrum of the resulting photoelectrons is monitored. In addition to the usual single-photon absorption, a two-photon absorption signal is observed, proportional to the intensity overlap of the pulses. Scanning their relative delay whilst recording the area of the two-photon signal yields an intensity autocorrelation trace. However, the signal is extremely weak and therefore requires long integration times. This, along with the fact that the intensity autocorrelation does not provide a complete characterization, has prevented these methods from becoming widely used.

6.4.3 Two-colour photoelectron spectroscopy

A broad class of techniques is based on the general principle of measuring the laser-dressed photoelectron spectrum as a function of laser-XUV delay. The methods differ in terms of their interpretation and the relative durations of the laser and XUV pulses.

If the XUV consists of a train of near-identical pulses, and the streaking laser is that used to produce the harmonics, then the period of the train is exactly half the period of the laser, and in the frequency domain the train is a series of discrete harmonics separated by $2\omega_L$. In this situation, the phase modulation produces sidebands separated by $\pm\omega_L$, above and below each harmonic. The upper sideband of harmonic q thus interferes with the lower sideband of harmonic $q+2$, and

their relative phase is $\phi_{q+2} - \phi_q - 2\phi_L$. Either by scanning the phase of the laser or sampling it appropriately, the relative phase of adjacent harmonics may be determined by the usual interferometric methods (section 2.3.3.2), yielding the temporal profile of the subpulses. This is the basis of the reconstruction of attosecond beating by interference of two-photon transitions (RABBIT) technique [320, 400].

At the other extreme, the XUV field consists of a single pulse, shorter than half the laser period. If the XUV pulse is cotemporal with a peak of the laser field, it experiences a quadratic phase modulation. A series of photoelectron spectra taken with different modulation amplitudes provide the necessary observations for temporal reconstructions based on a time-to-frequency mapping [342, 401, 402] or tomography [403].

A more general method is frequency-resolved optical gating (FROG)-complete reconstruction of attosecond bursts (CRAB), in which the photoelectron spectrum is sampled across the full range of pulse overlap. Similarly to the case of the spectrogram, the two-dimensional dataset may be inverted using the principal component generalized projections algorithm (PCGPA) [404, 405]. In principle, arbitrarily complex pulses may be retrieved using this method, although to date the measured pulses have been relatively simple, possessing smooth polynomial phases and little substructure. In addition to the generic challenges of measuring complex pulses using spectography discussed in section 3.1, FROG-CRAB traces have long acquisition times placing stringent demands on the stability of the system.

6.4.4 Spectral shearing photoelectron interferometry

If two time-delayed XUV pulses are used to generate two photoelectron wavepackets, an interference pattern is recorded on the photoelectron spectrometer. Spectral shearing interferometry may then be performed by choosing the dressing laser field such that the photoelectrons experience a spectral shear. One approach, suitable for XUV pulses that are much shorter than an optical period, is to align one XUV pulse with a zero-crossing of the laser field, where it experiences a linear temporal modulation which produces a spectral shear [406]. The other XUV pulse is aligned with a peak of the field, where the linear temporal phase is zero and the quadratic component assumed

to be negligibly slow. Another proposed method uses the time-varying ponderomotive energy of the laser as the temporal phase modulation [407]. Finally, when the XUV pulses are longer than an optical cycle so that the temporal phase modulation produces sidebands, a chirped laser will produce sidebands with different separations. The interference of these sidebands produces a spectral shearing interferometry trace [403].

Disadvantages of these methods include the requirement to generate two XUV pulses, the need for precise calibration, the need for good photoelectron spectrometer resolution, which increases the acquisition time, and the relative phase ambiguity between separate spectral components, which is particularly relevant for the highly structured spectra generated by HHG.

6.4.5 Generation of spectrally sheared replicas

Motivated by the experimental complexity and long acquisition times of photoelectron spectrometers, there have been several efforts towards all-optical characterisation using SSI in which the shear is performed using the generation process itself. Specifically, two harmonic sources are generated using spectrally sheared laser pulses. Intuitively, one expects the spectral shear to be carried over to the generated XUV pulses, and single-atom calculations for 30 fs drive pulses bear this out [403, 408]. There have been no systematic studies of the accuracy of this assumption for few-cycle pulses or with macroscopic effects included. There has been one experimental demonstration, in which collinear time-delayed spectrally sheared laser pulses are produced using an acousto-optic pulse shaper [409]. The disadvantage of this approach is that its temporal carrier imposes a high resolution requirement on the spectrometer, and the intensity of pulses must be low enough so that ionization produced by the first pulse does not affect the second. An proposed alternative is a spatially encoded arrangement, in which the two laser pulses are brought to spatially separate focii with nominally zero time delay [403, 408].

6.5 Spatial metrology of high-harmonic generation

At least two methods of HHG wavefront characterisation have been reported. Point-diffraction interferometry [410], described in section 2.4.1.2, involves a semitransparent membrane with a

small pinhole. The beam which diffracts out from the pinhole serves as a spherical reference for the beam which is transmitted through the membrane. An alternative is an adaption of the Shack-Hartmann sensor to the XUV [411]. However, neither of these techniques are spectrally resolved.

One spectrally resolved technique has been demonstrated: spectral wavefront optical reconstruction by diffraction (SWORD) [412], in which the diffraction pattern from a slit is measured using a spatially resolving spectrometer. The slit is scanned across the beam, and far-field spectra recorded at each point. The technique is approximately a spatial analog of spectography, with the slit and diffraction in SWORD playing the respective roles of the gate and spectrometer in spectography. As such it is inherently a multi-shot method.

6.6   Outline of this dissertation

The work presented in chapters 7 and 8 can now be introduced.

Both the single-atom (section 6.2) and macroscopic (section 6.3) theories of HHG predict non-trivial space-time coupling. One way of characterising space-time coupling is through a spatially resolved wavefront measurement. Currently, such techniques are limited (section 6.5). Chapter 7 presents a new method: spectrally resolved lateral-shearing interferometry for HHG. I develop the theory of this method and present a comprehensive set of experimental results, in good agreement with theory.

The quantum-orbits model (section 6.2.4) is a powerful description of HHG, combining classically intuitive dynamics with inherently quantum behaviour such as the interference between trajectories. In chapter 8, I present a means of using this interference to determine the relative amplitudes and phases of the contributions from individual quantum trajectories. The technique is analogous to multi-dimensional Fourier-transform interferometry. The relative phase between the interfering terms is modulated by varying the amplitude and phase of a weak control field, which induces a perturbative change in the action of each trajectory. This is completely different from the established forms of HHG metrology, which measure the generated field rather than disecting the process itself, and do not return enough information about the field to permit a complete quantum-path decomposition.

7 Lateral shearing interferometry for high-harmonic generation

This chapter presents a method for performing spectrally resolved spatial wavefront measurements on extended ultraviolet (XUV) and soft x-ray fields produced by high-harmonic generation (HHG). Its purposes are in common with all pulse metrology: to illuminate the physics of the generation process, and to characterize the source for use in applications. The method is based on spectrally resolved lateral shearing interferometry (LSI): the spatial wavefront derivative is measured as a function of frequency and position by comparing the phase of spatially displaced replicas using interferometry. This chapter derives the theory of the method and describes an experimental implementation in which the wavefronts of harmonics 13–25 in Krypton are studied as a function of the gas jet position relative to the laser focus. The data exhibit signatures of both the single-atom response and macroscopic effects: the intensity-dependence of the atomic dipole amplitude and phase, the imprint of the laser phase, and the propagation in the gas jet all have an influence.

The work of this chapter builds upon an initial set of LSI measurements [266, 413, 414] performed before my arrival in the Walmsley group. The experiments were conducted with Tobias Witting.

7.1 Motivations

Before discussing the reasons for performing spectrally resolved wavefront characterisation, it is worth discussing what these terms mean in the context of HHG. In fact, although the methods differ, the underlying principles and challenges are no different from the optical and infrared pulses considered in the first half of this dissertation. A wavefront is described by its amplitude and phase; in HHG, as ever, determining the latter is the greater obstacle. Any meaningful description of the wavefront of the highly space-time coupled pulses resulting from HHG is necessarily spectrally resolved; one therefore takes as given a spatially resolving spectrometer, capable of re-

turning the intensity $\tilde{I}(\omega, x, y)$[1]. The aim of a spectrally resolved wavefront measurement is to cause the phase $\phi(\omega, x, y)$, up to an unknown function of frequency $C(\omega)$, to be imparted onto the measured intensity. In the remainder of this section I will discuss some reasons for acquiring this information.

The first reason is exposing the physics of HHG itself. In the comparison of theory with experiment, phase information is particularly important in the context of HHG because the radiation is usually measured in the far-field, after significant diffraction from the generation region. However, the single-atom response applies directly in the spatial domain in the near-field, and its predictions of the amplitude and phase are largely decoupled. The amplitude depends on the ionization rate, electron diffraction in the continuum, and the recombination element, whilst the phase is dominated by the action integral in the continuum, with the laser phase also making a contribution. The near-field amplitude and phase is therefore the best basis for comparison of experiment with single-atom theories; using the far-field intensity distribution for this purpose is both indirect and incomplete. Although macroscopic effects may complicate this picture, they often act as trajectory "filters", favouring certain electron trajectories whilst suppressing others. The single-atom response, appropriately modified, is still a good description in such a situation.

Despite two decades of study, much remains unknown about HHG. The physics of HHG consists of three distinct conceptual units: the propagation of the driving laser field, the single-atom response at each point in the interaction region, and the macroscopic summation of the single-atom responses. Only of the third, which involves linear propagation in a tenuous dielectric, can there said to be a complete understanding. In both the laser pulse propagation and the single-atom response there is significant uncertainty in both the physical theory and in the practical determination of experimental parameters, especially when moving to more exotic regimes such as few-cycle pulses or molecular targets. For example, in many regimes, HHG is exquisitely sen-

---

[1]The limited spatial dimensionality of current detectors must intrude here: whilst a three-dimensional XUV spectrometer may be technically feasible using techniques borrowed from integrated-field spectroscopy in astronomy, such devices have not yet been developed. A single image from the spectrometer therefore returns $I(\omega, y)$, the field intensity at a distance y along a direction parallel to the entrance slit. The beam is either sampled or averaged in the direction perpendicular to the entrance slit. Three-dimensionality may be achieved by scanning the position of the entrance slit, but usually one renders this unnecessary by assuming cylindrical symmetry.

sitive to the spatio-temporal profile of the laser pulse, which is currently challenging to control and characterize. At higher powers, nonlinear propagation effects such as the ionization-induced plasma and possibly the Kerr effect play a role. However, significant uncertainty exists in both ionization rates and nonlinear scattering cross-sections.

Much remains to be understood of the single-atom response itself. Within the strong field approximation (SFA), the properties of the target that affect the result are the ionization rate and recombination dipole element. Even in the simplest possible system, a hydrogen atom, there have been recent theoretical developments in this regard [415, 416], and for molecules the topic is under active research [417]. Extensions of the SFA towards multi-electron systems, molecular systems, and more accurate incorporation of the ionic potential are also currently undergoing intensive research, and have been the subject of controversy [356].

The emerging role of HHG as a next-generation light source means that resolving the afore-mentioned uncertainties is of interest beyond strong-field physics. Section 2.1.2 outlined the symbiotic development of ultrafast sources and ultrafast metrology; there are no indications this trend will cease. For example, in applications of attosecond pulses, tight focusing may be desirable in order to obtain high spatial resolution and also to increase the peak intensity so as to access non-linearities, which have been limited thus far [398, 399]. A tight spatio-temporal focus requires the focal planes of the harmonics to coincide; however as the experiments and calculations of this chapter show, this will not generally be the case. Characterization and ideally control will be required to overcome this problem.

Ultraviolet and soft x-ray pulses from high harmonic generation have been proposed as a seed for the next-generation of ultrashort free-electron lasers [418, 419], although the term free-electron amplifiers is more apt in this context. Since the amplified pulse inherits the phase of the seed, and harmonics are generated in the amplifier itself due to nonlinearities, it is likely that wavefront characterization of the seed will assist in controlling and optimizing the output.
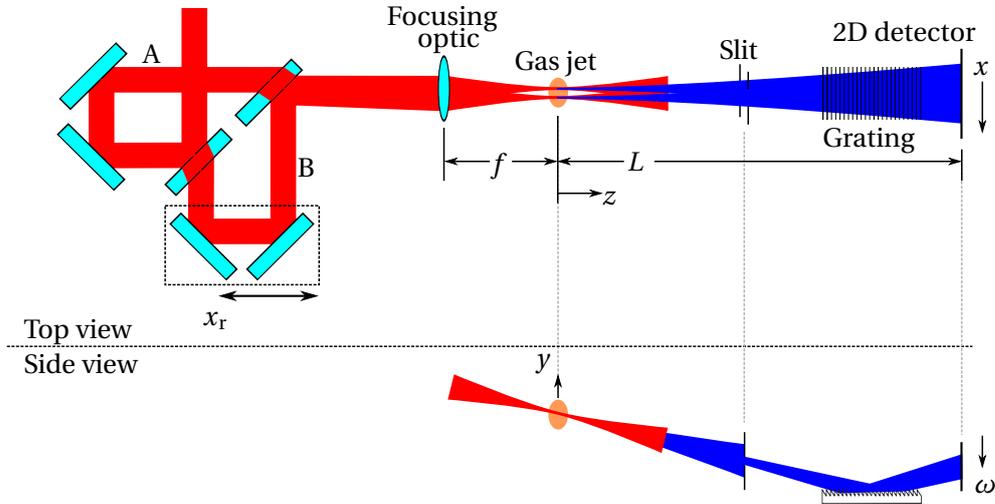
Figure 7.1: Setup for lateral shearing interferometry of high-harmonic generation; the focusing optic is drawn as a lens for clarity but a concave mirror may also be used.

7.2 Theory of the method

The technique is a type of LSI, and is depicted in Fig. 7.1. A Mach-Zehnder interferometer (MZI) produces two laser pulses with slightly different propagation directions in the horizontal (xz) plane. These are focused into the target, where they produce separate harmonic sources A and B. The harmonics propagate along the z-axis to a spatially resolving XUV spectrometer, which disperses the harmonics along the vertical (y) axis. In the xz plane, the beams undergo free propagation — no re-imaging occurs. Therefore, a two-source interference pattern is produced with spatial carrier $k_x = sk/L$ where s is the source separation in the focus and L the distance from the sources to the detector. The spatial carrier permits recovery of the phase difference between the fields using the standard Fourier-domain filtering algorithm [92], although a refinement that I describe below gives improved noise rejection.

We used multiple-shot LSI; this has practical advantages over a single-shot version. These are explained below. The retro-reflector of arm B is laterally displaced by an amount x_r to two or more positions, causing a corresponding displacement $x_d = 2x_r L/f$ on the detector. The mechanism is shown in Fig. 7.2. Interferograms patterns are recorded at each position. Arm A remains sta-
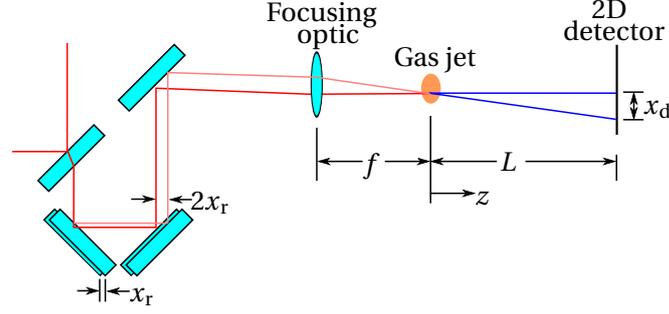
Figure 7.2: Mechanism for multiple-shot LSI. Two positions of the B arm are shown. The A arm remains fixed throughout the measurement and is omitted for clarity.

tionary and acts as a local oscillator throughout the measurement. Therefore, the phase of each interferogram is

$$\Phi(x, \omega; x_d) = \phi_B(x - x_d, L, \omega) + x_d kx/L - \phi_A(x, L, \omega). \tag{7.1}$$

where $\phi_A(x, z, \omega)$ and $\phi_B(x, z, \omega)$ are the phases of harmonic fields A and B, including that arising from propagation, and the middle term proportional to x describes the tilt which accompanies the lateral shear. The z origin is taken as the focal plane of the focusing optic. Upon subtraction of two phase differences from different shears $x_{d,1}$ and $x_{d,2}$, the phase of the A arm harmonics cancels out, leaving

$$\Gamma(x, \omega; X_d) = \Phi(x + x_{d,1}, \omega; x_{d,2}) - \Phi(x + x_{d,1}, \omega; x_{d,1}) \tag{7.2}$$

$$= \phi_B(x + X_d, L, \omega) - \phi_B(x, L, \omega) - \frac{X_d kx}{L}. \tag{7.3}$$

The first two terms on the right-hand side of (7.3) constitute a spectrally resolved LSI measurement of the field B at $z = L$. The shear is $X_d = x_{d,1} - x_{d,2}$ using the sign convention of (2.39).

The third term in (7.3) is a consequence of the fact that the shear is produced by a tilt. Because this term varies linearly with position, it produces a quadratic contribution in the reconstructed phase which must be taken into account. In principle it may be subtracted and subtracted. However, an elegant interpretation may be inferred from its contribution to the reconstructed phase,
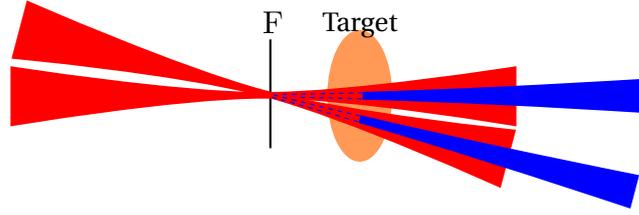
Figure 7.3: Effect on high-harmonic generation of a lateral displacement of the collimated laser beam, leading to a tilt of the focused beam, when the target is not in the focal plane. The beam is shown for two different shear stage positions. The dashed blue lines represent backpropagation of the generated harmonic field to the focal plane F.

found by integration:

$$\phi = \frac{1}{X_d} \int \frac{-X_d}{L} kx \, dx = -\frac{kx^2}{2L}. \tag{7.4}$$

This is equal and opposite to the phase associated with diffraction of a point source into a spherical wave (within the paraxial approximation) over a distance L. In the reconstruction, it cancels out the predominately quadratic phase of the diverging harmonics, leaving only the more interesting contribution that arises from deviations of the harmonics from an ideal point source at $z = 0$. Therefore, if one proceeds with reconstruction procedure using $\Gamma(x, \omega; X_d)$ as written in (7.3), the harmonic field that is reconstructed is that of arm B at $z = L$, minus the spherical phase resulting from propagation from $z = 0$ to $z = L$. This is true even if the gas jet is at $z > 0$, in which case the harmonics do not physically exist at $z = 0$. One must imagine the virtual sources at $z = 0$, found by back-propagating from the gas jet to $z = 0$. This is illustrated in Fig. 7.3

The primary rationale for performing this multiple-shot method, which is analagous to homodyne optical technique for SPIDER (HOT-SPIDER) (section 2.3.6.2), is that it does not require the A and B beams to be identical. The only requirement on the harmonics from arm A is that they span, both spatially and spectrally, those of arm B, and that they remain constant throughout the measurement apart from absolute phase fluctuations. An additional advantage is it enables the sequential acquisition of multiple shears.

Besides the aforementioned requirements on the A arm, the other formal requirements of the technique is that the generating medium be homogenous, and isotropic with respect to rotations about the y-axis; these conditions apply to the volume seen by the B arm over the full range of

shears. In the current set of experiments, the worst violation of these assumptions arose from the spatial wings of the A arm, which undergo coherent interference with the B arm, altering the HHG process.

Note that throughout this formalism, I have neglected terms constant with respect to x in the expressions for the phase because the experimental setup is not stable enough to detect them reliably. In particular, we found that the lateral displacement of the roof-hat in arm B resulted in random time delay, typically a few femtoseconds, and probably caused by stage translation errors. Because the resulting phase depends frequency but not position, this error does not affect quadratic and higher terms in the reconstructed spatial phase.

7.2.1 Far-field diffraction regime

A useful interpretation is revealed when the detector is distant enough from the source to be in the Fraunhofer diffraction regime. In that case, a stationary-phase approximation can be used to relate the field at the detector to the spatial Fourier transform of the (potentially virtual) field at $z = 0$:

$$E(x, y, L, \omega) \approx \frac{1}{i\lambda L} \exp\left[\frac{ik(x^2 + y^2)}{2L}\right] \tilde{E}\left(\frac{kx}{L}, \frac{ky}{L}, 0, \omega\right). \tag{7.5}$$

Here, $\tilde{E}(k_x, k_y, 0, \omega)$ denotes the spatial Fourier transform of the field at $z = 0$. If the slit is sufficiently small and located in the Fraunhofer diffraction regime, then it selects the $k_y = 0$ component. This is implied from here on. Under these conditions, if the beam is cylindrically symmetric then the zeroth order Hankel transform relates the near and far-fields. However, the symmetry is broken by the shear and so for the derivation I shall work with the spatial Fourier transform. Substituting (7.5) into (7.3), one obtains

$$\Gamma(x, \omega; X_d) = \tilde{\phi}_B\left(k\frac{x + X_d}{L}, 0, \omega\right) - \tilde{\phi}_B\left(k\frac{x}{L}, 0, \omega\right) \tag{7.6}$$

where $\tilde{\phi}$ denotes the phase of $\tilde{E}$. Note that neither the quadratic phase factor of (7.5) nor the linear term in (7.3) appear in (7.6) — they cancel each other exactly. Therefore in the Fraunhofer regime the experiment represents LSI in the spatial Fourier domain, with the shear being $K_x = kX_d/L$.

The reconstructed phase combined with the amplitude returns the spatial Fourier transform of the field at $z = 0$.

7.2.2 Linear multiple-shear algorithm

To enable consistency-checking of the results, we acquired interferograms at multiple shears. For the reconstruction, I could have used the multishear algorithm of chapter 3. However, it was not ideally suited to the data for several reasons. i) The intensity profiles were smooth with no intensity nulls. This meant that the phase could be unwrapped without risk of inconsistency between the shears. The unwrapping part of the preprocessing in the general multishear algorithm was therefore unnecessary. ii) In some instances, there were significant differences between the phase gradients obtained using different shears, well beyond that expected from shot-to-shot instability. This was attributed to interference between the laser focii, preventing them from acting as independent sources. As a workaround, we took many different shears, and treated the discrepancies as a random error, to be minimized using a fitting procedure. However, the multiple shear algorithm of chapter 3 is not well suited to this sort of error because it works on each shear in sequence, rather than performing a global optimization across them.

To alleviate these issues, I developed a simplified multiple shear algorithm which performs the entire reconstruction in a single linear least-squares minimization. The restriction, compared to the general algorithm, is that all the measured phase differences are assumed to be unwrapped with mutual consistency. I shall now provide a derivation of this algorithm.

For consistency with the current experiment I will work in the spatial domain; otherwise the notation is the same as the original multishear derivation in section 3.5.2. The sampled points are $x_n = x_0 + nX$ where the sampling rate X is the absolute value of the size of the smallest shear. The other shears $\{X_k\}$ are assumed to be sampled at integer multiples $\{C_k\}$ of the sampling rate i.e. $X_k = C_k X$. In discrete form, the phase difference measurements are

$$\Gamma_{k,n} = \phi_{n+C_k} - \phi_n + \eta_k \tag{7.7}$$

where η_k is the absolute phase of interferogram k. Equation set (7.7) may then be solved using standard linear least squares methods. In matrix notation, the unknown phase is the column vector $\boldsymbol{\phi} = [\phi_0, \phi_1, \ldots]^{\mathrm{T}}$, whilst the phase difference column vector for shear k is $\boldsymbol{\Gamma}_k = [\Gamma_{k,0}, \Gamma_{k,1}, \ldots \ldots]^{\mathrm{T}}$. The finite differences are represented by the matrices $\mathbf{G}_k$, defined explicitly in (3.20). One also introduces a diagonal weights matrix $\mathbf{W}_k$ with entries equal to the inverse of the standard deviation on the measured phase differences. For each shear, (7.7) is then written with weighting included as

$$\mathbf{W}_k \boldsymbol{\Gamma}_k = \mathbf{W}_k (\mathbf{G}_k \boldsymbol{\phi} + \eta_k). \tag{7.8}$$

where the scalar η_k is implicitly multipled by a column vector of ones. The shears are then combined by vertically stacking (7.8) for each k, so that the entire problem is summarized as

$$\mathbf{B}\mathbf{d} = \mathbf{F} \tag{7.9}$$

where

$$\mathbf{B} = \begin{bmatrix} \mathbf{W}_1(\mathbf{G}_1 + \mathbf{H}_1) \\ \mathbf{W}_2(\mathbf{G}_2 + \mathbf{H}_2) \\ \vdots \end{bmatrix}, \qquad \mathbf{d} = \begin{bmatrix} \boldsymbol{\phi} \\ \boldsymbol{\eta} \end{bmatrix}, \qquad \mathbf{F} = \begin{bmatrix} \mathbf{W}_1 \boldsymbol{\Gamma}_1 \\ \mathbf{W}_2 \boldsymbol{\Gamma}_2 \\ \vdots \end{bmatrix}. \tag{7.10}$$

I have introduced the vector of absolute phases $\boldsymbol{\eta} = [\eta_1, \eta_2, \ldots]^{\mathrm{T}}$, which must be included in the unknown vector although one is not interested in their values, and the matrix $\mathbf{H}_j$ which is all zeroes except for the j-th column:

$$\mathbf{H}_j = \begin{bmatrix} 0 & \ldots & 0 & 1 & 0 & \ldots \\ 0 & \ldots & 0 & 1 & 0 & \ldots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \ddots \end{bmatrix}. \tag{7.11}$$

The number of rows of $\mathbf{B}$ and $\mathbf{F}$ is equal to the total number of "observations"

$$P = KN - \sum_{k=1}^{K} C_j. \tag{7.12}$$

(The "missing terms" represented by the negative term in (7.12) are those phase differences at the boundaries of the sampled frequency interval.) Since the zeroth- and first-order phases are not defined, one sets (arbitrarily) $\phi_0 = 0$ and $\eta_1 = 0$. This is achieved by removing the corresponding columns from **B**. The number of unknowns is therefore $N' = N - 1 + K - 1$.

Equation (7.9) may be solved using standard linear least squares methods. The absolute phase terms couple together all of the unknowns, meaning that (7.9) is not banded unlike the generalized concatenation step (3.22) of the original multishear algorithm. It may therefore present some numerical challenges for large problems. However, for the reconstructions in this chapter, with typically 30 unknowns and 4 shears, the solution time was less than 100 ms on my laptop computer.

In the experiments of this chapter, consistency between the shears was a significant issue in some cases. Because the reconstruction is a simple matrix operation, standard methods are available for evaluating both the "goodness-of-fit" and the uncertainty in the retrieved parameters. Here I use the reduced chi-squared as a convenient quantification of the goodness-of-fit:

$$\chi^2_{\text{red}} = \frac{|\mathbf{Bd} - \mathbf{F}|^2}{\nu} \tag{7.13}$$

where $\nu = P - N'$ is degrees of freedom. If the errors have been estimated correctly, one expects $\chi^2_{\text{red}} \approx 1$. A higher value indicates that the data is inconsistent with the model or the errors have been under-estimated. In the present context, I estimated the errors from statistics of shot-to-shot fluctuations. Reduced chi-squared values significantly in excess of unity indicate inconsistency between the shears, presumably caused by violation of the underlying assumptions of the method, mentioned in section 7.2. For such reconstructions, the least-squares minimization is implicitly assuming that the errors resulting from the violation of the assumptions are independent and normally-distributed.

For $\chi^2_{\text{red}} < 1$, the uncertainty in the solution **d** is determined by the covariance matrix $\mathbf{V} = (\mathbf{B}^\mathrm{T}\mathbf{B})^{-1}$ which contains all second-order correlations between retrieved parameters. Here, only the unknown phase coefficients are of interest; one has $\Delta\phi_i\Delta\phi_j = (\mathbf{V})_{ij}$ where Δ denotes fluc-
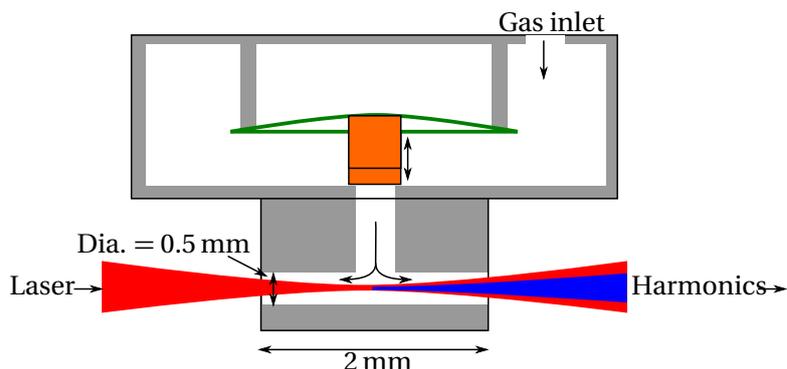
Figure 7.4: Pulsed gas jet target used in the longitudinal gas jet scan experiment.

tuations about the mean value. For example, $\sqrt{(\mathbf{V})_{ii}}$ is a meaningful size for an error bar on the inferred ϕ_i.

The definition of the covariance matrix in the previous paragraph describes fluctuations due to the shot-to-shot fluctuations. For $\chi^2_{\mathrm{red}} > 1$, the dominant errors are due to inconsistency between the shears. In the present experiment, I did not attempt to estimate empirically the magnitude of these errors. Instead, I used the standard technique of *assuming* that these errors were such that the fit was "good but not too good" i.e. $\chi^2_{\mathrm{red}} = 1$ [420]. Effectively, this means replacing the definition of the covariance matrix with $\mathbf{V} = \chi^2_{\mathrm{red}}(\mathbf{B}^{\mathrm{T}}\mathbf{B})^{-1}$.

7.3 Experiment: longitudinal gas jet scan

We constructed an experimental implementation at the attosecond beamline at Imperial College London [421, 422]. The laser source was a hollow core fibre compressor [60] pumped by a Ti:Sapphire chirped-pulse amplifier (CPA). For a convincing agreement with theory, we sought to scan a parameter of the high-harmonic generation process, taking LSI measurements at each point. An attractive choice would have been the peak laser intensity. However, we needed the full ouput of our laser system just to get any measurable signal at all. Therefore, we choose to scan the gas jet z-position. The gas jet target is shown in Fig. 7.4.

The procedure was the following: we scanned the gas jet z-position over a distance of 8 mm in 0.5 mm steps, centred around the position of maximum harmonic yield. At each gas jet z-position, we used 7 different shear stage positions. At each shear stage position, we took a several far-field

Table 7.1: Gas jet z-scan parameters.

Parameter	Value
Laser	
Laser centre wavelength	785 nm
Laser pulse duration	13.6 fs
Laser waist	50 μm
Laser peak intensity	1.4×10^{14} W/cm^2
Target	
Gas species	Krypton
Gas jet effective thickness	1.25 mm
Gas peak pressure	0.1 atm
Gas jet z-scan range	8 mm in 16 steps
Lateral shearing interferometer	
Mach-Zhender lateral shear range	0.6 mm
Focal spot separation	130 μm
Focal length	40 cm
Focii-to-detector distance	118 cm
Camera exposure time	5 ms

spectra: background (5 images), A arm (10 images), B arm (10 images), and AB interferogram (50 images). A servo-driven beam blocking system automated the acquisition of the different arms, whilst a linear actuator adjusted the shear stage position. We moved the gas jet z position manually. The entire run took 51 minutes.

The parameters are listed in Table 7.1. The broadened laser spectrum was measured using an Ocean Optics HR2000 spectrometer, and the centre wavelength was taken as the centroid of the spectrum. The pulse duration was taken by assuming a transform-limited pulse. The waist was taken from the focal spot measurements of the beams. The peak intensity was inferred from power measurements before the vacuum chamber. The gas jet thickness was taken as the tube length in the target. The gas target pressure was rather uncertain, with original estimations based on experience with different designs [421]. The middle of the gas jet (along the z-axis) was located to a precision of ± 2 mm with respect to the focus. The values of the target thickness, pressure and location with respect to the focus were refined using the the simulations described below.

Figure 7.5 shows the position and shape of the two laser focii for the different shears and gas jet positions. We took care to adjust the collimation of the incoming laser beam such that its waist coincided with the geometric focus of the focusing optic used for HHG. Our method was to
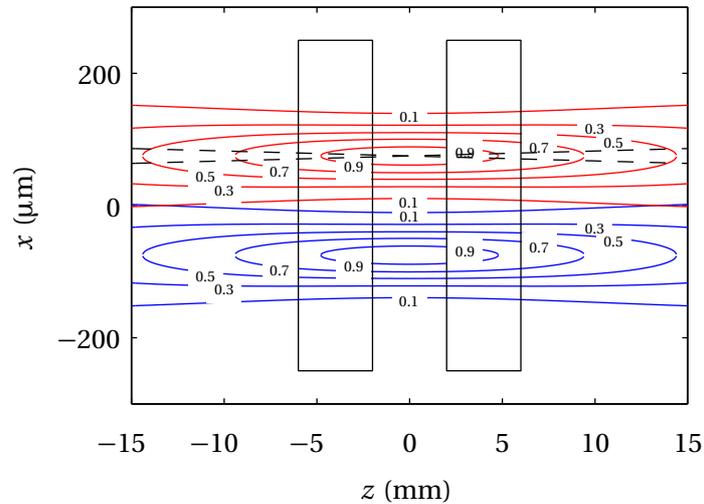
Figure 7.5: Shape of A (blue) and B (red) laser focii in the gas jet; different scales are used for the z and x directions but otherwise the drawing is to scale. Contour lines indicated the intensity. The extreme ends of the target scan range are indicated by the black rectangles. The central axes of beam B for the two extreme shear stage positions are indicated by the dashed black lines.

image the focal plane onto a camera and scan the shear stage in the MZI back and forth — in the geometric focus, the beam does not move.

7.4 Data processing

7.4.1 Extraction of phase differences

The phase of each interferogram was retrieved using a minor variation on the usual Fourier-domain filtering procedure. Figure 7.6(a) shows a typical raw interferogram. As usual with beams crossing at an angle, the fringe period is proportional to the wavelength, and this is particularly evident with the broadband harmonic spectra. To obtain the best noise rejection, I first performed a discrete Fourier transform (DFT) along the x-axis, yielding the transform shown in Fig. 7.6(b). I then applied the filter shown by the white lines Fig. 7.6(b). This isolated the sidebands and removed noise at high and low vertical spatial frequencies. Then, I applied a DFT along the y-axis, yielding the transform shown in Fig. 7.6(c). A second filter, indicated by the white lines, removed noise at high horizontal spatial frequencies. Finally, a 2D inverse DFT brought the data back to the spatial domain. Taking the complex argument yielded the fringe phase, shown in Fig. 7.6(d).
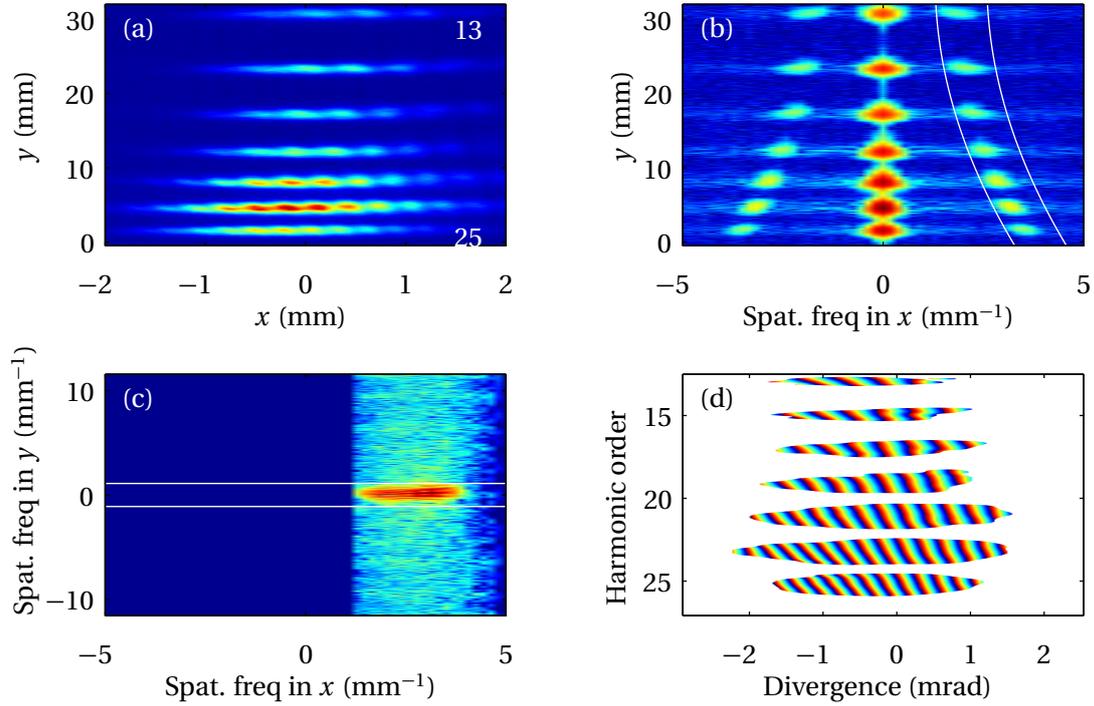
Figure 7.6: Extraction of the interferogram phase. (a) Raw data, linear colour scale. The 13th and 25th harmonics are indicated. (b) 1D-DFT along x-axis, and sideband filter (white lines). Colour scale is logarithmic in amplitude over a range of 10^3. (c) 2D-DFT, with noise filter (white lines). Colour scale is logarithmic in amplitude over a range of 10^3. (d) Phase of the resulting sideband. Linear colour scale over $[-\pi, \pi]$. The spectral and spatial calibration has been applied and the phase is only shown in regions with significant signal.

One shear, for which I defined $x_\text{d} = 0$, was taken as a common reference. The phase of this shear was subtracted from all the others to give the phase differences. Figure 7.7(a) shows a typical phase difference between two interferograms. The tilt on the fringes shows that they have a spectral component, possibly caused by a time delay resulting from a stage translation error. I assumed the time delay to be completely random, and did not try to infer any information from it. The result is that the measurement did not retrieve the linear spatial phase component in the far field, or equivalently the source transverse position in the near field. This is analogous to the random absolute phase in a spectral phase interferometry for direct electric-field reconstruction (SPIDER) interferogram preventing the technique from retrieving the linear spectral phase, or equivalently the arrival time of a pulse. Instead, I set the absolute phase of every frequency, defined by an
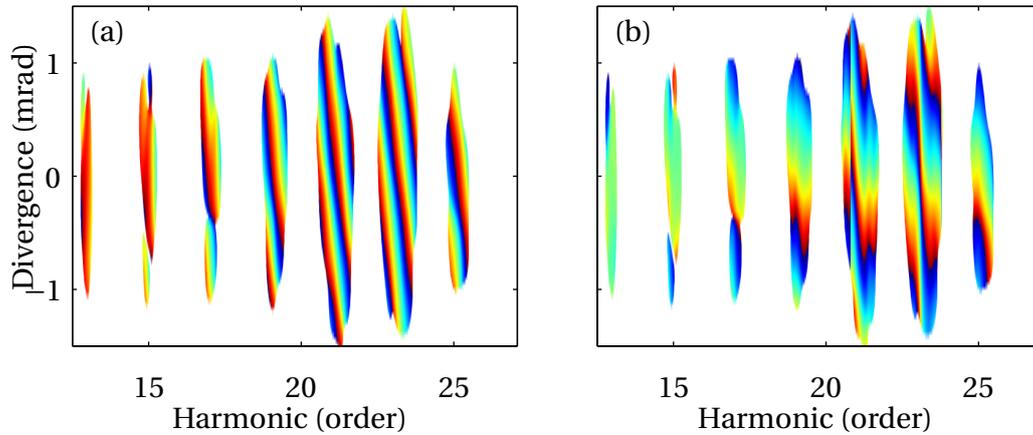
Figure 7.7: Typical phase difference between two interferograms. (a) As obtained from interferograms. (b) After zeroing absolute phase at every frequency to zero.

intensity-weighted average, to be zero, effectively placing the centroid of all sources at $x = 0$. This causes the fringe phase to have predominately spatial variations, as shown in Figure 7.7(b).

As Fig. 7.7(b) suggests, the resulting phase differences were fairly homogenous across each individual harmonic. (The irregularities of the 21st and 23rd harmonics were not reproduced from shot-to-shot.) To simplify the analysis and improve the signal-to-noise ratio (SNR), I averaged the phase across each harmonic, again using an intensity weighting. Figure 7.8 shows a typical set of results for a subset of 7 out the 50 shots taken at each combination of shear and gas jet position. Significant shot-to-shot fluctuation is evident. I performed shot averaging on the complex-valued interferometric products, and then took the phase of the result. This produces better results than taking the phases, then averaging.

I also estimated the uncertainty on the shot-averaged phase differences. Without some knowledge of their uncertainties, it would have impossible to quantify the consistency between the different shears. I decomposed the fluctuations about the mean into the in-phase and quadrature components. Only the latter cause phase fluctuations. Examination of the quadrature fluctuations showed little shot-to-shot correlation, so I assumed that their average converged with the square root of the number of samples. The standard error of the quadrature fluctuations was then converted into phase noise using the methods of appendix A. The end result of this process was
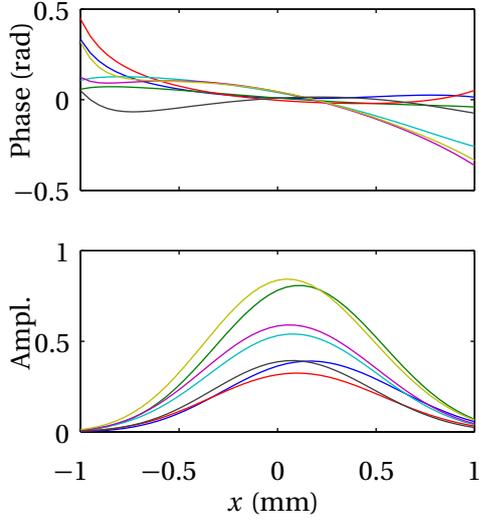
Figure 7.8: Phase (upper plot) and amplitude (lower plot) of interferometric products for 7 out of the 50 shots taken at one specific gas jet z and shear position.
.
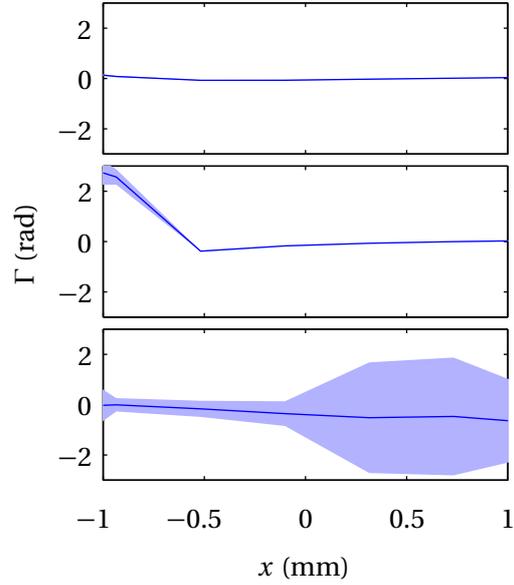
Figure 7.9: Three typical shot-averaged phase-difference measurements and their uncertainties, each at a different shear and harmonic number.

$\Gamma(x,q)$ and its uncertainty $\sigma_\Gamma(x,q)$. A sample of these is shown in Fig. 7.9. In some cases, the uncertainty (compared to the size of the phase-difference signal) after shot-averaging was acceptably small, whilst in others, principally at the higher harmonics towards the extreme ends of the gas jet scan range, the uncertainty was large.

7.4.2 Phase reconstruction

I reconstructed the fields using the algorithm of section 7.2.2. Four shears of size $X_d = \pm0.42, \pm0.83$ mm were used. The minimum shear size was dictated by the spatial correlation length of the k_x-filter used in the Fourier processing procedure — there was little point in using a shear that is smaller than this size because structure on a finer scale is removed by the filter. Furthermore, the filtered noise contributions would become correlated, complicating the error analysis. The maximum shear size was dictated by the apertures of the beamline.

Most measurements had reduced-χ^2 values between 50 and 150, indicating very poor consistency between shears if the shot-to-shot fluctuations are taken as the dominant measurement

error. Only for the highest harmonics at the extremes of the gas jet scan, where the signal was extremely weak, did shot-to-shot fluctuations dominate the error. Therefore, in estimating the uncertainty of the reconstructed phase, it was necessary to include χ^2_{red} in the variance as described in section 7.2.2 for most calculations.

7.4.3 Extraction of intensity profiles

Far-field spatial profiles were obtained by subtracting the background image from the B arm images, which were taken for every shear stage setting and gas jet z-position. A typical raw image is shown in Fig. 7.10. I averaged the profiles across each harmonic, and then averaged across the 10 shots taken at each gas jet position and shear combination. I also moved each image by the opposite of its displacement x_{d} to undo the effect of the shear. A sample set of results, taken at different shears but at the same gas jet z position, are shown in Fig. 7.11.

As with the phase reconstructions, there were significant shear-to-shear variations. However, unlike in the phase reconstructions, these cannot be attributed to the presence of the A beam, since it was blocked for the intensity measurements. Instead, I attribute these fluctuations to drifts in the laser power on the approximately 30 s timescale between shear measurements. Similar fluctuations were observed in the harmonics from A arm, which was also recorded at every shear position but remained stationary throughout the measurement. For consistency with the phase measurements, I averaged the intensity measurements across the different shears to produce a "grand average" for each harmonic at each gas jet position. The standard error (standard deviation divided by square root of number of samples) of the shear average was taken as the uncertainty.

7.5 Comparison of results with simulation

Figure 7.12 presents the entire set of amplitude and reconstructed phase profiles.

To interpret these results, I shall first show that they largely agree with the results of a simulation. With the help of the simulation, a physical interpretation will then be given. The simulation details are described in Appendix B. Because the spectral response of the spectrometer was unknown, I applied a scaling factor to each harmonic to obtain best agreement with theory. It must
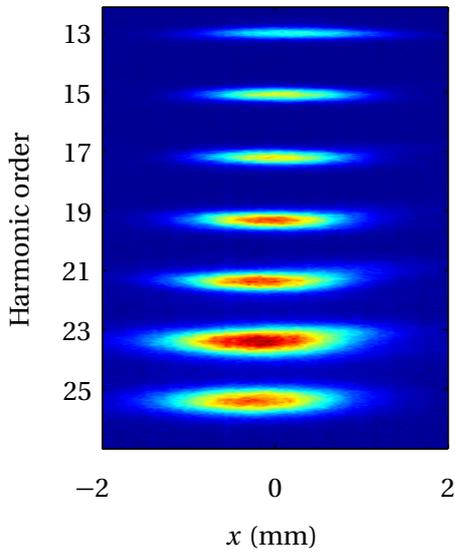
Figure 7.10: A typical raw intensity image.
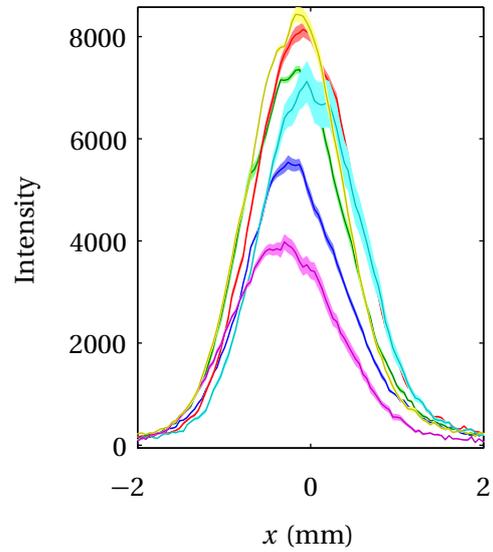
Figure 7.11: Shot-averaged spatial intensity profiles at different shears at a single gas jet z position. The uncertainty, taken as the standard error, is indicated by the shaded error.

be emphasized that this factor did not vary with the gas jet z position, meaning that any agreement between the measured and theoretical values in the z-dependence of the intensity is genuine.

Figure 7.13 shows the amplitude and phase profiles for harmonics 13, 19, and 25 at gas jet positions $-4\,$mm, $-1\,$mm and $2\,$mm. For harmonic 25, the agreement is good: the beam profile is close to Gaussian and experiment and theory agree in terms of the peak intensity, beam width, and phase curvature. For harmonic 19, the phase agreement is poor at $-4\,$mm but reasonable at the other values, whilst the intensity agrees reasonably. For harmonic 13, there is some disagreement in the beam width, but overall the measured intensity profiles are in reasonable agreement with the simulation. For the phase, the agreement is once again poor at $-4\,$mm. At $-1\,$mm, the simulated phase differs greatly from quadratic, being largely flat across the majority of the beam and then flipping by $\approx \pi$ in the wings. The flat central phase is captured by the measurements, along with a phase flip on one side. Finally at $2\,$mm, the agreement is good for the phase.

These observations are consistent with three general trends in the agreement between theory
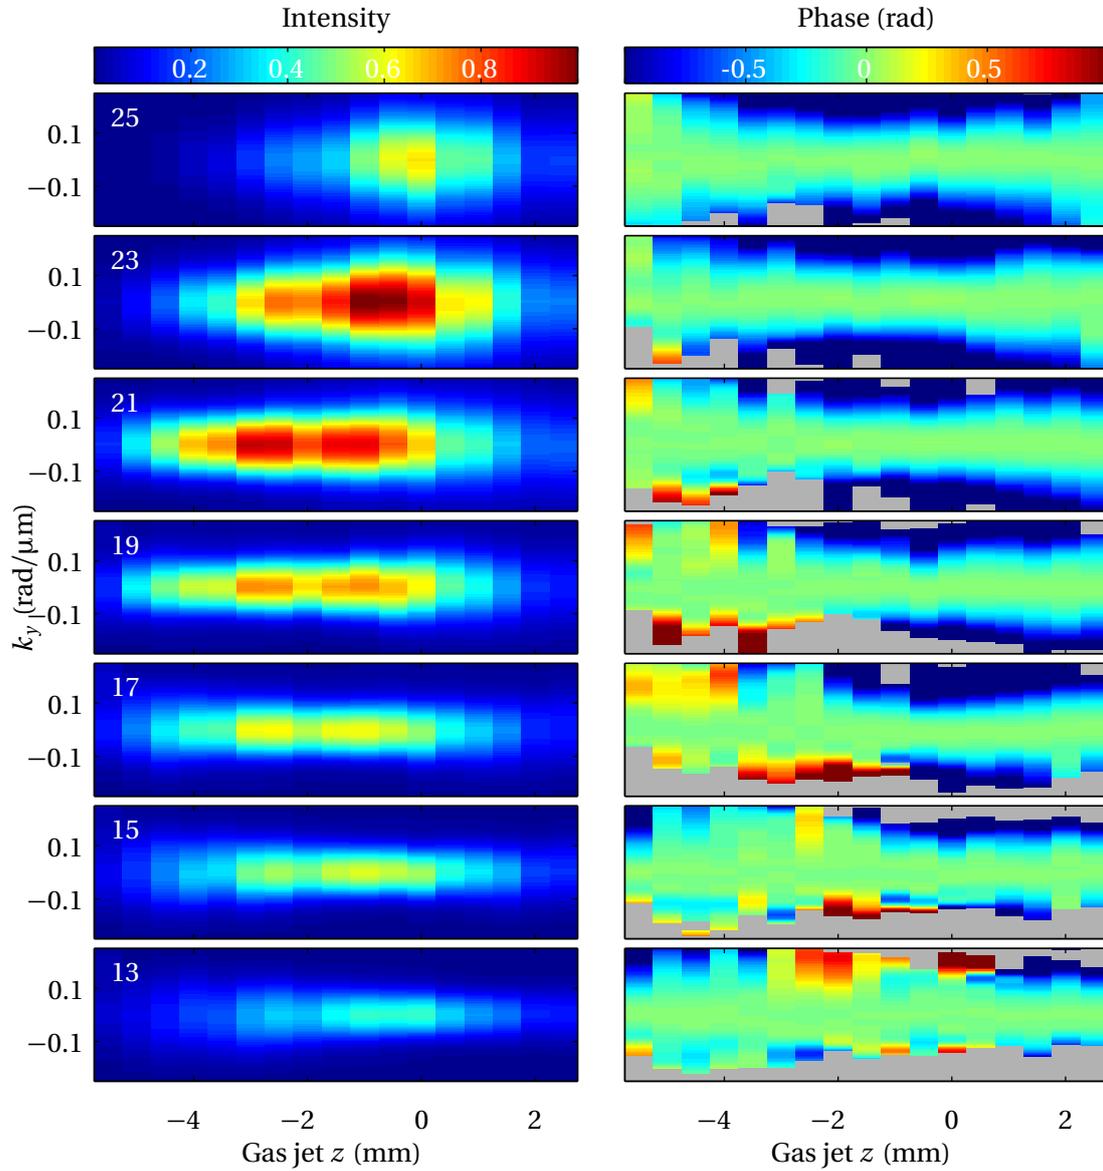
Figure 7.12: Measured far-field intensity (left) column and phase (right) for harmonics 13–25 across the entire z-scan range. The intensity is as recorded — no scaling has been applied. The phase has been greyed out where its uncertainty exceeds 0.4 rad.

and experiment evident from an examination of the entire data set. First, the higher harmonics tend to agree better. Second, the agreement tends to deteriorate towards negative z values i.e. as the gas jet moves upstream of the focus. Third, in cases where the simulated phase is flat in the centre but changes sharply in the wings, the experiments capture the flat middle but not, in general, the sharp changes.

7.6 Physical interpretation

Because the simulation and experiment are largely in agreement, the former can be used to disect the physics leading to the observed spatial profiles. Section 7.6.1 describes the spatial distribution of the single-atom response leading to the transverse macroscopic behaviour — that is, the profiles that would be observed if the target were infinitely thin. Then, section 7.6.2 summarises how phase-matching and other longitudinal macroscopic effects modify the profiles.

7.6.1 Single-atom response and transverse macroscopic effects

Figure 7.14 shows the laser intensity profile at the focus that is used in the simulations. Harmonics 13–23 are below the classical cutoff at the peak intensity. The first step in calculating the spatial profiles is computing the single-atom response at each point in the focus. Since the calculation ignores all spatial variation in the laser field except for peak intensity $I_L(r, z)$ and phase $\phi_L(r, z)$, the single-atom response may be completely described by its intensity dependence. This is shown in Fig. 7.15 for harmonics 13, 19 and 25.

Figure 7.15 shows the amplitude and phase of the harmonics over the range of laser intensities in the focus. The usual linear dependence on intensity is observed for the long and short trajectories below cutoff and the dominant (long in this case) trajectory above cutoff. The sum, given by the uniform approximation, exhibits constructive and destructive interference, with rapid phase jumps accompanying the latter when the field goes through a node. The spatial distribution of the single-atom response is shown in Fig. 7.16. In harmonics 13 and 19, quantum-path interference produces an annular amplitude distribution. Phase discontinuities accompany the points of destructive interference. The phase has been set to zero on axis so that only radial variations are
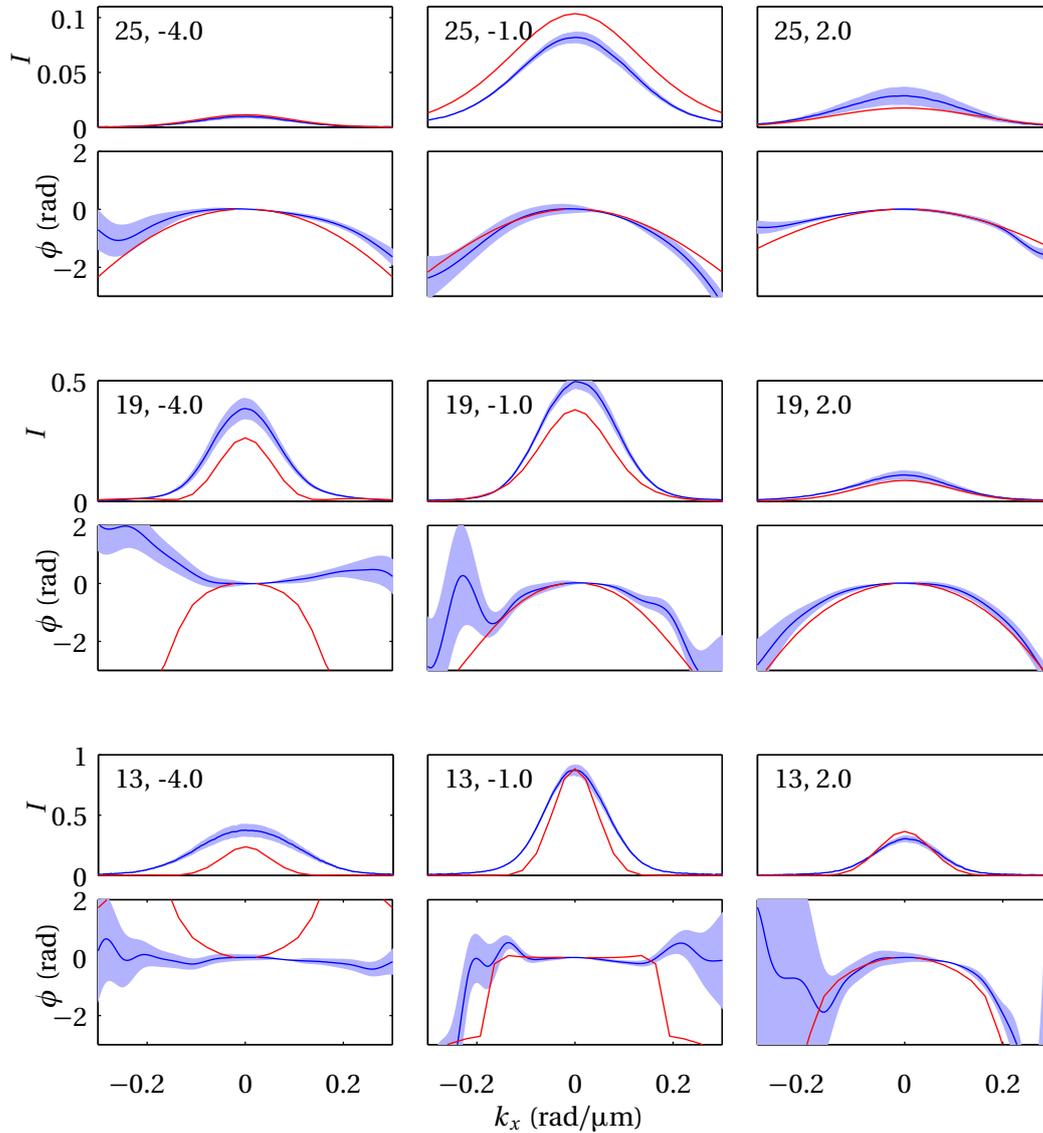
Figure 7.13: Measured (blue) and simulated (red) spatial profiles, for harmonics 13, 19 and 25, and gas jet z positions $-4\,$mm, $-1\,$mm and $2\,$mm. Each combination of harmonic number and gas jet z position is depicted in a pair of plots. In each pair, the upper plot is the intensity and the lower plot is the phase. The experimental uncertainty is indicated by the shaded blue area.
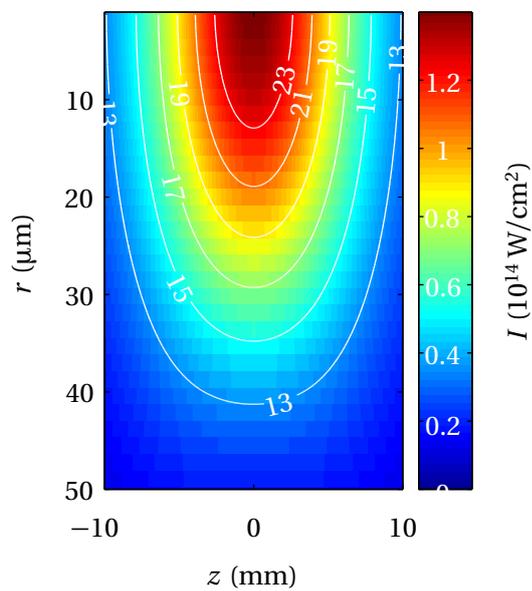
Figure 7.14: Laser intensity in the focus for the simulation. The contour lines show the classical cutoffs of harmonics 13–23.
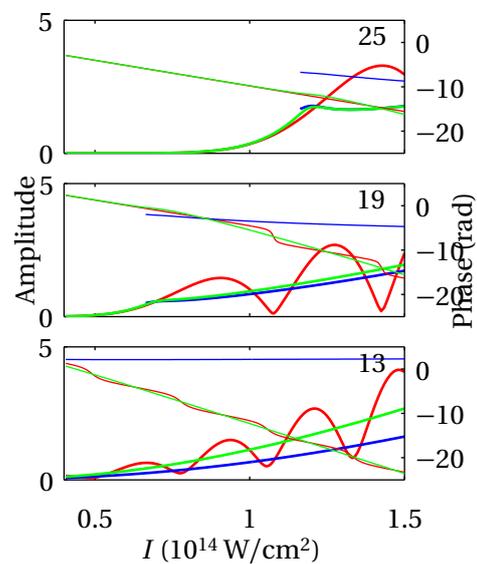
Figure 7.15: Theoretical amplitude (thick lines) and phase (thin lines) of short trajectories (blue), long trajectories (green) and the uniform approximation to their sum (red). The 13th, 19th and 25th harmonics are shown.
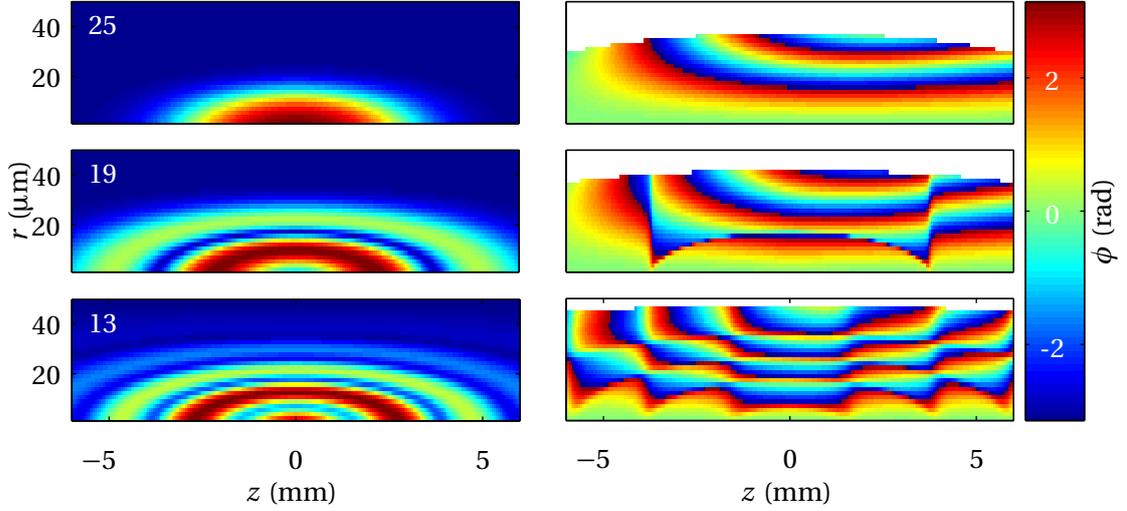
Figure 7.16: Simulated spatial profile of the single-atom response for a single half-cycle for harmonics 13, 19 and 25. The amplitude (left column) is normalized and on a linear scale. The phase (right column) has been set to zero on axis and is wrapped to $[-\pi, \pi]$, with colour scale shown. It is set to white where the intensity is so low that numerical artifacts dominate the phase.

shown. Across most of the z range, the phase increases off-axis due to the intensity-dependence of the single-atom response. However, it is not symmetric about $z = 0$, because of the phase $q\phi_L$ imparted by the phase-fronts of the laser. For $z < 0$, the laser wavefront is converging and opposes the intensity-dependent effect, and for $z > 0$ the opposite holds true.

The far-field profile of the single-atom response, obtained using the Hankel transform, is shown in Fig. 7.17. The phase includes the diffraction term $-\frac{k_T^2}{2qk_L}(0-z)$, accounting for propagation of the harmonics from their generation at z to the measurement plane at $z = 0$. In harmonic 25, the intensity profile is narrower for $z < 0$ because in the near-field distribution the intensity-dependent dipole phase is approximately cancelled by the converging laser phase. The opposite occurs for $z > 0$. The phase of harmonic 25 decreases away from the axis, and the curvature changes with z. This trend will be examined in detail below. In harmonics 13 and 19, quantum-path interference effects are observed, accompanied by phase discontinuities. The profile has a strong on-axis component, corresponding to the short trajectories, and a broad but weak off-axis component produced by the long trajectories, the difference in divergence being caused by the different intensity dependence of the phases.
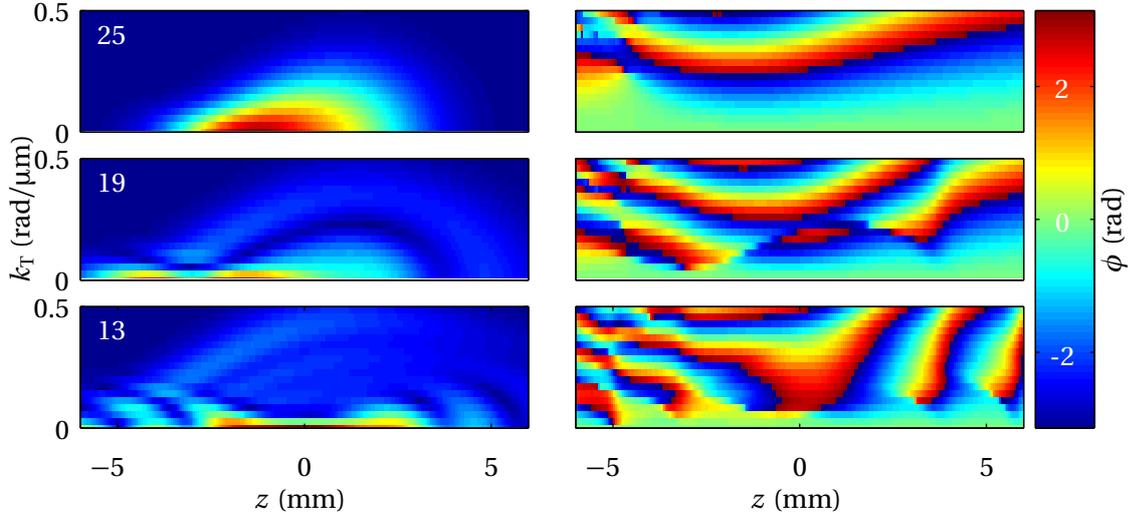
Figure 7.17: Simulated far-field profile of the single-atom response from a single half-cycle for harmonics 13, 19 and 25. The amplitude (left column) is normalized and on a linear scale. The phase (right column) has been set to zero on axis and is wrapped to $[-\pi, \pi]$, with colour scale shown.

The previous plots have considered only the strongest half-cycle of the pulse. I now bring in the other half-cycles. Figure 7.18 shows the far-field spatial profiles with all the half-cycles included; otherwise it is equivalent to Figure 7.17. The main effects of the other half cycles are a smoothing of the fine structure and reduction in the relative amplitude of the off-axis components.

7.6.2 Longitudinal macroscopic effects

The previous section examined the far-field profiles expected from the single-atom responses at individual transverse planes, as would be observed if the gas jet were infinitely thin. This section introduces the gas jet profile and discusses the influence of longitudinal macroscopic effects. In fact, I will show that because of the large absorption and dispersion of the gas, longitudinal macroscopic effects do not substantially alter the response.

The gas jet profile is shown in Fig. 7.19(a). Absorption plays a significant role. For example, a component of the 13th harmonic generated at the start of the gas jet is attenuated by a factor of 25 (in amplitude) during its propagation through the gas jet. The interplay between absorption and the rate of harmonic generation, which is proportional to the gas density, means that harmonics are generated most effectively in the second half of the gas jet. This *gas jet response*, the product
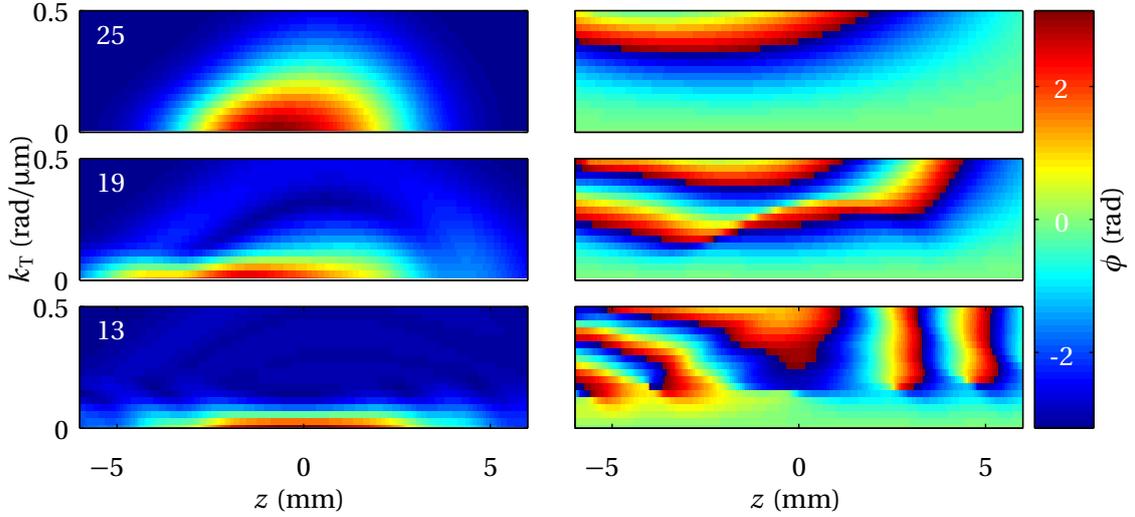
Figure 7.18: Simulated far-field profile of the single-atom response including every half-cycle for harmonics 13, 19 and 25. The amplitude (left column) is normalized and on a linear scale. The phase (right column) has been set to zero on axis and is wrapped to $[-\pi, \pi]$, with colour scale shown.

of the local pressure and the transmission to the end of the jet, is shown in Fig. 7.19(b). Its peak moves upwards and to the centre of the gas jet with increasing harmonic number due to the increasing transmission of krypton. Overall, absorption has a narrowing effect on the gas jet profile, particularly for the lower harmonics.

The components of the longitudinal phase mismatch of harmonics 13, 19 and 25 are shown in Fig. 7.20. The contributions from the gas are shown at the peak pressure at the centre of the jet. For all three, dispersion of the gas at both the laser and harmonic wavelengths combine to produce a net positive phase mismatch which dominates the other components. Its smallest value, $k_z \approx 5\,\text{rad/mm}$, occurs at $z \approx -3\,\text{mm}$, where the intensity dependence causes a negative contribution, particularly for the long trajectories. The corresponding coherence length is $\pi/k_z \approx 0.6\,\text{mm}$, approximately half the effective gas jet length. This shows that the phase matching causes an $\approx 50\%$ attenuation of the overall amplitude. However, for predicting the spatial profiles, what is most relevant are the relative efficiencies of the long and short trajectories. For $z < 0$, one expects the long trajectories to be favoured, and vice versa for $z > 0$. However, this effect is only moderate, with the coherent lengths differing by no more than a factor or two.
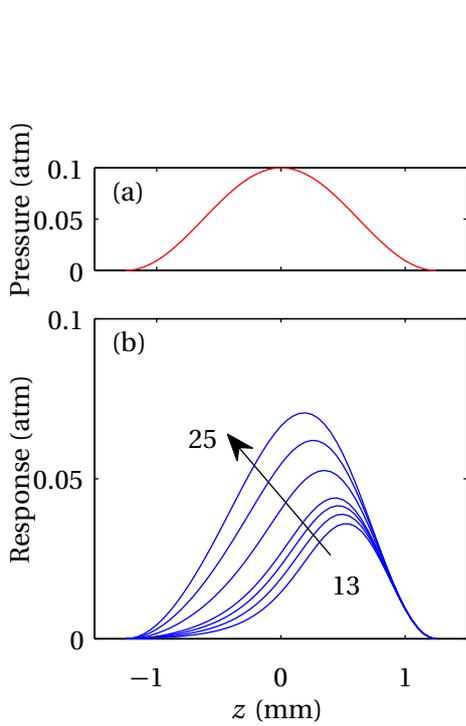
Figure 7.19: (a) Gas pressure. (b) Amplitude of gas jet response (product of gas pressure and transmission to $z = \infty$) for harmonics 13–25.
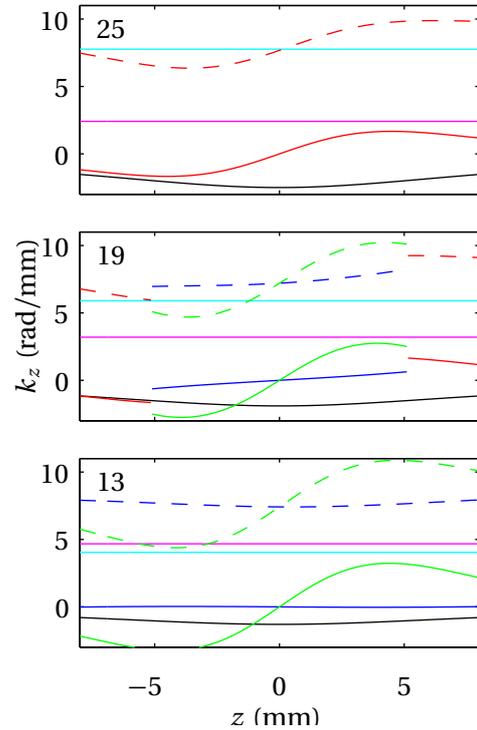
Figure 7.20: Components of the longitudinal phase mismatch from the refractive index at the laser wavelength (cyan solid), refractive index of the harmonics (magenta solid), Gouy phase (black solid), intensity-dependence of short and long (blue and green solid) trajectories below cutoff, and the dominant trajectory (red solid) above cutoff. The total phase mismatch is also shown (blue, green and red dashed).

One may now interpret the simulated far-field macroscopic profiles which were used for the comparison in Fig. 7.12. These are shown in Fig. 7.21. They are in fact very similar to the single-atom response (Fig. 7.18), the main differences being a small increase (decrease) in the apparent amplitude of the off-axis components for $z < 0$ ($z > 0$) due to the phase matching, and a movement of 1–2 mm of the entire response to negative z. The latter effect is caused by the shifted peaks of the gas jet response, shown in Fig. 7.19, and is greater for the lower harmonics.
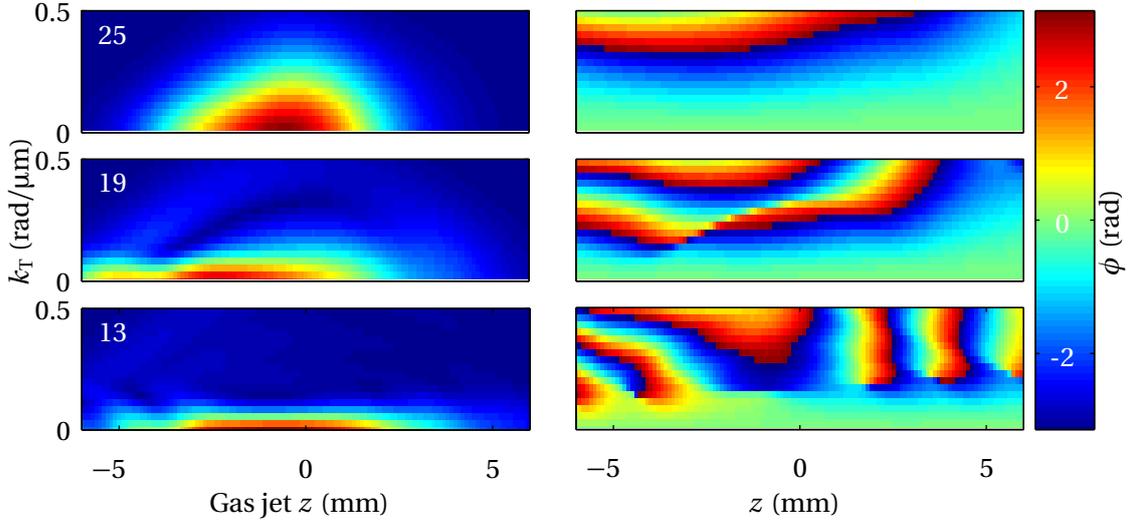
Figure 7.21: Simulated far-field profiles of the full macroscopic response for harmonics 13, 19 and 25. The amplitude (left column) is normalized and on a linear scale. The phase (right column) has been set to zero on axis and is wrapped to $[-\pi, \pi]$, with colour scale shown.

7.6.3   Trends in the gas jet z scan

Having deconstructed the physical processes leading to the macroscopically observed spatial profiles, I shall now discuss some trends observed in both the measurements and the simulations.

Harmonic 25 is perhaps the simplest to understand because the intensity is always below the cutoff so there is no significant splitting into the long and short trajectories. There is no quantum path interference and the phase of the single-atom response is always linear with the intensity. The spatial profiles are Gaussian-like, and may be characterized by their on-axis amplitude and phase curvature

$$\sigma_{k_{\mathrm{T}}}^{-2} = -\frac{1}{E(0)} \left.\frac{\partial^2 E(k_{\mathrm{T}})}{\partial k_{\mathrm{T}}^2}\right|_{k_{\mathrm{T}}=0}, \qquad \phi_{2,k_{\mathrm{T}}} = \left.\frac{\partial^2 \phi(k_{\mathrm{T}})}{\partial k_{\mathrm{T}}^2}\right|_{k_{\mathrm{T}}=0}. \tag{7.14}$$

Applied to a chirped Gaussian profile, $\sigma_{k_{\mathrm{T}}}$ is the width and $\phi_{2,k_{\mathrm{T}}}$ is the chirp. Corresponding definitions σ_r^{-2} and $\phi_{2,r}$ may be written in the near field. The far-field curvatures extracted from simulation and experiment are plotted in Fig. 7.22. The far-field width is smaller for negative z, because of the cancellation of the intensity-dependent dipole phase by the converging laser wave-front in the near field. For positive z, the near field acquires a strongly divergent wavefront, producing a

181
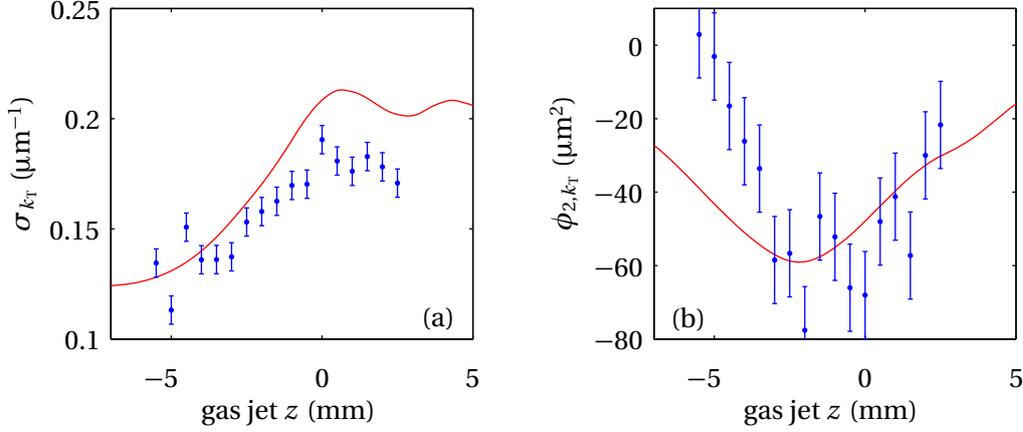
Figure 7.22: Properties of the 25[th] harmonic extracted from simulation (red line) and experiment (blue); (a) far-field width; (b) on-axis phase curvature.

broad angular spectrum. In the phase, several effects are at play. For $z \approx -5$, the near cancellation of the intensity-dependent phase by the laser wavefront produces a near-zero phase curvature in both the near- and far-fields. One possible cause of the disagreement between theory and experiment in this region is deviations of the laser beam profile from Gaussian. With increasing z, the near-field becomes more curved as the laser curvature decreases, and this initially (-5 – -2 mm) leads to curvature of the far-field. However, the relationship between the on-axis curvatures in the near- and far-fields is nonmonotonic. Rather for a Gaussian profile,

$$\phi_{2,k_{\mathrm{T}}} = \frac{-\phi_{2,r}}{\sigma_{2,r}^{-4} + \phi_{2,r}^2} \tag{7.15}$$

i.e. the absolute value of the far-field curvature initially increases with increasing near-field curvature, but then reaches a maximum before asymptotically decreasing to zero. This is observed for $z > -2$ in Fig. 7.22(b). An additional effect is the propagation from the generation plane to the measurement plane which produces a phase curvature $-(0-z)/(2qk_{\mathrm{L}})$ i.e. a linear increase with z.

Another trend evident in both simulation and theory is the dependence of the amplitude and phase profiles on the harmonic number. This is presented in Fig. 7.23 at $z = -1.5\,\mathrm{mm}$. Generally, the amplitudes become broader with increasing harmonic number. This can be attributed to the
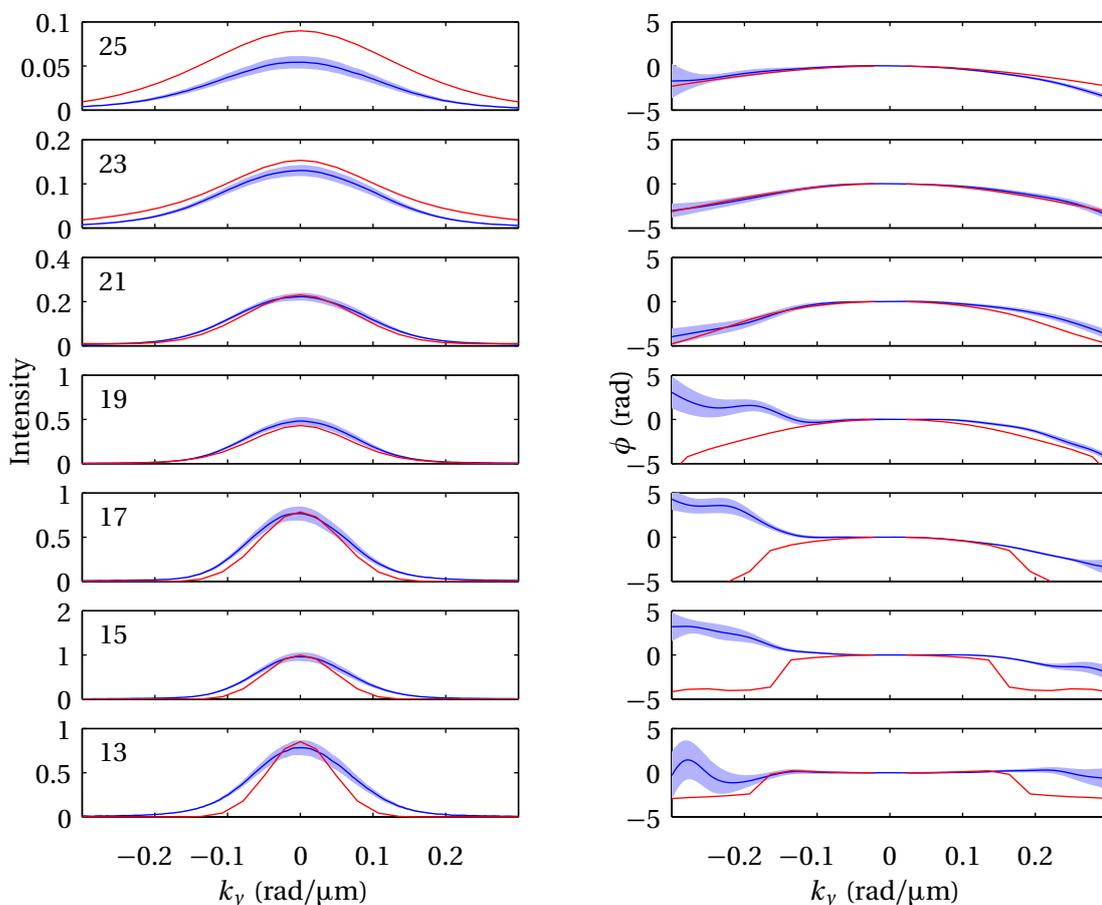
Figure 7.23: Measured (blue) and simulated (red) amplitude (left column) and phase (right column) profiles for harmonics 13–25 at gas jet $z = -1.5\,\text{mm}$. The uncertainty is indicated by the shaded blue area.

increasing intensity-dependent dipole phase of the short trajectories, which dominate the measurements close to the axis. This trend is also evident in the phase profiles: at low harmonics, the phase is quite flat because the short trajectories have a very weak intensity-dependence. Far into the wings, sharp phase discontinuities appear due to interference with the long trajectories. This lies below the sensitivity of the measurements.

7.7   Summary and outlook

Spectrally resolved wavefront characterisation is beneficial to both the science and applications of HHG. This chapter described one such method based on LSI: two harmonic sources are produced

and interfered on a spatially resolved spectrometer. One of the sources is rotated, producing a lateral displacement on the spectrometer. The lateral shearing interferogram so recorded encodes the spatial phase of the field at the detector, minus the spherical wavefront introduced by propagation to this plane from the source. I presented a comprehensive set of measurements of the spatial phase and amplitude profiles of harmonics 13–25 in krypton as a function of the gas jet z position. Over most of the parameter range, the results agreed with a simulation based on the quantum-orbits model. Various single-atom and macroscopic effects play a role in producing the observed profiles.

This work raises many avenues for future research. Constrained by the available laser power, the experiments were limited to harmonics 13–25 in krypton, with the entire system being optimized just to get any signal at all. Dispersion and absorption dominated the macroscopic effects, preventing direct observation of the interplay of the single-atom response and the geometry. Novel effects arising from nonlinear propagation of the laser were also not exhibited at these power levels. A higher power source would provide for more options in terms of the target gas species, intensity scan range and focusing geometry. The other noble gases also have lower absorption and dispersion, and a greater range of harmonics would be produced.

8 Quantum-path interferometry in high-harmonic generation

This chapter proposes quantum-path interferometry (QPI) mediated by a weak *control field*, a new technique to distinguish the contributions of different quantum paths in high-harmonic generation (HHG). Specifically, it measures the relative amplitudes and phases of the individual terms in the sum over paths that results from the stationary-phase approximation (SPA) to the single-atom response. The role of the weak control field is to cause a phase shift on each of the quantum paths, leaving them otherwise unaltered. I use the strong field approximation (SFA) to predict the value of this phase shift which, crucially, can be made different for each of the quantum paths. This allows their coherent sum to be decomposed into its individual summands using methods of carrier-encoded interferometry. The technique is novel in that it could potentially be a robust method of *experimental quantum-path analysis*, extending established techniques like spectrographic analysis of simulated HHG. This chapter develops the theory and presents a numerical example.

8.1 Motivations

The quantum orbits picture is a powerful and intuitive description of HHG and other strong field processes such as optical field ionization. It may be derived by applying the SPA to the multi-dimensional integrals that in turn arise by applying the SFA to the time-dependent Schrödinger equation (TDSE). The integrals are reduced to sums, and often only a few summands are required for an accurate description of the process. Furthermore, the summands usually have clear classical interpretations. Much of the phenomenology of HHG — the time-frequency structure, macroscopic phase-matching and attosecond pulse generation — can be understood using the quantum paths model.

Despite this, the quantum orbits model is an incomplete description of HHG. Given the temporal profile of the laser field, the quantum orbits picture predicts the birth times, momenta, and recombination times of the electrons, and hence the time-frequency distribution of the emitted

radiation. However, the probability amplitudes for the birth and recombination are not described by the SFA. Instead, they are given by the atomic and/or molecular structure of the target species, which may potentially be in a highly perturbed and dynamic state. When HHG is viewed as a probe of atomic and/or molecular structure and dynamics (as opposed to as a source of radiation for other purposes), these amplitudes are the very information that one is seeking. In this context, the quantum orbits model is a means of connecting the (potentially rather complex) features of the emitted radiation to the atomic/molecular physics of the target species. However, when multiple quantum orbits are present, their contributions must be disentangled. Furthermore, it may be useful to acquire amplitudes from multiple quantum orbits simultaneously.

A specific example is the ultrafast dissociation of hydrogen discussed in section 6.1.2. Here, macroscopic phase matching was used to isolate the short trajectory, whilst a repetition of the measurement with deuterium provided a calibration of the returning electron wavepacket. If one could detect the short and long trajectories simultaneously, and possibly even higher-order trajectories, it would extend the time range of the experiment and possibly offer a self-referenced means of characterizing the returning electron wavepacket without using a deuterated substitute. This is the first motivation for experimental quantum path analysis.

Despite its power, in many ways the SFA, upon which the standard quantum orbits model is based, is a gross approximation. It ignores the effect of the laser on the ground state and the effect of the atomic/molecular potential on the continuum states, even during recombination when the electron overlaps with the core. Whilst good quantitative agreement with the TDSE has been achieved for very high photon energies (compared to the ionization potential) and simple systems, the concordance breaks down at photon energies approaching the ionization potential, and is also expected to become inaccurate as the molecular size increases. However, the breakdown of the SFA does not necessarily prevent the use of the SPA and hence a quantum orbits approach, because the two sets of approximations are distinct and in many ways independent. The SFA is an approximation to the quantum mechanics: its essence is the partitioning of the wavefunction into core and continuum components, and the application of approximate, easily solvable Hamiltonians to each component. By contrast, the SPA, which leads to the quantum-orbits model, is a

general technique for recovering classical mechanics, or semi-classical interpretations, from the path-integral formulation of quantum mechanics. Mathematically, it is an asymptotic approximation to the path integral which becomes exact in the limit $\hbar \to 0$. Whether or not the underlying quantum mechanics is exact or an approximation is in principle immaterial for its application.

The implications of this for HHG and other strong-field processes is that going beyond the SFA does not necessarily mean losing the quantum orbits picture. It simply means that the nature of the quantum orbits will change. For example, including the effect of the core potential on the orbits may change their birth and return times, and will require corrections to the action integral. More drastic corrections in some hypothetical theory may change the parameterization (currently birth time, return time, and canonical momentum) itself. Regardless of the details, provided that i) the result is expressed as an integral over paths and ii) some of the important processes are amenable to classical interpretation, the SPA will be applicable and accurate, yielding both useful physical insight and computational efficiency by reducing the integral to a sum. This leads to the second motivation for experimental quantum path analysis: the ability to simultaneously observe these stationary paths at once independently and as a whole.

8.2   Related work

Quantum path analysis has been a key plank in the theory of HHG since the development of the Lewenstein model [368]. Almost all numerical models which do not involve direct solution of the TDSE involve some form of quantum path analysis, whilst experimental methods exist to restrict efficient harmonic production to only a single quantum path. However, the technique described in this chapter distinguishes the contributions of quantum paths in the single-atom response itself (as opposed to theoretical analyses starting from the laser field). One technique, similar in intent, is spectrography, which separates quantum paths based on their different emission times and frequencies. The limitation of spectrography is that it requires full amplitude and phase characterization of the harmonics which is generally not achieveable for experimental data. Spectrography is therefore primarily used in interpreting the results of simulations. Furthermore it cannot distinguish between trajectories with similar return times and emission frequencies.

Another technique for distinguishing trajectories lies in the different intensity dependence of their phases [381, 423]. This can be seen as a special case of the method described in this chapter, obtained when the control field is set equal to the drive field. In simulations, this technique has played an important role in understanding HHG, but the difficulty of eliminating macroscopic effects in experimental situations has made experimental observation challenging. Recently, experimental conditions with reduced macroscopic effects has provided access to the single-atom response, leading to the observation of intensity-dependent interference [424, 425]. This provides estimates of the relative amplitude and phase of the quantum paths, and the proposed method of this chapter can be seen as a generalization offering greater precision and flexibility.

The use of superposed drive and control fields of different wavelengths is essential to the present technique. Such *two-color* HHG experiments have an extensive history, with results including steering electronic motion [367, 426, 427] and temporal gating for isolated attosecond pulse generation [428–430]. For many-cycle drive fields, one ubiquitous effect of adding a control field is the appearance of even harmonics due to the broken symmetry between the positive and negative harmonics. This effect has been used to diagnose the emission times of the quantum paths [431]. Another concept related to the work of this chapter is *heterodyne mixing*, in which the effect of the weak control field is multiplied by the strong drive field. This has been demonstrated through a shift in the classical cutoff [432] due to a control field.

8.3 Action-shift induced by a weak control-field

This section commences the theoretical analysis of control-field mediated quantum-path analysis. The aim of this section is to derive, within in the framework of the SFA, the change in the action integral, and hence the change in phase, of a quantum path in response to a weak control field.

8.3.1 Single orbit analysis

One writes the electric field as the sum of a strong drive field $\mathcal{E}_{\mathrm{D}}(t)$ and a weak control field $\theta \mathcal{E}_{\mathrm{C}}(t)$, where θ parameterizes the strength of the control field:

$$\mathcal{E}(t) = \mathcal{E}_{\mathrm{D}}(t) + \theta \mathcal{E}_{\mathrm{C}}(t). \tag{8.1}$$

I shall calculate the first order change in the action of a quantum path caused by the control field. The action is written as

$$S^{(j)}(\omega) \approx S_{\mathrm{D}}^{(j)} + \theta S_{\mathrm{C}}^{(j)}. \tag{8.2}$$

The aim is to calculate $S_{\mathrm{C}}^{(j)}$, assuming that a set of control-field-free solutions (with subscript D) are already known. All of the other variables — for example the birth and return times of each quantum path, and the vector potential — may be written in a similar fashion.

To begin, one calculates the total derivative of the action (6.14) with respect to θ:

$$
\begin{aligned}
S_{\omega,\mathrm{C}}^{(j)} &= \frac{\mathrm{d}S_\omega}{\mathrm{d}\theta} \\
&= -\int_{t_{\mathrm{b}}^{(j)}}^{t_{\mathrm{r}}^{(j)}} \left[\mathbf{p}^{(j)} + \mathbf{A}(t)\right] \frac{\partial}{\partial\theta} \left[\mathbf{p}^{(j)} + \mathbf{A}(t)\right] \mathrm{d}t - \frac{\partial t_{\mathrm{b}}^{(j)}}{\partial\theta} \left\{ \frac{\left[\mathbf{A}(t_{\mathrm{b}}^{(j)}) + \mathbf{p}^{(j)}\right]^2}{2} + I_{\mathrm{p}} \right\} + \\
&\quad \frac{\partial t_{\mathrm{r}}^{(j)}}{\partial\theta} \left\{ \frac{\left[\mathbf{A}(t_{\mathrm{r}}^{(j)}) + \mathbf{p}^{(j)}\right]^2}{2} + I_{\mathrm{p}} \right\} + \omega \frac{\partial t_{\mathrm{r}}^{(j)}}{\partial\theta}.
\end{aligned}
\tag{8.3}
$$

The second and third summands arise from the dependence of the action integral on its start- and end-points. The second summand is identically zero due to the birth-time saddle-point condition (6.17), whilst the third and fourth summands cancel due to the recollision-time saddle-point condition (6.18). Equation (8.3) therefore reduces to

$$S_{\omega,\mathrm{C}}^{(j)} = -\int_{t_{\mathrm{b}}^{(j)}}^{t_{\mathrm{r}}^{(j)}} \left[\mathbf{p}^{(j)} + \mathbf{A}(t)\right] \frac{\partial\mathbf{A}(t)}{\partial\theta} \mathrm{d}t - \frac{\partial\mathbf{p}^{(j)}}{\partial\theta} \int_{t_{\mathrm{b}}^{(j)}}^{t_{\mathrm{r}}^{(j)}} \left[\mathbf{p}^{(j)} + \mathbf{A}(t)\right] \mathrm{d}t. \tag{8.4}$$

The second integral is identically zero due to the momentum saddle-point condition (6.15). Writing the total field as a sum of drive and control fields, so that $\mathbf{A}(t) = \mathbf{A}_{\mathrm{D}}(t) + \theta \mathbf{A}_{\mathrm{C}}(t)$, one obtains the result

$$S_{\omega,\mathrm{C}}^{(j)} = -\int_{t_{\mathrm{b}}^{(j)}}^{t_{\mathrm{r}}^{(j)}} \left[\mathbf{p}^{(j)} + \mathbf{A}_{\mathrm{D}}(t)\right] \mathbf{A}_{\mathrm{C}}(t) \mathrm{d}t. \tag{8.5}$$

Equation (8.5) is invariant to the gauge of the control field vector potential — the addition of a constant to $\mathbf{A}_{\mathrm{C}}(t)$ does not change its result because of the momentum saddle-point condition.

Therefore, one may replace $\mathbf{A}_C(t)$ with $\mathbf{v}_C(t) = \mathbf{A}_C(t) - \mathbf{A}_C(t_b^{(j)})$, the velocity of the electron due to the control field. Noting that the electron velocity due to the drive field is $\mathbf{v}_D^{(j)}(t) = \mathbf{p}^{(j)} + \mathbf{A}_D(t)$, one obtains

$$S_{\omega,C}^{(j)} = -\int_{t_b^{(j)}}^{t_r^{(j)}} \mathbf{v}_D^{(j)}(t)\mathbf{v}_C(t)\,\mathrm{d}t. \tag{8.6}$$

Equation (8.6) represents heterodyne mixing between the drive and control fields. It has been presented for a more specific case in ref. [431]. The electron velocity due to the weak control field is amplified by that due to the strong drive field. Heterodyne mixing in high-harmonic generation has previously been discussed in a classical framework [432], leading to shifts in the classical cutoff.

8.3.2 Multiple control fields and multiple orbits

I now incorporate the previous result (8.5) into an expression for the total harmonic emission in the presence of multiple control fields using the quantum orbits model. I begin by writing the dipole response using as a sum over quantum paths, each given by (6.19):

$$\mathscr{D}(\omega) = \sum_j \mathscr{B}^{(j)}(\omega)e^{iS_\omega^{(j)}}. \tag{8.7}$$

Below the cutoff, the sum is over all trajectories, whilst above the cutoff, only the dominant trajectories are included. The trajectories are labelled by $j = \alpha\beta m$ as explained in section 6.2.4. I shall not attempt to make predictions around the cutoff because the sum is given by the uniform approximation (6.20), a perturbative treatment of which is beyond is scope of this dissertation. However, the simulations presented below show that the interaction of the control field with the uniform approximation does not cause any unexpected effects.

The applied electric field is the drive field plus an arbitrary number of control fields, labelled by subscript:

$$\mathscr{E}(t) = \mathscr{E}_D(t) + \theta_1\mathscr{E}_1(t) + \theta_2\mathscr{E}_2(t) + \dots \tag{8.8}$$

I now introduce another key assumption which is that the control fields are so weak that their effect on the trajectory amplitudes $\{\mathscr{B}^{(j)}(\omega)\}$ may be neglected, and only their effect on $S_\omega^{(j)}$ needs

to be considered. This is justified because the action is typically many multiples of 2π, so that only a small relative change in the parameters causes a significant effect upon the interference. I also test this approximation numerically in section 8.5.1.1 below.

For conciseness I shall introduce a column-vector notation for the control fields and their associated quantities. For example,

$$\bar{\theta} = [\theta_1, \theta_2, \ldots]^{\mathrm{T}} \tag{8.9}$$

where superscript T denotes transpose. One may therefore abbreviate (8.8) as $\mathscr{E}(t) = \mathscr{E}_{\mathrm{D}}(t) + \bar{\theta}^{\mathrm{T}}\bar{\mathscr{E}}_{\mathrm{C}}(t)$ where $\bar{\mathscr{E}}_{\mathrm{C}}(t) = [\mathscr{E}_1(t), \mathscr{E}_2(t), \ldots]^{\mathrm{T}}$.

Through linearity of (8.5), the phase shift due to multiple control fields is the sum of the individual phase shifts:

$$\begin{aligned} S_\omega^{(j)} &= S_{\omega,\mathrm{D}}^{(j)} + \theta_1 S_{\omega,1}^{(j)} + \theta_2 S_{\omega,2}^{(j)} + \ldots \\ &= S_{\omega,\mathrm{D}}^{(j)} + \bar{\theta}^{\mathrm{T}} \bar{S}_{\omega,\mathrm{C}}^{(j)}. \end{aligned} \tag{8.10}$$

Substitution of (8.10) into (8.7) shows that dipole response at any given frequency has the form of a sum of plane waves in the control-field amplitude space $(\theta_1, \theta_2, \ldots)$:

$$\mathscr{D}(\omega; \bar{\theta}) = \sum_j \mathscr{B}^{(j)}(\omega) e^{iS_{\omega,\mathrm{D}}^{(j)}} e^{i\bar{\theta}^{\mathrm{T}} \bar{S}_{\omega,\mathrm{C}}^{(j)}}. \tag{8.11}$$

The amplitude of each of these plane waves is $\mathscr{B}^{(j)}(\omega) e^{iS_{\omega,\mathrm{D}}^{(j)}}$, whilst their wave-vector is the *control-field sensitivity* vector $\bar{S}_{\omega,\mathrm{C}}^{(j)} = [S_{\omega,1}^{(j)}, S_{\omega,2}^{(j)}, \ldots]^{\mathrm{T}}$.

The key idea of this chapter is that one may treat the $\mathscr{B}^{(j)}(\omega) e^{iS_{\omega,\mathrm{D}}^{(j)}}$ as unknowns, and then fit (8.11) to an experimentally or numerically acquired sample of $\mathscr{D}(\omega; \bar{\theta})$, taken over some set of control-field amplitudes. As long as each trajectory possesses a unique control-field sensitivity, one may expect that such a fit will be possible and unique. Indeed, as I shall show, the distribution of the control-field sensitivity vectors is crucial in determining the success of the proposed method.

To visualize the control-field sensitivity vectors, it is useful to plot them as points in *control-field sensitivity* space $(\alpha_1, \alpha_2, \ldots)$. One can connect such a diagram with the multi-dimensional Fourier transform of $\mathscr{D}(\omega; \bar{\theta})$ taken along $\theta_1, \theta_2, \ldots$:

$$
\begin{aligned}
\tilde{\mathscr{D}}(\omega; \bar{\alpha}) &= (2\pi)^{L/2} \sum_j \mathscr{B}^{(j)}(\omega) e^{iS_{\omega,\mathrm{D}}^{(j)}} \delta(\alpha_1 - S_{\omega,1}^{(j)}, \alpha_2 - S_{\omega,2}^{(j)}, \ldots) \\
&= (2\pi)^{L/2} \sum_j \mathscr{B}^{(j)}(\omega) e^{iS_{\omega,\mathrm{D}}^{(j)}} \delta(\bar{\alpha} - \bar{S}_{\omega,\mathrm{C}}^{(j)})
\end{aligned}
\tag{8.12}
$$

where L is the number of control fields. In this control-field-sensitivity space, each quantum trajectory is Dirac-delta function located at $\bar{S}_{\omega,\mathrm{C}}^{(j)}$.

8.4 Example

I shall consider quasi-monochromatic control fields — that is, pulses which are significantly longer than the drive field. The control-field frequency is considered to be fixed, and one may parameterize its remaining degrees of freedom by the (real-valued) amplitude of the in-phase and quadrature components (with respect to the drive field), denoted θ_I and θ_Q respectively. I use this rectangular representation because the resulting control field is a linear combination of basis functions, and it thus fits naturally into the perturbative formalism. Of course, in the laboratory it may be more convenient to scan the amplitude and phase of the control field, but this provides the same information and may be digitally converted into rectangular coordinates.

Mathematically, the control field is

$$
\mathscr{E}_\mathrm{C}(t) = E_0 \left[\theta_\mathrm{I} \cos(\omega_\mathrm{C} t) + \theta_\mathrm{Q} \cos\left(\omega_\mathrm{C} t - \frac{\pi}{2}\right) \right].
\tag{8.13}
$$

The control field amplitudes $(\omega_\mathrm{Q}, \omega_\mathrm{I})$ are scaled by the peak amplitude of the drive field E_0. I will consider the effect of this control field on two drive fields: a monochromatic field and a few-cycle pulse.

8.4.1 Monochromatic drive field

In this section I consider a monochromatic drive field with frequency ω_{D} and amplitude E_0. I consider first-, second- and third-order trajectories (i.e. $\beta = 1, 2, 3$) born in first half cycle and the one previous to it (i.e. $\gamma = m - \beta = -1, 0$). A quantum orbit analysis for this field was given in Fig. 6.6.

For each of these orbits, I calculated the control field sensitivities $S_{\omega,Q}^{(j)}$, $S_{\omega,I}^{(j)}$ using (8.5). For $\omega_{\mathrm{C}} > \omega_{\mathrm{D}}$, I found that the control-field sensitivities decreased, bringing them closer to the origin and hence less distinct. Therefore, I shall concentrate on control fields $\omega_{\mathrm{C}} < \omega_{\mathrm{D}}$. A sampling of the results is shown in Fig. 8.1. Long and short orbits begin separately (at low harmonic orders) and move towards each other, coming together at cutoff. For the first-order orbits ($\beta = 1$) the short trajectories have smaller action sensitivity than long trajectories, but for higher-order trajectories this behavior is more complex, depending on the control field frequency. This carries implications for the goal of separating the orbits and will be a consideration in choosing a control field frequency.

Another fact evident from Fig. 8.1 is that orbits of constant birth event γ form continuous paths in action-sensitivity space, with the long trajectory of one merging into the short trajectory of another at low harmonic orders. This is not surprising as classically these correspond to the "same" electrons.

In a monochromatic field, the orbits are the same from one half-cycle to the next, except for a change in sign. However, the relative phase of the control and drive field advances by $\Delta\phi_{\mathrm{CD}} = \omega_{\mathrm{C}}\pi/\omega_{\mathrm{D}} - \pi = (\omega_{\mathrm{C}}/\omega_{\mathrm{D}} - 1)\pi$ every half-cycle. An advance of the control-to-drive field relative phase results in a rotation in action-sensitivity space. This is reflected in Fig. 8.1: the action-sensitivities of orbits with birth event $\gamma = 0$ are simply an anticlockwise rotation, by $\Delta\phi_{\mathrm{CD}}$, of the corresponding orbits with $\gamma = -1$. This carries implications for distinguishing the orbits of different half-cycles. For n half-cycles, the orbits will span an angle $n\Delta\phi_{\mathrm{CD}}$ in action-sensitivity space. If $n\Delta\phi_{\mathrm{CD}} > 2\pi$, then the orbits may "wrap around" and overlap one another. However by careful choice of $\omega_{\mathrm{C}}/\omega_{\mathrm{D}}$, so that $\Delta\phi_{\mathrm{CD}}$ does not divide evenly into 2π, one may angularly interleave the
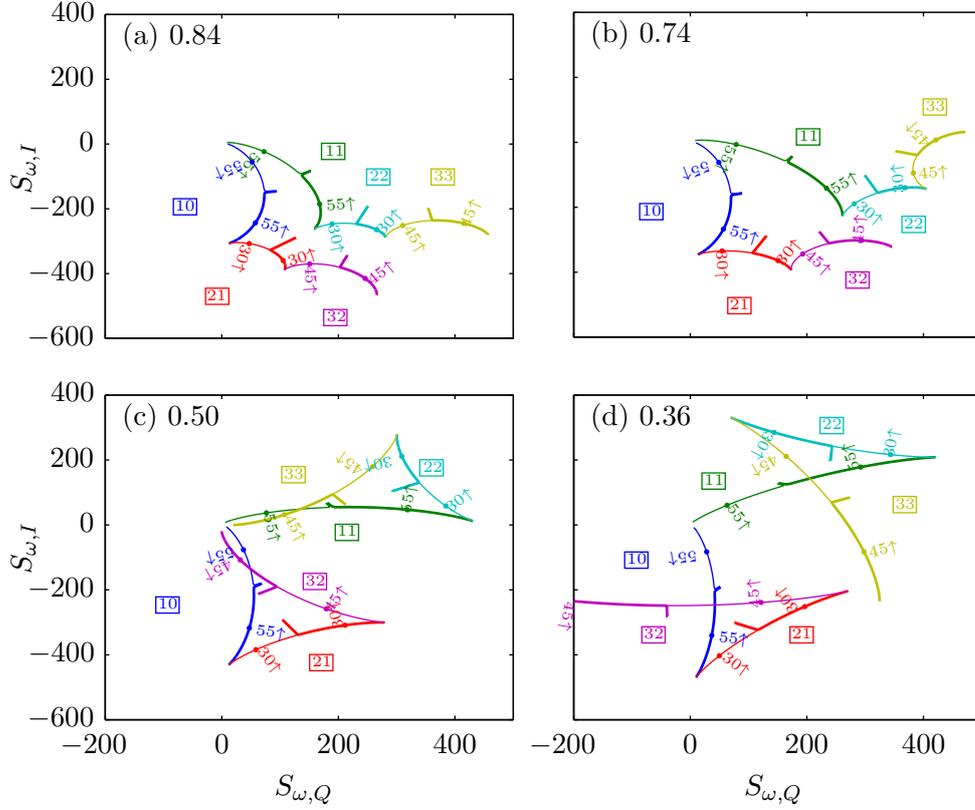
Figure 8.1: Action sensitivities of the quantum orbits of a monochromatic drive field of intensity $5 \times 10^{14}\,\mathrm{W/cm^2}$ in Argon subject to in-phase and quadrature monochromatic control fields. The long/short orbit pairs are identified by their βm labels (boxed) as well as a color code identical to that in Fig. 6.6. Within each pair, thin (thick) lines represent short (long) orbits. For each orbit, the position of one harmonic is indicated; the arrow indicates the direction of increasing harmonic number. Each subfigure has fractional control frequency ω_C/ω_D indicated at the top left.

half-cycles. For a pulsed drive field the orbits are not identical for each half-cycle. Nonetheless, even for a few-cycle pulse the similarity between the half-cycles is sufficient for a similar rotation principle to apply.

8.4.2 Pulsed drive field

Having established some basic properties of the orbits in action-sensitivity space, I now consider a few-cycle pulsed drive field, with 5.0 fs duration, 800 nm centre wavelength and peak intensity
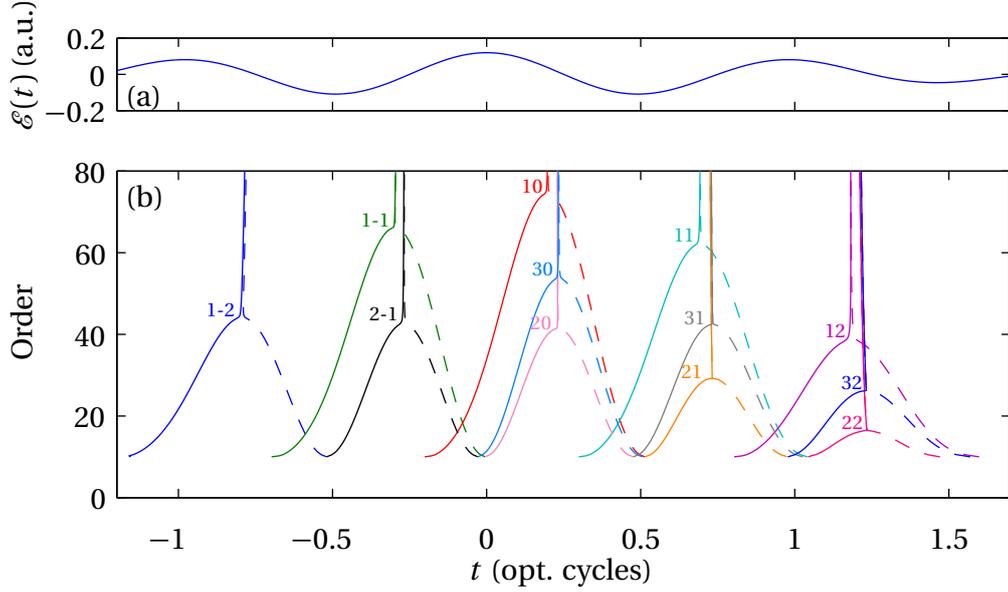
Figure 8.2: (a) Pulsed drive field, zoomed in on central region. (b) Harmonic order versus recombination time of the quantum orbits of a 5.0 fs, 800 nm, $5 \times 10^{14}\,\text{Wcm}^{-2}$ drive field in Argon. The orbits are labelled by excursion length and recombination event βm, and short (long) orbits are solid (dashed).

$5 \times 10^{14}\,\text{Wcm}^{-2}$, in Argon. Mathematically, the drive field is

$$\mathscr{E}_{\text{D}}(t) = E_0 \cos^2 \left[\frac{2\cos^{-1}(2^{1/4})}{T_{\text{FWHM}}} t \right] \cos(\omega_{\text{D}} t) \tag{8.14}$$

where I use a cosine-squared pulse for analytical tractability. The quantum path analysis is shown in Fig. 8.2.

For these orbits, I evaluated the action-sensitivities to the in-phase and quadrature control fields using (8.5). A sample of the results are shown in Fig. 8.3. For subfigures (a) and (b), the control-to-drive phase advance per half-cycle $\Delta\phi_{\text{CD}}$ is small enough so that over the five half-cycles the net rotation is less than 2π. The first-order trajectories therefore do not overlap. Also, for these values of $\omega_{\text{C}}/\omega_{\text{D}}$ the higher-order trajectories have larger action-sensitivity and do not overlap the first-order trajectories. In particular, for $\omega_{\text{C}}/\omega_{\text{D}} = 0.6$ one has $\Delta\phi_{\text{CD}} = 2\pi/5$, so the 5 trajectories are equally spread around the circle. However, for $\omega_{\text{C}}/\omega_{\text{D}} = 0.5$ (not shown) $\Delta\phi_{\text{CD}} = -\pi/2$ so that orbits separated by 4 half-cycles coincide azimuthally. For Fig. 8.3(c), the control-
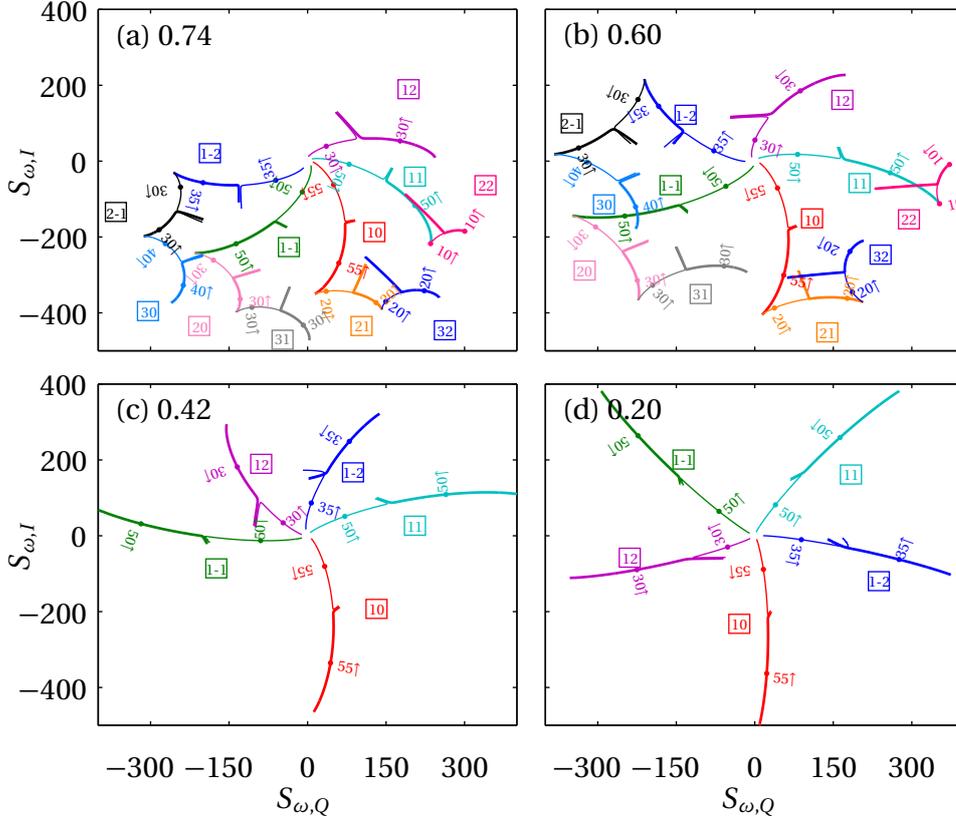
Figure 8.3: Action sensitivities of the quantum orbits of Fig. 8.2 to in-phase and quadrature monochromatic control fields. The long/short orbit pairs are identified by their βm labels (boxed) as well as a color code identical to that in Fig. 8.2. Within each pair, thin (thick) lines represent short (long) orbits. For each orbit, the position of one harmonic is indicated; the arrow indicates the direction of increasing harmonic number. In subfigures (c) and (d) the second- and higher-order trajectories have been dropped for clarity. Each subfigure has fractional control frequency ω_C/ω_D indicated at the top left.

drive phase slip is large enough so that first and last half-cycles have exchanged angular positions, whilst for Fig. 8.3(d) $\Delta\phi_{CD} = \frac{2}{5}2\pi$ so the orbits circle around the origin twice but fall neatly in between one another. For subfigures (c) and (d), the higher-order trajectories move radially inward and hence overlap with the first-order trajectories, so they have been dropped for clarity.

8.5 Limitation of the perturbative quantum path analysis

The previous section discussed the positions of the trajectories in control-field-sensitivity space, and showed that for certain control fields the trajectories can be well separated, without introducing any precise definition of the "separateness". However, in common with all Fourier-transform

interferometry, the size of the components in the Fourier domain is just as important as their location in determining whether the interferometric component can be isolated. Since the perturbative analysis of section 8.3 culminating in (8.12) indicates that the trajectories are delta functions i.e. infinitely small, their actual size in the Fourier domain must be related to the accuracy of the pertinent assumptions — namely the linearity of the change in action and the invariance of the trajectory amplitudes with respect to the control field amplitudes. The resolution in the Fourier domain is also limited by the range of control-field amplitudes over which the dipole response is known. In this section I explore these limitations and relate them to the size of the trajectories in the Fourier domain — and hence the success or failure of the present endeavor.

8.5.1 Brute-force study

In this section I verify the assumptions and predictions of the perturbative quantum orbit formalism using exact quantum orbit calculations. I used the CW drive and control fields presented in section 8.4, and performed a quantum orbit analysis over a two-dimensional grid of control field amplitudes $\theta_Q, \theta_I \in [-0.1, 0.1]$ in 0.01 increments. The zero control field case is compared with several extreme cases in Fig. 8.4. In most cases, the control fields cause only a minor change in the time-frequency structure and the excursion times. However, in one of the cases ($\theta_1 = 0, \theta_2 = -0.1$), the cutoff is suppressed by approximately 15 harmonic orders by the control field.

This results set enabled me to examine in turn the assumptions of a constant dipole amplitude, and the linear action variation.

8.5.1.1 Variation of dipole amplitude

From the quantum orbits analysis, I calculated the control-field dependent dipole amplitude $\mathscr{B}^{(j)}(\omega; \bar{\theta})$ using (6.19). Figure 8.5 (a) and (b) shows the results for harmonic order 40, well below cutoff. It is apparent that the variation of the dipole amplitude and phase with control field amplitude is small, and predominatly linear, as will be quantified below. The conclusions are similar well above the cutoff.

Figure 8.5 (c) and (d) shows the same results for harmonic order 70, around cutoff. Here, the dependence is not linear or smooth, with the in-phase control field causing a pronounced en-
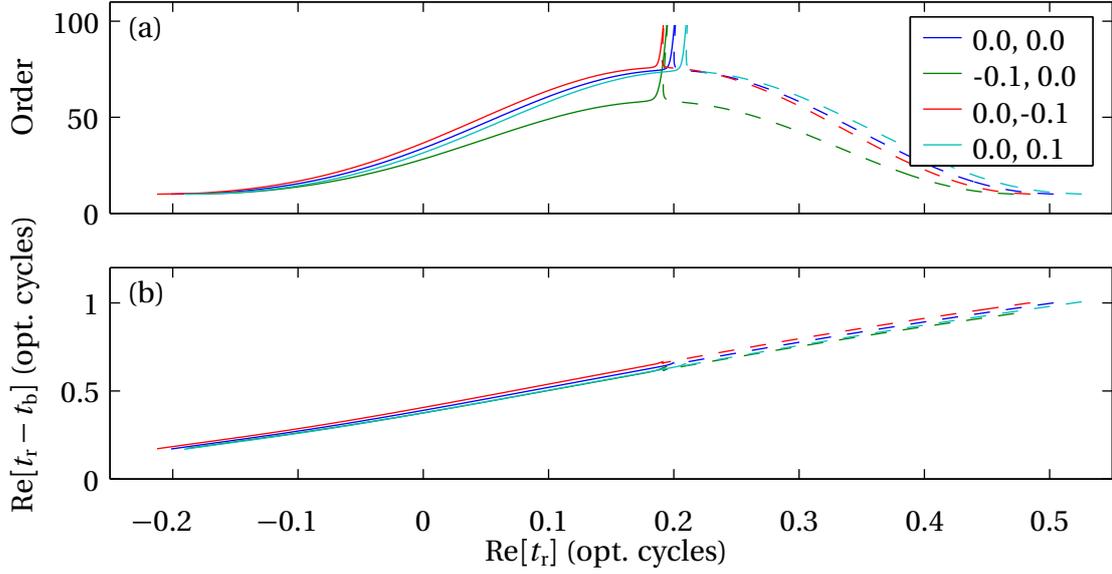
Figure 8.4: First-order quantum orbits subject to various monochromatic control fields. (a) Harmonic order versus real part of recombination time. The in-phase and quadrature control field amplitudes (θ_1, θ_2) are shown in the legend. Short (long) trajectories are solid (dashed). (b) Real part of excursion time versus real part of combination time.

hancement around $\theta_I = -0.02$, accompanied by a sharp phase discontinuity. This is associated with the cutoff being suppressed, as was seen in Fig. 8.4. This will not necessarily cause a problem because around the cutoff the net dipole response is not simply the sum of each trajectory's response as in (8.7); the uniform approximation is needed. Incorporating the perturbative formalism into the uniform approximation is beyond the scope here so I simply note this as a potential hazard and move on.

Linear phase variation of the dipole amplitude $\mathcal{B}^{(j)}(\omega; \bar{\theta})$ will result in a shift of the trajectories in control-field-sensitivity space away from the predicted $\bar{s}_\omega^{(j)}$. To quantify this, one writes the dipole amplitude as a plane wave in $\bar{\theta}$:

$$\mathcal{B}^{(j)}(\omega; \bar{\theta}) \approx \mathcal{B}^{(j)}(\omega; 0) \exp[i\bar{\theta}^{\mathrm{T}} \bar{\mathcal{C}}^{(j)}(\omega)]. \tag{8.15}$$

By fitting (8.15) to the amplitudes calculated by brute force, the coefficients $\bar{\mathcal{C}}^{(j)}(\omega)$ are obtained.
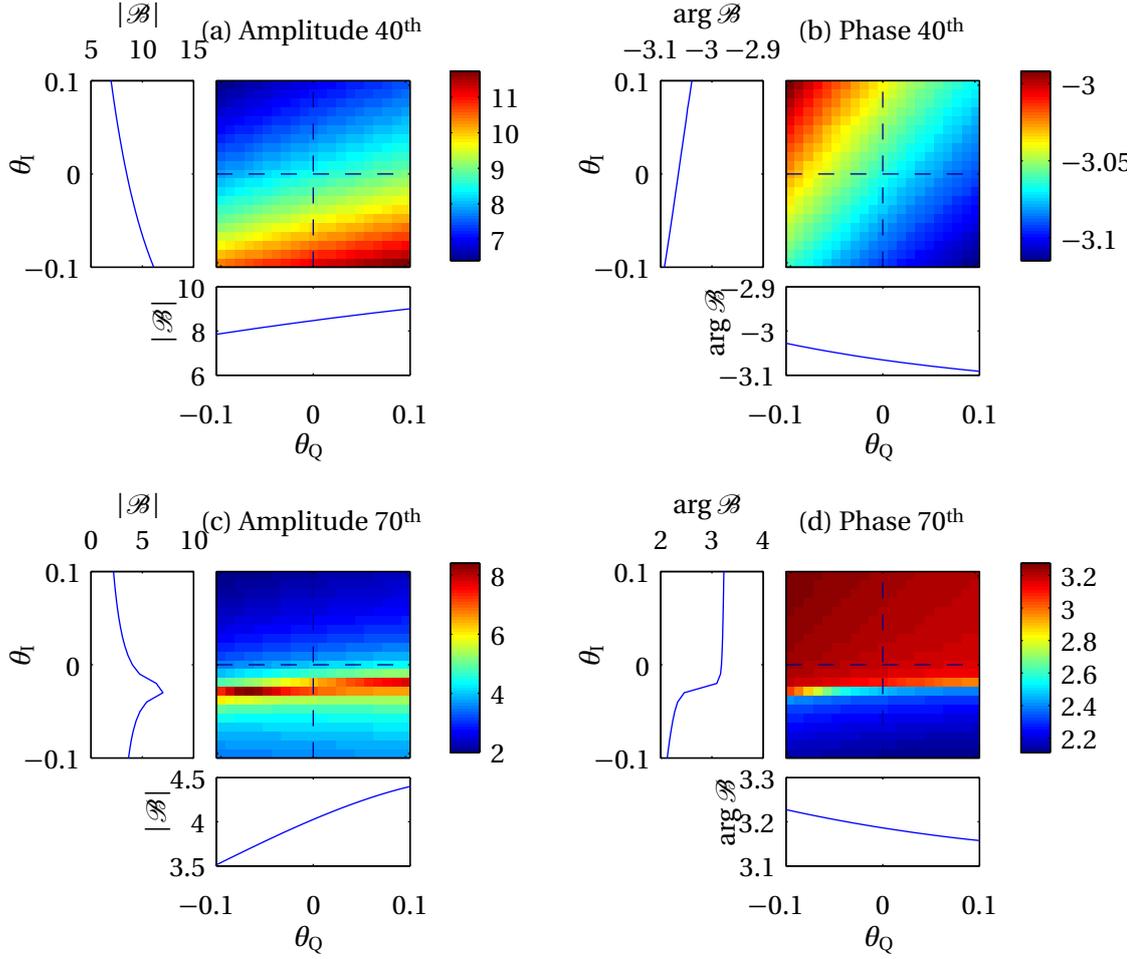
Figure 8.5: Amplitude ((a),(c)) and phase ((b),(d)) of the dipole amplitude $\mathscr{B}^{(j)}(\omega; \bar{\theta})$ for the long trajectory at the 40th ((a), (b)) and 70th ((c),(d)) harmonics. The drive field is monochromatic with identical parameters to those of Fig. 8.1 — an intensity of $5 \times 10^{14}\,\text{W/cm}^2$ in Argon. The colour plots are over the whole (θ_Q, θ_I) space, whilst the line plots show cross sections at $\theta_Q = 0$ and $\theta_I = 0$ respectively.
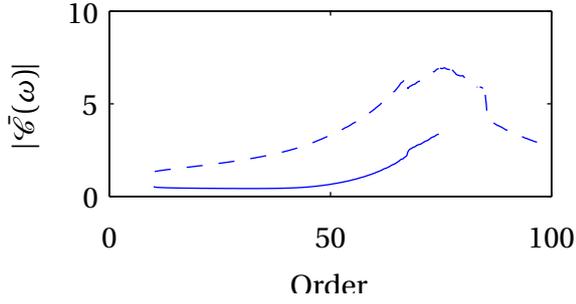
Figure 8.6: Strength of linear variation of dipole amplitude versus frequency for a first-order trajectory of the CW drive field of section 8.4. The short (solid) and long (dashed) trajectories are shown. The short trajectory is recessive and is dropped after cutoff.
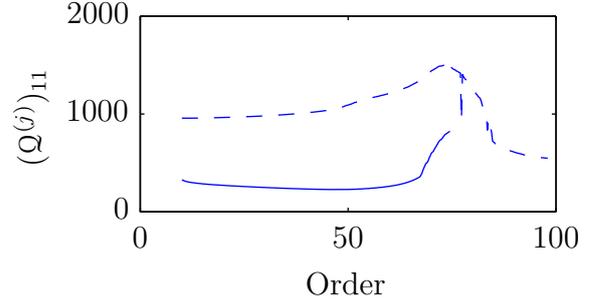
Figure 8.7: Curvature of the action versus harmonic order for short (solid) and long (dashed) trajectories. The short trajectory is recessive and is dropped after cutoff.

The magnitude of this quantity,

$$|\bar{\mathscr{C}}^{(j)}(\omega)| = \left[[\mathscr{C}_1^{(j)}(\omega)]^2 + [\mathscr{C}_2^{(j)}(\omega)]^2 + \ldots \right]^{1/2}, \tag{8.16}$$

is the scalar wavenumber of the plane wave in (8.15) and is a useful indicator of the strength of the linear variation of $\mathscr{B}^{(j)}(\omega; \bar{\theta})$ with $\bar{\theta}$. This is plotted in Fig. 8.6. The maximum value is less than 10, an order of magnitude smaller than the predicted values of $\bar{S}_\omega^{(j)}$. Therefore, I conclude that in the regions of validity of (8.7) — i.e. well above or below the cutoff of a trajectory — the approximation of a constant dipole amplitude is sound.

8.5.1.2 Nonlinearity of action change

From the exact quantum orbit analysis, I computed the control-field dependent action $S_\omega^{(j)}(\bar{\theta})$ using (6.14) directly. Figure 8.8 compares a representative subset of the results with the linear approximation, given by (8.5) and (8.2), and is organized as follows: the first row (subfigures (a)–(c)) shows the effect of the in-phase control field (with zero quadrature control field), and the second row (subfigures (d)–(f)) shows the effect of the quadrature control field. The first column (subfigures (a) and (b)) are at harmonic 40, significantly below cutoff, whilst the second and third columns are at harmonics around cutoff and above cutoff respectively. The real parts are shown
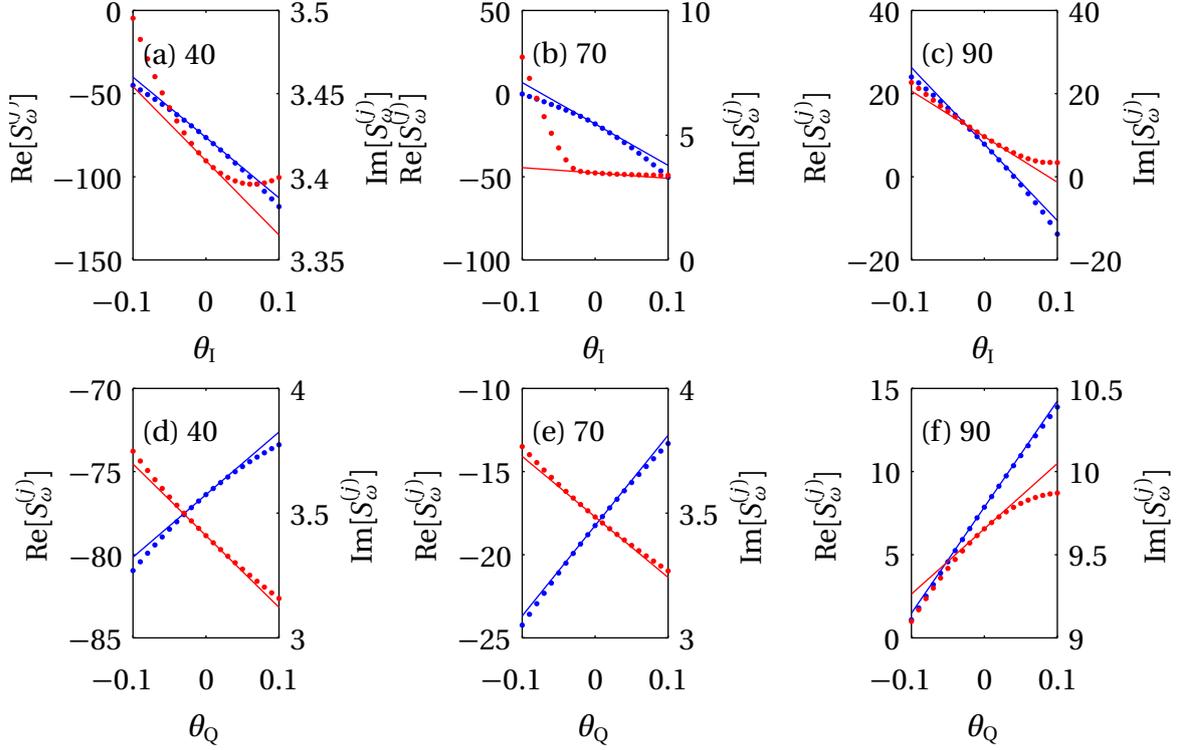
Figure 8.8: Real (blue, left axis) and imaginary (red, right axis) parts of the action of the long (dominant) trajectory, subject to in-phase (subfigures (a)–(c) in the top row), and quadrature (subfigures (d)–(f) in the top row) control fields. The harmonic orders are 40 (subfigures (a) and (d), first column), 70 (subfigures (b) and (e), second column), and 90 (subfigures (c) and (f), third column).

in blue scaled according to the left vertical axes; the imaginary parts are in red and use the right vertical axes. Moving to the interpretation of the figure, it is evident that in all cases, the exact results are tangent to the linear approximation at $\theta = 0$, verifying the correctness of the formalism. Concerning the real parts, the linear approximation appears by eye to be reasonable for all the cases shown. The imaginary parts show similar agreement except for the in-phase control field at harmonic 70. Here, the exact results show a marked departure from the linear approximation at $\theta_I = -0.02$. This corresponds to the suppression of the cutoff by the control field and highlights the potential problems of this approach around cutoff.

Figure 8.8 suggests that the linear approximation is accurate for most cases. However it in no way covers the entire dataset and doesn't show the complete picture of the nonlinearity, which is a two-dimensional surface dependent on θ_I and θ_Q. I required a quantitative measure of the curva-

ture which could be systematically applied to the entire dataset. To do so, I included a quadratic term in the power series expansion of the control-field dependent action:

$$S^{(j)}(\omega) \approx S_D^{(j)} + \bar{\theta}^T \bar{S}_C^{(j)} + \frac{1}{2} \bar{\theta}^T \mathcal{Q}^{(j)} \bar{\theta}. \tag{8.17}$$

Here $\mathcal{Q}^{(j)}$ is a matrix which describes all combinations of quadratic self- and cross- couplings between the control fields. I then obtained $\mathcal{Q}^{(j)}$ by fitting (8.17) to the exact quantum orbit analysis. In general, $\bar{\theta}^T \mathcal{Q}^{(j)} \bar{\theta}$ describes a paraboloid, the contours of which are arbitrarily aligned to the coordinate axes. The direction and magnitudes of the principal axes are given by diagonalization. Here, I not concerned with the details of the geometry, but only the magnitude of the curvature. So, without loss of generality one may choose a co-ordinate system in which $\mathcal{Q}^{(j)}$ is diagonal with largest eigenvalue $(\mathcal{Q}^{(j)})_{11}$. This eigenvalue is equal to the largest quadratic coefficient over all possible directions in $\bar{\theta}$ space, and therefore forms a convenient measure of the curvature of the action. This quantity is plotted for the long and short first-order trajectories in Fig. 8.7, and for the long trajectory (the worst case) is around 1000. One obtains a feeling for the significance of this number by noting that the erroneous phase introduced by the nonlinearity is equal to $1/2(\mathcal{Q}^{(j)})_{11}\theta_1^2$. The control field amplitude at which the error equals one radian is therefore $\sqrt{2/(\mathcal{Q}^{(j)})_{11}} \approx 0.045$ for the monochromatic example used throughout this chapter.

8.5.2 Choice of control-field amplitude scan range

The previous section showed that, away from cutoff, nonlinearity of action change is the dominant distortion. The other factor that influences the width of the orbits in control-field-sensitivity space is the intrinsic resolution in that space, determined by the range of control fields over which $D(\omega; \bar{\theta})$ is known. One may represent this mathematically via a window function $\mathcal{W}(\bar{\theta})$ which is unity at the origin $\bar{\theta} = 0$ and smoothly decays when approaching the edge of the sampled range of control field amplitudes. For the discussion which follows it is necessary to choose a particular form for the window function — here I use an isotropic Gaussian function

$$\mathcal{W}(\bar{\theta}) = \exp(-\frac{|\bar{\theta}|^2}{2w^2}) \tag{8.18}$$

where $w = w_{\text{FWHM}}/(2\sqrt{2\log 2})$ parameterises the width of the window. Incorporating the window function and the nonlinearity in the action into (8.11) gives

$$\mathscr{D}(\omega;\bar{\theta}) = \sum_j \mathscr{B}^{(j)}(\omega)\mathscr{G}(\bar{\theta};\mathscr{Q}^{(j)})e^{iS_{\omega,\text{D}}^{(j)}}e^{i\bar{\theta}^{\text{T}}\bar{S}_{\text{C}}^{(j)}}. \tag{8.19}$$

where I have incorporated the window function and the effect of the nonlinearity into a single function

$$\mathscr{G}(\bar{\theta};\mathscr{Q}) = \exp(-\frac{|\bar{\theta}|^2}{2w^2} + i\frac{1}{2}\bar{\theta}^{\text{T}}\mathscr{Q}\bar{\theta}). \tag{8.20}$$

The Fourier transform is then

$$\tilde{\mathscr{D}}(\omega;\bar{\alpha}) = (2\pi)^{-L}\sum_j \mathscr{B}^{(j)}(\omega)\tilde{\mathscr{G}}(\bar{\alpha} - \bar{S}_{\omega,\text{C}}^{(j)};\mathscr{Q}_{\omega,\text{C}}^{(j)})e^{iS_{\omega,\text{D}}^{(j)}} \tag{8.21}$$

where the Fourier transform of $\mathscr{G}(\bar{\theta};\mathscr{Q})$ is

$$\tilde{\mathscr{G}}(\bar{\alpha};\mathscr{Q}) = \sqrt{\frac{i^L}{\det\mathscr{Q}}}\exp\left[-\frac{1}{2}\bar{\alpha}^{\text{T}}(w^{-2}I - i\mathscr{Q})^{-1}\bar{\alpha}\right]. \tag{8.22}$$

To find the spread of $\tilde{\mathscr{G}}(\bar{\alpha};\mathscr{Q})$ one may temporarily adopt a co-ordinate system in which $\mathscr{Q}$ is diagonalized. For simplicity I also assume that $\mathscr{Q}$ is real. Then

$$\tilde{\mathscr{G}}(\bar{\alpha};\mathscr{Q}) = \sqrt{\frac{i^L}{\det\mathscr{Q}}}\exp\left[-\frac{1}{2}\sum_{l=1}^{L}\alpha_l^2\frac{w^{-2} + i\mathscr{Q}_{ll}}{w^{-4} + \mathscr{Q}_{ll}^2}\right] \tag{8.23}$$

Here, $\mathscr{Q}_{ll}$ are the eigenvalues of $\mathscr{Q}$. The spread $\Delta\alpha_l$ of $\tilde{\mathscr{G}}(\bar{\alpha};\mathscr{Q})$ along α_l is the square root of the inverse of the real part of the corresponding coefficient in (8.23)

$$\Delta\alpha_l = \sqrt{\frac{w^{-4} + \mathscr{Q}_{ll}^2}{w^{-2}}}. \tag{8.24}$$

Minimization of (8.24) provides a natural choice of window function width

$$w = \max_l \mathscr{Q}_{ll}^{-1/2} \tag{8.25}$$
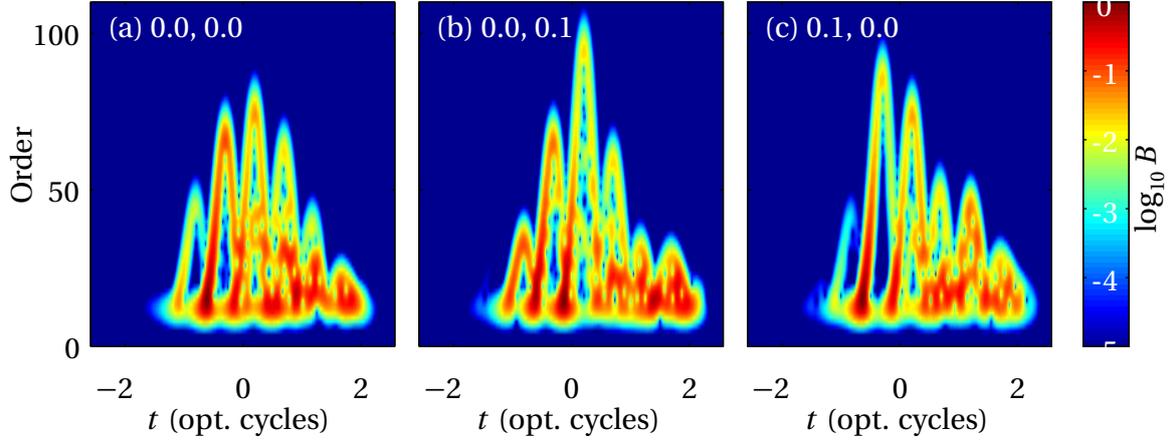
Figure 8.9: Spectrogram of dipole emission with 5 fs, $5 \times 10^{14}\,\mathrm{Wcm^{-2}}$, 800 nm pulse with various monochromatic control fields (θ_Q, θ_I) with relative frequency $\omega_C/\omega_D = 0.45$.

where one chooses the largest nonlinear coefficient as it is most likely to cause a problem.

For the examples of section 8.5.1.2, the largest eigenvalues of the action nonlinearity are about 1000 across most of the spectrum. The application of (8.25) suggests that a window width of $w = 0.03$ is therefore appropriate. However, $D(\omega; \bar{\theta})$ must be known over a window several times larger to eliminate artifacts from truncation.

8.6 Simulated experiment

Having used a perturbative extension to the quantum orbits model to make predictions of the effect of applying control fields, I now test these predictions against simulations of the single-atom response using the 1D TDSE.

The drive and control fields were those of section 8.4.2. The control field amplitudes were scanned over $\theta_I, \theta_Q \in [-0.1, 0.1]$ in 0.01 increments, and at each step, the dipole response was calculated using a one-dimensional TDSE code. I denote the results $\mathscr{D}(\omega; \theta_Q, \theta_I)$. Figure 8.9 shows a sample of the resulting dipole responses for different control field combinations, represented using spectrograms. In subfigures (b) and (c), the control fields are at the extremes of the scan range, and the cutoffs have been modified significantly from the zero control field case. However, the window function introduced in section 8.5.2 will suppress these extreme cases.

Figure 8.10 shows the amplitude and phase of the entire results set for the 45th harmonic.
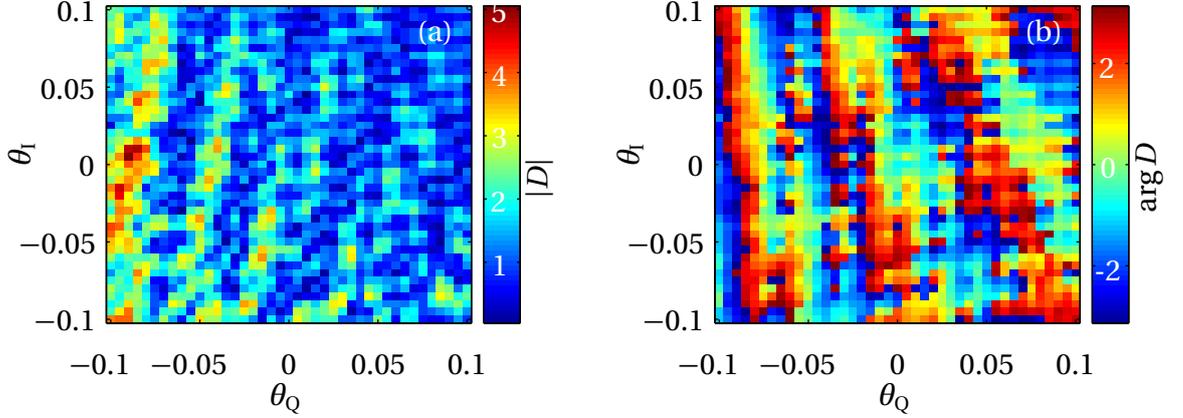
Figure 8.10: (a) Amplitude and (b) phase of the computed dipole response $\mathscr{D}(\omega; \theta_Q, \theta_I)$ at the 49th harmonic.

Whilst the eye can discern some regular patterns in both, a normal fringe pattern is not observed because there are many Fourier components present. Around $(\theta_Q, \theta_I) = (-0.1, 0)$, an enhancement of the amplitude over many periods of the pattern shows that control field is starting to distort the amplitude of the quantum paths, rather than just their phases.

The results are best interpreted in control-field sensitivity space, found by taking discrete Fourier transforms along both control field amplitude co-ordinates yielding $\tilde{\mathscr{D}}(\omega; \bar{\alpha})$. This is plotted for four harmonics in Fig. 8.11, along with the control-field sensitivities $\bar{S}_{\omega,C}^{(j)}$ expected from the perturbative analysis. The $\alpha\beta m$ labelling system for the orbits has been extended slightly: above cutoff, $\alpha = C$ refers to "coalesced" long and short trajectories. For the 30th harmonic, there are many peaks and it is difficult to assign the orbits. However, for the higher harmonics, almost all of the strong peaks can be assigned to an orbit, showing that the quantum-path analysis has succeeded.

The next section addresses some technical issues which arise during the discrete Fourier transform to control-field sensitivity space. These do not affect the physics and it is possible to skip to section 8.8 without losing much understanding.
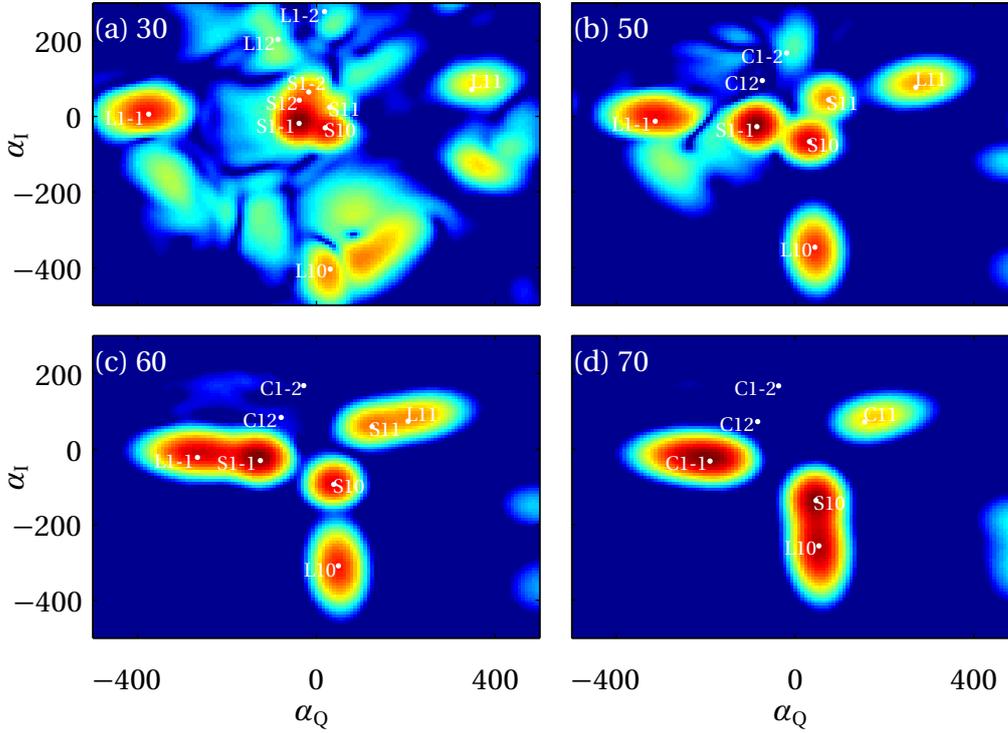
Figure 8.11: Dipole response in control-field sensitivity space i.e. $|\tilde{\mathscr{D}}(\omega;\bar{\alpha})|^2$, the 2D DFT along control-field amplitudes of the dipole response. Harmonics 30, 50, 60 and 70 are shown in subfigures (a)–(d) respectively. The colour scale is normalized and logarithmic over a $10^{-4} - 1$ intensity range. The dipole response has been zero-padded and windowed as described in section 8.7. The predicted control-field sensitivities $\bar{S}_{\omega,C}^{(j)}$ (white dots) are labelled in $\alpha\beta m$ notation.

8.7 Fourier transform issues

8.7.1 Effect of padding

The sampling period of a discrete-Fourier-transformed signal is inversely proportional to the interval over which the signal is defined. This means that by padding a signal with zeros, an arbitrarily fine sampling in the Fourier domain is possible. Such *Fourier-domain oversampling* smooths the discrete Fourier transform, easing visual interpretation and enabling the peaks to be identified easily. It is important to note that despite the apparent improvement in resolution in the Fourier domain, no new information is added. In the discrete Fourier transforms presented above, Fourier-domain oversampling was performed so that the shape and position of the trajectories could be accurately discerned by eye. In the interests of full disclosure, Fig. 8.12 shows the trans-
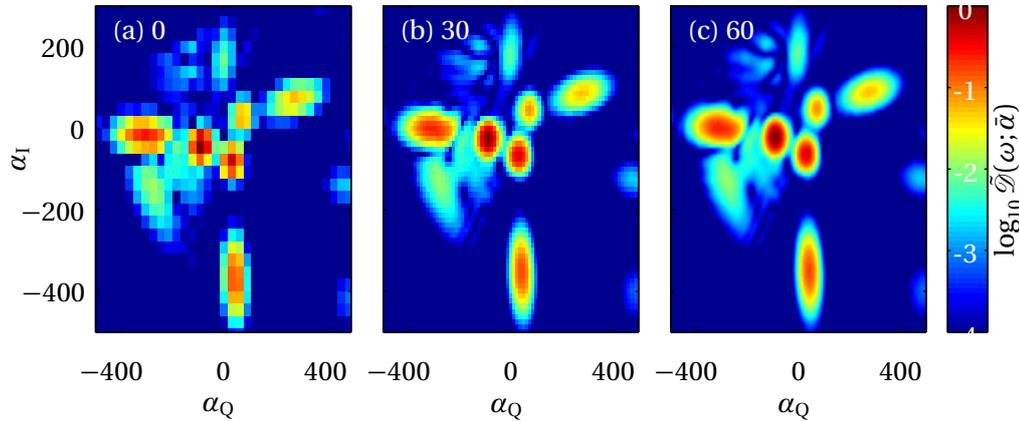
Figure 8.12: Intensity of the two-dimensional discrete Fourier transform, in control field amplitude, of the dipole response, for zero-padding widths (on all four sides) of (a) 0, (b) 30, and (c) 60.

formed data with different amounts of oversampling. I used 60 rows/columns of zeros on all four sides throughout this chapter.

8.7.2 Effect of windowing

The implicitly cyclic boundary conditions of the discrete Fourier transform require a windowing, or apodization function to avoid introducing spurious sidelobes. However, when choosing an apodization function, there is a tradeoff between resolution and dynamic range in the Fourier domain. The colour plots of Fig. 8.11 are shown with an intensity dynamic range of 10^4, an optimistic estimation of what may be experimentally achievable, and the chosen window was a Gaussian with $w = 0.035$ as defined in (8.18). In this section I discuss the choice of window and illustrate the effects of changing it.

The choice of window function is necessarily somewhat subjective. In the present context, a Gaussian has several advantages. i) It is consistent with the formalism of section (8.5). ii) It possesses a width parameter which can be adjusted to locate an optimum choice. iii) The Fourier transform of a Gaussian function of infinite support is itself a Gaussian, and thereby possesses no sidelobes. Any spurious sidelobes introduced by a Gaussian window are due to truncation of its theoretically infinite support. Their amplitudes are therefore directly related to the width w, which may be adjusted to suit a given dynamic range.
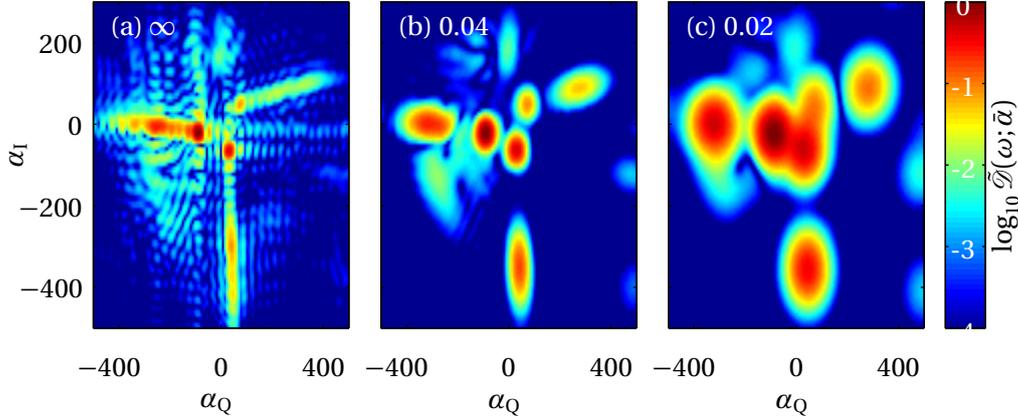
Figure 8.13: Intensity of the two-dimensional discrete Fourier transform, in control field amplitude, of the dipole response, for Gaussian window widths w of (a) ∞, (b) 0.04, and (c) 0.02.

Figure 8.13 illustrates the effect of varying w. For $w = \infty$, the window is solely due to truncation and is effectively a rectangular function. Many spurious sidelobes are visible around the main peaks. At $w = 0.04$ there is some loss of resolution and only a hint of mutually interfering sidelobes around some of the weaker peaks. Finally for $w = 0.02$, no sidelobes are visible but there is a drastic loss of resolution. Throughout this chapter, I used $w = 0.035$.

8.8 Extraction of orbit amplitudes using filtering

Section 8.6 qualitatively showed that the orbits were distinct in the Fourier domain, which I have here denoted "control-field sensitivity space". In this section I extend this by extracting the complex amplitudes of the orbits using Fourier-domain filtering. For each orbit j, one defines a frequency dependent filter function $\mathscr{F}^{(j)}(\omega; \bar{\alpha})$ which isolates the orbit in control-field sensitivity space, in an identical fashion to the Fourier-transform interferometry used elsewhere in this dissertation. The Fourier transform of the total dipole response $\tilde{\mathscr{D}}(\omega; \bar{\alpha})$ is multiplied by this filter to yield the Fourier transform of the response from the selected orbit:

$$\tilde{\mathscr{D}}^{(j)}(\omega; \bar{\alpha}) = \tilde{\mathscr{D}}(\omega; \bar{\alpha}) \mathscr{F}^{(j)}(\omega; \bar{\alpha}). \tag{8.26}$$

Inverse Fourier transforming $\tilde{\mathcal{D}}^{(j)}(\omega; \bar{\alpha})$ gives $\mathcal{D}^{(j)}(\omega; \bar{\theta})$. Evaluating this at $\bar{\theta} = 0$ returns the trajectory amplitude. In fact, using elementary properties of the Fourier transform one has

$$\mathcal{D}^{(j)}(\omega; 0) = \int \tilde{\mathcal{D}}^{(j)}(\omega; \bar{\alpha}) \mathbf{d}\bar{\alpha} \tag{8.27}$$

where the integral is over all of the control-field sensitivity space.

The filters, which depend on the trajectory and the frequency, must be chosen with some care. Ideally, one requires the filter to completely contain all the energy corresponding to its orbit, whilst excluding that of the other orbits. The infinite support of the orbits clearly makes this an impossible requirement to fulfill exactly, and instead one must work down to an effective dynamic range. Here I retain the choice of 10^{-4} (in intensity) used above. If the orbits are separate down to this level, then it is possible to choose suitable filters. However, there are two situations where this will not occur. The first is if orbits are coalescing long and short trajectories. In this case, there is some arbitrariness in the assignment of energy between the two orbits. It may be argued that this is actually a problem of the definition, rather than the method, since as one approaches cutoff the distinction between the long and short trajectories begins to fade. What is *not* arbitrary is the combined amplitude of the coalescing orbits — this is a well-defined quantity that one may "expect" the analysis to accurately retrieve. This condition is satisfied if one chooses the filters such that

$$\mathcal{F}^{(S\beta m)}(\omega; \bar{\alpha}) + \mathcal{F}^{(L\beta m)}(\omega; \bar{\alpha}) = 1, \tag{8.28}$$

i.e. regardless of the arbitrary assignment of energy into long and short trajectories, the total is conserved.

The second situation is if the orbits are not coalescing, but are simply not sufficiently well separated in control-field sensitivity space, perhaps due to improper choice of control field. In this situation, the quantum path interferometry with the chosen dynamic range has failed — one cannot separate the trajectories objectively. This is a limitation of the method.

For nonoverlapping orbits, the filter is a closed curve which encloses the orbit. The filter value
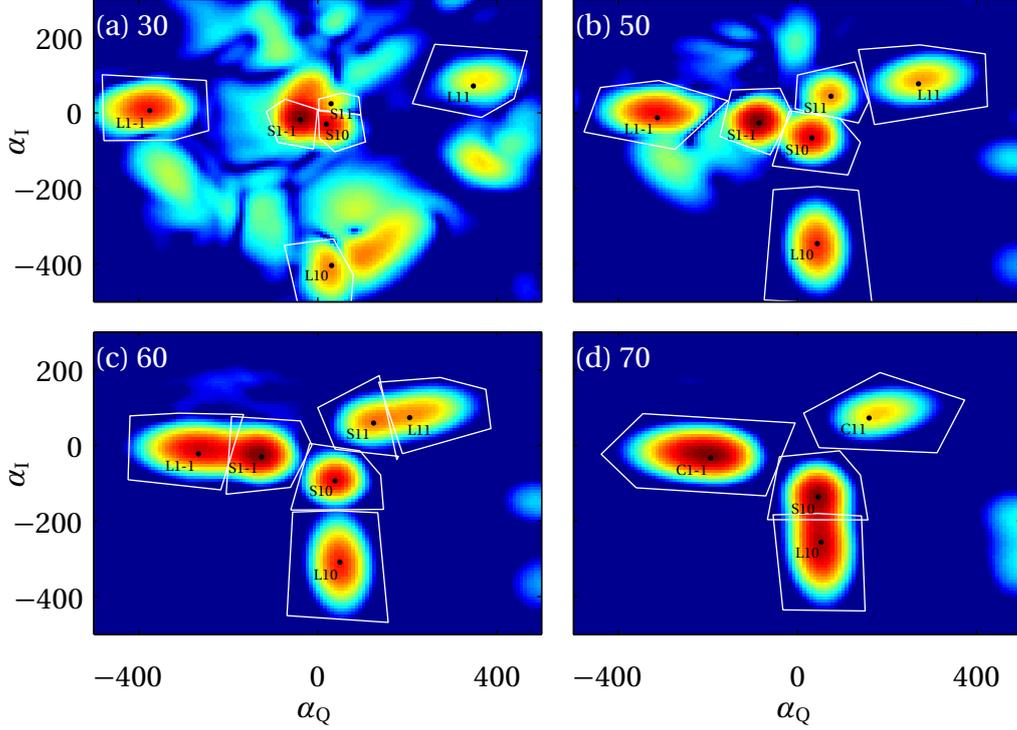
Figure 8.14: Intensity of dipole response in control-field sensitivity space $|\tilde{\mathcal{D}}^{(j)}(\omega;\bar{\alpha})|^2$ showing orbit selection filters (white polygons) and perturbative control-field sensivities $\bar{S}^{(j)}_{\omega,C}$ (black dots) labelled in $\alpha\beta m$ notation. The colour scale is logarithmic with dynamic range 10^4.

is 1 inside the curve and 0 outside. A smooth decay is not needed since the curve traces a line of zero energy. For simplicity, I used polygons, drawn with the mouse using a simple interactive computer program. For overlapping orbits, the polygon is drawn approximately through the saddle which separates them. The filters themselves never overlap — if the polygons overlap then the computer program arbitrarily sets one of the filters to zero in the overlap area. In this way I maintain condition (8.28). Because the orbits evolve gradually with frequency and it is laborious to draw a filter polygon for each frequency, I only drew polygons at every 10th harmonic, and used a simple linear interpolation routine for in-between frequencies. Figure 8.14 shows the filter polygons at four different frequencies, overlaid upon the dipole response.

The filtered orbit amplitudes, obtained using (8.26) and (8.27), are shown in Fig. 8.15. Approaching cutoff, there is a discontinuity as one switches from describing the dipole response as the sum of long and short trajectories to a single dominant trajectory. However below cutoff, the
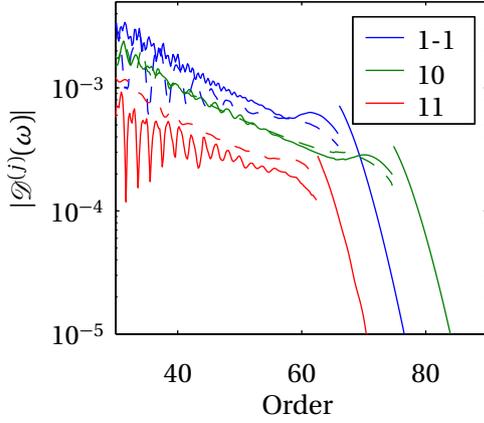
Figure 8.15: Amplitudes $|\mathscr{D}^{(j)}(\omega)|$ of orbits recovered from Fourier-domain filtering. Below cutoff, the short (solid) and long (dashed) orbits are shown, whilst above cutoff the dominant (solid) orbit is shown. The rest of the orbit label βm is denoted by colour as indicated in the legend.
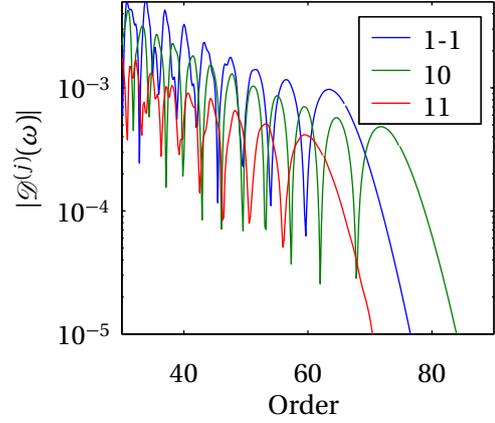
Figure 8.16: Amplitudes of orbits from Fourier-domain filtering. Below cutoff, the short and long trajectories have been combined yielding $|\mathscr{D}^{(S\beta m)}(\omega) + \mathscr{D}^{(L\beta m)}(\omega)|$. Above cutoff, only the dominant orbit is shown. The orbit label βm is denoted by colour as indicated in the legend.

more well-defined quantity is the sum of the long and short trajectories $\mathscr{D}^{(S\beta m)}(\omega) + \mathscr{D}^{(L\beta m)}(\omega)$, which does not depend on the allocation of energy between the two. This is plotted in Fig. 8.16 below cutoff, whilst above cutoff the dominant trajectory amplitude is displayed. The discontinuity is almost completely removed by this process. This shows that the filtering procedure distinguishes between the long and short trajectories where such a distinction is possible; otherwise, it returns the combined response.

8.8.1 Comparison with temporal quantum path analysis

Because the orbits have different emission times, an alternative form of quantum path analysis may be performed by temporally filtering the dipole response $\mathscr{D}(t)$ (with zero control field). Of course, this requires full characterization of the dipole response, information which is likely to be available in computer simulations only, and the ability to perform quantum path analysis without full characterization of the dipole response is one of the major advantages of the method proposed in this chapter. Temporal filtering does, however, offer a straightforward way to verify the accuracy of the proposed method.

I performed such an analysis. First, I computed a set of temporal gates $\mathscr{G}^{(j)}(t)$ based on the real part of recombination times yielded by the quantum orbit calculation. For a pair of short and long orbits $S\beta m$ and $L\beta m$, the gates were

$$\mathscr{G}^{(S\beta m)}(t) = \mathrm{csw}(t; t_{\mathrm{AS}}^{(\beta m)} - t_{\mathrm{w}}, t_{\mathrm{AS}}^{(\beta m)}, t_{\mathrm{s}}) \tag{8.29}$$

$$\mathscr{G}^{(L\beta m)}(t) = \mathrm{csw}(t; t_{\mathrm{AS}}^{(\beta m)}, t_{\mathrm{AS}}^{(\beta m)} + t_{\mathrm{w}}, t_{\mathrm{s}})$$

where $t_{\mathrm{AS}}^{(\beta m)}$ is the real part of the recombination time at the anti-Stokes transition, t_{w} is the window size, and $\mathrm{csw}(t; t_{\mathrm{a}}, t_{\mathrm{b}}, t_{\mathrm{s}})$ is a window with smooth cosine-shaped edges of transition width t_{s} located at t_{a} and t_{b} [1]. Here I take $t_{\mathrm{w}} = 0.23$ optical cycles, and $t_{\mathrm{s}} = 0.09$ optical cycles. Equation (8.29) has the property that in the vicinity of $t_{\mathrm{AS}}^{(\beta m)}$, where the short trajectory transitions into the long, $\mathscr{G}^{(S\beta m)}(t) + \mathscr{G}^{(L\beta m)}(t) = 1$, ensuring that the total energy is conserved, regardless of the vagaries of the actual allocation between the short and long trajectories. Figure 8.17 shows the gates superimposed on a spectrogram of the dipole response.

Each trajectory was then isolated using $\mathscr{D}^{(j)}(t) = \mathscr{D}(t)\mathscr{G}^{(j)}(t)$, and then Fourier transformed to yield $\mathscr{D}^{(j)}(\omega)$. The results are compared with the control-field mediated QPI in Fig. 8.18. In all cases good agreement is observed above the 40th harmonic.

8.9   Summary and outlook

Experimental quantum path analysis promises benefits to the understanding of HHG and other strong field processes and to dynamic and structural imaging using HHG, even when the SFA begins to break down. This chapter proposed control-field mediated quantum-path interferometry, a new technique for experimental quantum path analysis. Using the framework of the SFA, I derived perturbative expressions for the change in the action integral produced by a weak control

---

[1]Mathematically,

$$\mathrm{csw}(t; t_{\mathrm{a}}, t_{\mathrm{b}}, t_{\mathrm{s}}) = \frac{1}{2} \begin{cases} 0 & t < t_{\mathrm{a}} - t_{\mathrm{s}}/2 \\ 1 + \sin(\pi \frac{t - t_{\mathrm{a}}}{t_{\mathrm{s}}}) & t_{\mathrm{a}} - t_{\mathrm{s}}/2 \leq t < t_{\mathrm{a}} + t_{\mathrm{s}}/2 \\ 2 & t_{\mathrm{a}} + t_{\mathrm{s}}/2 \leq t < t_{\mathrm{b}} - t_{\mathrm{s}}/2 \\ 1 - \sin(\pi \frac{t - t_{\mathrm{b}}}{t_{\mathrm{b}}}) & t_{\mathrm{b}} - t_{\mathrm{s}}/2 \leq t < t_{\mathrm{b}} + t_{\mathrm{s}}/2 \\ 0 & t_{\mathrm{b}} + t_{\mathrm{s}}/2 \leq t \end{cases} . \tag{8.30}$$
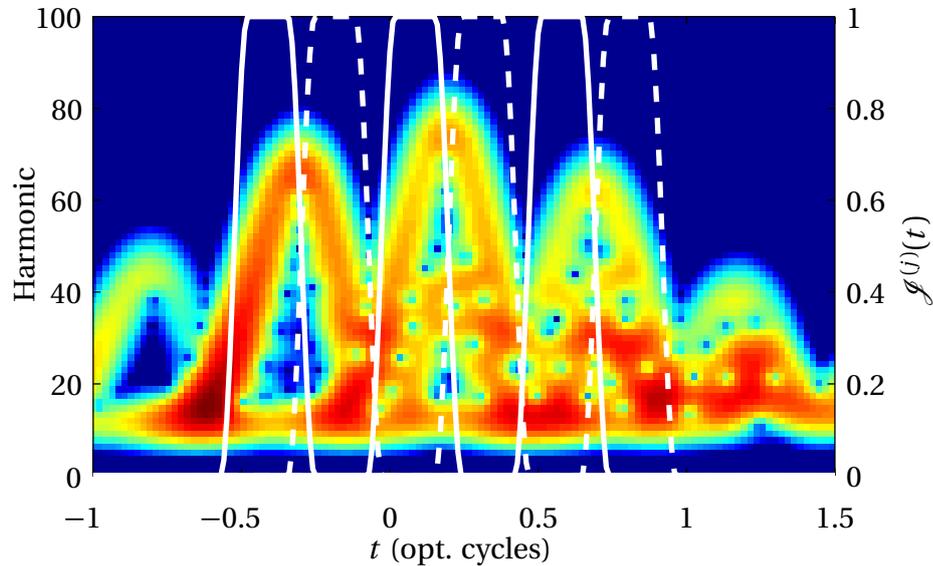
Figure 8.17: Temporal gates $\mathscr{J}^{(j)}(t)$ for quantum orbit analysis superimposed on the spectrogram. Short (long) trajectories are solid (dashed). The colour scale is logarithmic over a range of 10^4 in intensity.
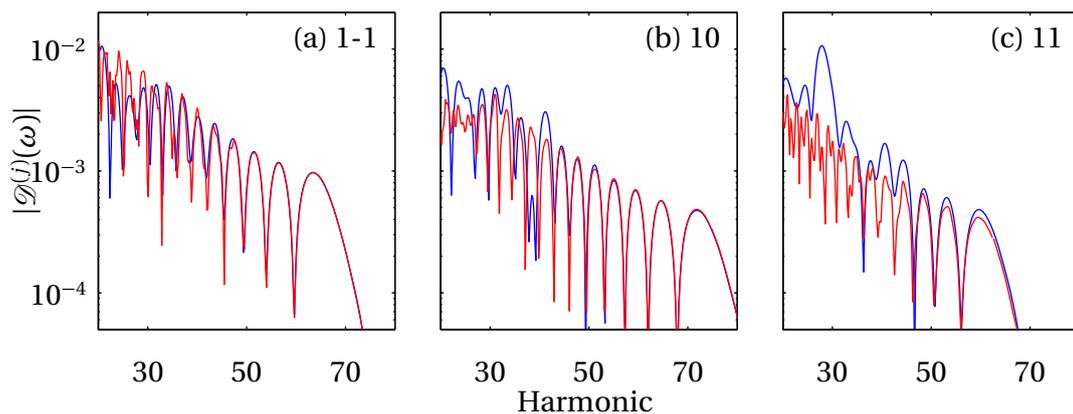
Figure 8.18: Comparison of quantum path amplitudes returned by control-field mediated QPI (red) and temporal filtering (blue) for three different pairs of long/short trajectories (a), (b) and (c). The βm labels of the pairs are given in the plots. For each pair, the sum of the long and short trajectories is shown.

field. Physically, this may be interpreted as a heterodyne between the electron velocities produced by the drive and control field. I examined the range of validity of these expressions and found them adequate for typical experimental situations. I then performed a "numerical experiment", using the TDSE to compute the single-atom response whilst scanning the amplitude and phase of an applied control field. The response of the quantum paths to the control field was in agreement with the perturbative expressions. A Fourier filtering procedure extracted the amplitudes and phases of the quantum paths, which were shown to agree with those obtained directly using temporal filtering of the single-atom response.

The developments of this chapter pose several new questions. Since direct experimental quantum-path analysis has thus far been limited, it remains to be determined what the best use of the new information shall be. For example, one of the motivations mentioned in section 8.1 is the use of multiple paths to help characterize the returning electron wavepacket, an essential step in imaging and dynamical studies. The basic idea is that when comparing different contributions to a fixed frequency, the recombination matrix element is a common factor, so that any differences in the amplitudes are due to the structure of the returning electron wavepacket. The details of this procedure are yet to be determined.

A promising line of research is to search for conditions where higher-order trajectories (i.e. emission from the second return and beyond) may be detected. The perturbative analysis revealed that for control field frequencies $\omega_C/\omega_D \approx 0.7$, the higher-order trajectories lay further from the origin than the first-order trajectories (Figures 8.1 and 8.3 for CW and pulsed drive fields respectively). However, at the control field frequency used in the simulations, $\omega_C/\omega_D = 0.45$, the higher-order trajectories overlap the first-order trajectories. (I chose this frequency because it gave the best spread of first-order trajectories to assist with what was an initial proof-of-principle calculation.) At the lower harmonics (e.g. the 30[th] harmonic shown in Fig. 8.11(a)), many peaks appeared in the Fourier domain, and it is possible that some of them belonged to higher-order trajectories. The ability to detect such trajectories would be significant because experimental evidence to date has been limited.

The most important next step is the development of a plausible experimental implementation.

For the phase-sensitive detection, one could use an interferometer similar to that of chapter 7, with the control field inserted into one arm using a dichroic mirror. The control field itself could be produced using an optical parametric amplifier (OPA). Current designs based on the nonlinear crystal BBO are strong candidates [433–441, 441–443]. The amplitude and phase of the control field could be scanned in a variety of ways, with the most direct being a half-waveplate/polarizer variable attenuator and a rotating thin glass plate. Shutters could rapidly block and unblock the reference arm of the interferometer, enabling interferograms as well as intensity-only images to be acquired. Shutters could also repeatedly block the control field, providing control-field-free images which would provide the interferometer reference phase and thus keep track of slow drifts during the scan. It would be necessary to eliminate phase-matching effects and access, as closely as possible, the single-atom response. A thin gas jet, combined with a flat-top beam profile [444, 445] would achieve this end.

Another possibility is noninterferometric experimental implementation, in which only the intensity of the harmonics from the drive+control field is obtained. In principle, the filtering step should still be possible, since beating between the quantum paths will appear as sidebands in the Fourier domain. However, with n quantum paths there are $n(n-1)$ sidebands in the Fourier domain (self-terms are at the origin), which may complicate the identification and isolation steps.

It is also worthwhile to consider more complex control fields, which may offer better quantum-path selectivity. One possibility is a much shorter control field, which could select only a single half-cycle of the drive field.

215

9 Summary and outlook

This chapter summarises the achievements of this dissertation and describes potential future research.

Multiple shear spectral interferometry. Chapter 3 developed the theory of multiple-spectral-shearing interferometry. This is a generalisation of the concatenation algorithm for spectral shearing interferometry (SSI) which combines data taken at different spectral shears. The benefits are many: improved precision, consistency checking and estimation of the magnitude of systematic errors. The method also partially overcomes the relative phase ambiguity for spectrally disjoint subpulses. In single-shear SSI, even a small spectral null causes a relative phase ambiguity. In multiple-shear SSI, one may "cross" spectral gaps up to the size of the largest contiguous spectral lobe. This is the upper limit for all self-referenced methods based on sum or difference frequency generation [81].

Future theoretical work may include extending the capabilities of the algorithm to deal with noninteger shear multiples, and working in two or more dimensions. Perhaps the most pressing experimental application is characterising highly structured spectra from nonlinear pulse compression systems pushed to their limits. In general, if a variable shearing interferometer is available then there is no reason for not performing a multiple shear reconstruction. I hope that this becomes part of "good practice" in the future.

Experimental demonstration of multishear spectral interferometry. Chapter 4 presented two experimental implementations of multiple-shear spectral interferometry. In the first, the multiple shears were taken sequentially using a spatially encoded arrangement SPIDER (SEA-SPIDER), whilst in the latter they were taken simultaneously using a spatially encoded arrangement CAR-SPIDER (SEA-CAR-SPIDER). Both improved precision and resolution of the relative phase ambiguity were demonstrated.

A promising future research path is simplifying and miniaturising spectral shearing interferometers capable of acquiring multiple shears. The recent commercial release[1] of a two-dimensional spectral shearing interferometer capable of acquiring variable shears [446] shows that robust im-

---

[1]http://www.idestaqe.com/products/2dsi/

plementations are possible. This motivates a search for a compact single-shot arrangement. The original chirped-arrangement for SPIDER (CAR-SPIDER) demonstration made use of a long crystal arrangement, but nonetheless required a noncollinear Mach-Zehnder interferometer (MZI) and an imaging spectrometer. Perhaps these components may be done away with using some additional ingenuity.

Compact space-time SPIDER. Chapter 5 presented an interferometer capable of simultaneously recording both spatial and spectral phase gradients, thus enabling spatio-temporal ultrashort pulse measurement. The interferometer itself is actually very similar to that of a conventional spectral phase interferometry for direct electric-field reconstruction (SPIDER) device, the main addition being an imaging spectrometer. The new capability of measuring the spatial phase gradient was primarily due to careful interpretation of lateral shearing interferometry (LSI) with a lens between the shearing element and the detector. The device does require careful calibration and therefore I developed a precise method and gave a thorough description.

One limitation of this method is the high spectral resolution requirement of temporally encoded SPIDER. Unfortunately, a spatially encoded version is difficult to imagine, because SEA-SPIDER has a completely different design — two ancillae are used and the unknown pulse is not split. Instead, one could consider augmenting SEA-SPIDER designs to measure the spatial phase gradient as well, perhaps using the fundamental beam after upconversion in the crystal.

Another issue is the nonintuitive nature of the acquired trace, which requires Fourier filtering and carrier subtraction before being interpretable to the eye. Specially-written software is therefore required for on-line measurements.

With these caveats, the new device may be used wherever a conventional SPIDER based on a Mach-Zehnder interferometer is used. One useful application may be in monitoring the output from chirped-pulse amplifier (CPA) systems, which often feature smoothly varying spatio-temporal distortions due to the Kerr nonlinearity and misaligned stretchers and compressors.

Lateral shearing interferometry for high-harmonic generation. Chapter 7 presented the first systematic demonstration of spatially resolved wavefront sensing using LSI for high-harmonic generation (HHG). The advantages of the method is that all manipulation is performed before

the vacuum chamber on the laser pulses, and that it requires only a few acquisitions (one plus the number of shears). I derived a simple interpretation for the method: it returns the phase of the harmonics at the detector plane minus the spherical phase resulting from their propagation from the focus, and in the far-field this may be interpreted as the spatial Fourier transform of the harmonics in the focal plane of the focusing optic. I presented a comprehensive set of results, in good agreement with theory.

The requirements are that sufficient laser power must be available to divide the pulse energy four ways and the spectrometer must have sufficient resolution to resolve the fringes produced by separate harmonic sources. With the wide deployment of commercial CPAs and the maturing of extended ultraviolet (XUV) metrology, both these requirements are likely to be fulfilled in many laboratories. Therefore, wavefront sensing can enter the set of standard diagnostics available for probing HHG and will be a useful asset for the exploration of new regimes, such as long-wavelength drive lasers and molecular targets.

Quantum-path interferometry for high-harmonic generation. Chapter 8 showed how a weak control field could be used to distinguish between quantum paths that contribute to the single-atom response in HHG. By judicious choice of control field, the phase shift imparted on each quantum path depends differently on the amplitude and phase of the control field. By scanning these parameters and recording the harmonic emission in a phase-sensitive fashion, the complex amplitude of each quantum path may be determined. The method has a simple interpretation as multi-dimensional Fourier-transform interferometry. I presented a numerical experiment which illustrated the principle.

Many groups are currently developing multi-wavelength sources with a view to both probing and controlling HHG. Quantum-path interferometry will be a useful framework for understanding mixing processes in the case of a strong drive field and a weak control field. Potential highlights of future research include the possibility of observing high-order trajectories and extending the time-range of chirped-encoded recollision experiments such as probing attosecond dynamics by chirp encoded recollision (PACER) [363].

A   Noise and uncertainty

This appendix derives several results concerning noise and uncertainty in interferometry and pulse characterisation.

A.1   Correlations of filtered additive white noise

This section derives some properties of additive detector noise after the Fourier-domain filtering algorithm used in interferometry. I shall treat one-dimensional spectral interferometry for concreteness. The results may be easily adapted to other domains and/or higher dimensions. The additive real detector noise $\zeta(\omega)$ is filtered to produce complex noise $\xi(\omega)$, which contaminates the true value of the interferometric product $\bar{D}(\omega) = E_2(\omega)E_1^*(\omega)$. One therefore writes $D(\omega) = \bar{D}(\omega) + \xi(\omega)$, where $D(\omega)$ is the value of the interferometric product obtained by Fourier filtering. I shall determine the two-frequency correlations $\langle \xi(\omega_1)\xi^*(\omega_2) \rangle$ and $\langle \xi(\omega_1)\xi(\omega_2) \rangle$. These will be used in the next section to derive the statistics of the fluctuations in the extracted phase $\Gamma(\omega) = \arg D(\omega)$.

One writes $\xi(\omega)$ as being the result of Fourier filtering the detector noise

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

where $\tilde{\zeta}(t)$ is the Fourier transform of the detector noise and $F(t)$ is the response of the filter in the quasi-time domain. Then

$$\xi(\omega_1)\xi(\omega_2) = \frac{1}{2\pi} \int_{-\infty}^{\infty} dt \int_{-\infty}^{\infty} dt'\, \tilde{\zeta}(t)\tilde{\zeta}(t')F(t)F(t')e^{i\omega_1 t + i\omega_2 t'}. \tag{A.2}$$

One now applies the expectation brackets and uses $\tilde{\zeta}(t) = \tilde{\zeta}^*(-t)$ since $\zeta(\omega)$ is real, yielding

$$\langle \xi(\omega_1)\xi(\omega_2) \rangle = \frac{1}{2\pi} \int_{-\infty}^{\infty} dt \int_{-\infty}^{\infty} dt'\, \langle \tilde{\zeta}(t)\tilde{\zeta}^*(-t') \rangle F(t)F(t')e^{i\omega_1 t + i\omega_2 t'}. \tag{A.3}$$

Assuming that the noise intensity in the frequency domain is uniform, the two-time correlation is a delta function. I also assume the detector noise is white i.e. it is uncorrelated in the frequency

domain, so that its Fourier-domain amplitude is uniform, and so one can write

$$\langle \tilde{\zeta}(t)\tilde{\zeta}^*(t') \rangle = |\tilde{\zeta}|^2 \delta(t - t'). \tag{A.4}$$

Substitution into (A.3) gives

$$\langle \xi(\omega_1)\xi(\omega_2) \rangle = \frac{|\tilde{\zeta}|^2}{2\pi} \int_{-\infty}^{\infty} \mathrm{d}t\, F(t)F(-t)e^{i(\omega_1 - \omega_2)t}. \tag{A.5}$$

Since the filter isolates the positive or negative sideband in the quasi-time domain, it satisfies the property $F(t)F(-t) = 0$. Therefore $\langle \xi(\omega_1)\xi(\omega_2) \rangle = 0$.

The other two-frequency correlation $\langle \xi(\omega_1)\xi^*(\omega_2) \rangle$ can be written using (A.1) and (A.4) as

$$\langle \xi(\omega_1)\xi^*(\omega_2) \rangle = \frac{\langle |\tilde{\zeta}|^2 \rangle}{2\pi} \int_{-\infty}^{\infty} \mathrm{d}t\, |F(t)|^2 e^{i(\omega_1 - \omega_2)t}. \tag{A.6}$$

The correlation between frequencies ω_1 and ω_2 is therefore proportional to the Fourier transform of the intensity of the filter response at frequency $\omega_2 - \omega_1$.

In SSI, the trace is resampled with spacing Ω, and so the correlation between frequencies separated by integer multiples of Ω is of interest. The correlation turns out to be zero for a rectangular filter with passband $2T = 4\pi/\Omega$, and will remain zero if this rectangular filter's response is smoothed by a convolution. Other well-designed filter profiles will produce a small but not identically zero correlation. Equation (A.6) is actually a standard result for the cross-spectral density of a random process subject to a linear system [447, 448].

Ignoring correlations in (A.6), the filtered noise is completely characterized by its variance

$$\langle |\xi|^2 \rangle = \frac{\langle |\tilde{\zeta}|^2 \rangle}{2\pi} A \tag{A.7}$$

where $A = \int_{-\infty}^{\infty} \mathrm{d}t\, |F(t)|^2$ is the "area" of the filter.

The extension of these arguments to a two-dimensional data trace is straightforward, relying only on the noise being real-valued and on the filter transfer function having a property analogous

to the aforementioned condition for the one-dimensional case.

A.2   Estimation of the filtered noise amplitude

This section explains a convenient means of estimating $\langle|\xi|^2\rangle$ without applying (A.7) directly. This is useful because estimating $\langle|\zeta|^2\rangle$ straight from the raw data may be difficult if the detector is completely filled by the signal. However, the signal rarely occupies the entirety of the Fourier domain. Therefore, one may use choose a "noise filter" which picks out some region of the Fourier domain which has no significant contribution from the signal. The intensity of filtered noise is then integrated, providing an estimate of $\langle|\tilde{\zeta}|^2\rangle A_{\mathrm{n}}/(2\pi)$ where A_{n} is the area under the noise filter. Multiplying this by A/A_{n} provides an estimate of $\langle|\xi|^2\rangle$ provided that the noise is uniform in the Fourier domain.

A.3   Amplitude of the phase fluctuations

Here, the aim is to determine the statistics of the extracted phase $\Gamma(\omega) = \arg D(\omega)$. Taylor expansion of the complex logarithm gives

$$
\begin{aligned}
\Gamma(\omega) &= \arg\left[\bar{D}(\omega) + \xi(\omega)\right] & \text{(A.8)}\\
&= \arg\bar{D}(\omega) + \mathrm{Im}[\frac{\xi(\omega)}{\bar{D}(\omega)}] + O\left(\left[\frac{\xi(\omega)}{\bar{D}(\omega)}\right]^2\right). & \text{(A.9)}
\end{aligned}
$$

Note that the lower-case arg denotes the local analytic continuation, so that branch cuts of the argument function are ignored. I henceforth work to first order in $\frac{\xi(\omega)}{\bar{D}(\omega)}$. I shall obtain the two-frequency correlation

$$
\langle\Delta\Gamma(\omega_1)\Delta\Gamma(\omega_2)\rangle = \left\langle \mathrm{Im}[\frac{\xi(\omega_1)}{\bar{D}(\omega_1)}]\,\mathrm{Im}[\frac{\xi(\omega_2)}{\bar{D}(\omega_2)}]\right\rangle. \tag{A.10}
$$

Throughout this section, Δ denotes departure from the expected value, and angle brackets denote averaging over an ensemble of measurements, each identical except for random detector noise.

223

Using the identity $\text{Im}(z_1)\text{Im}(z_2) = \frac{1}{2}\text{Re}\left[z_1 z_2^* - z_1 z_2\right]$, one writes

$$\langle\Delta\Gamma(\omega_1)\Delta\Gamma(\omega_2)\rangle = \frac{1}{2}\text{Re}\left(\frac{\langle\xi(\omega_1)\xi^*(\omega_2)\rangle}{|\bar{D}(\omega_1)\bar{D}(\omega_2)|} - \frac{\langle\xi(\omega_1)\xi(\omega_2)\rangle}{\bar{D}(\omega_1)\bar{D}(\omega_2)}\right) \tag{A.11}$$

factoring $\bar{D}(\omega)$ out of the expectation brackets because it is constant. From section A.1 below, the second term is zero and for a choice of filter appropriate for SSI, correlations are small. The noise intensity is also assumed to be frequency-independent. One therefore has

$$\sigma_\Gamma^2(\omega) = \frac{\langle|\xi|^2\rangle}{2|\bar{D}(\omega)|^2}. \tag{A.12}$$

Equation (A.12) also assumes that $\langle|\xi|^2\rangle \ll |\bar{D}(\omega)|$. Although this is false around spectral nulls, it still leads to an appropriate de-weighting which is the main purpose of this formalism.

A.4   Minimizing root-mean-square variation in a set of fields in the presence of trivial ambiguities

It is often necessary, when evaluating the variation of an ensemble of pulse reconstructions, to remove zeroth and first order phase (some of the "trivial ambiguities") so that only fluctuations of higher order would affect the statistics. One method of accomplishing this is by minimizing the root-mean-square (RMS) field variation over the ensemble [283, 284]

$$\epsilon = \left\langle\int_{-\infty}^{\infty}|\Delta E(t)|^2 \, dt\right\rangle^{1/2}. \tag{A.13}$$

Here, Δ denotes fluctuation about the mean value. The norm may also be equivalently defined via Parseval's theorem in the frequency domain. A procedure for assigning the zeroth- and first-order phase to minimize ϵ for a pair of pulses is given in Appendix A in ref. [283]. To generalize this procedure to a set of pulses, I used an iterative algorithm. Each iteration consists of two steps. In the first step, the average pulse $\langle E(t)\rangle$ is computed. In the second step, the procedure of ref. [283] is individually applied between each pulse and the average pulse. Generally, only a few iterations were required for convergence.

B Simulation codes for high-harmonic generation

Chapters 6–8 make use of several codes for simulating HHG. I wrote all of them in MATLAB. This appendix provides some details on each.

B.1 Time-dependent Schrödinger equation

The time-dependent Schrödinger equation (TDSE) is analytically intractable for HHG, neccessitating numerical methods. One common approach is the split-operator method [375], which alternates between the spatial domain and the momentum domain using appropriate transforms. The potential and electric field components of the length-gauge Hamiltonian are multiplicative in the spatial domain, whilst the momentum operator is multiplicative in the momentum domain. The calculation is often performed with reduced spatial dimensionality. For example, in one spatial dimension, taken parallel to the electric field polarisation, a Fourier transform links the two domains.

My TDSE solver treated one spatial dimensional and used the split-operator method. One iteration consists of the application of the core and laser potential, a Fourier transform to the momentum domain, application of the kinetic energy operator, and an inverse Fourier transform. A linear absorbing boundary located beyond the ponderomotive motion amplitude prevented back-reflections. For convenience and efficiency, I used an adaptive step-sizing method: in each iteration, the result of two half steps is compared with one full step. If the difference exceeds a certain tolerance, then the step size is reduced and the iteration repeated. Otherwise, the step size is lengthened slightly and the simulation advances. Typical HHG simulations for a 7 fs, 800 nm 5×10^{14} W/cm^2 pulse took 15 minutes on my laptop computer. In all calculations, I used the smoothed "soft-core" potentials $V(x) = -1/\sqrt{x^2 + a^2}$. The parameter a is chosen so that the ionization potential matches that of the target atom. I obtained the ground-state wavefunctions by expressing the time-independent Schrödinger equation as a matrix and solving the resulting eigenvalue problem using MATLAB's linear algebra routines.

B.2 Classical recollision solver

The classical recollision solver finds all solutions (t_r, t_b) to the recollision equation i.e. the roots of (6.1). Its inputs are the vector potential $A(t)$ and a set of initial guesses of (t_r, t_b) pairs, one for each half-cycle in which an electron "birth" is required. As guesses, I use the cutoff birth and return times for a sinusoidal field — this is close enough even for few-cycle pulses. The algorithm moves through a range of return times in regular steps of size Δt_r. At each step, the birth time is computed using the MATLAB `fzero` function with the previously calculated birth time as a guess. Starting from the initial guess, close to the cutoff, the algorithm steps backwards in return time, moving through the short trajectories. It stops when the computed birth time equals the return time. It then returns to the initial guess but steps forward in return time, passing through the long trajectories and then on to the multiple-return trajectories. It stops when a given maximum excursion time is reached.

B.3 Quantum-orbits model

The quantum orbits model solves for stationary points (t_r, t_b) of the action i.e. (6.17) and (6.18). (The stationary momentum is easily calculated using (6.16) and not therefore not treated as an unknown.) Its inputs are the vector potential $A(t)$, a set of initial guesses at a given emission frequency ω and the ionization potential. Each guess corresponds to one trajectory $\alpha\beta m$. The algorithm steps through the desired range of emission frequencies, with step size $\Delta\omega$. At each frequency (t_r, t_b) are found using Newton's method. A linear extrapolation of the previous two solutions is used for the initial guess at each frequency.

Near the cutoff, the solutions of the long and short trajectories become very close, and the Newton's method solver tends to jump between them. I overcame this with an adaptive step-size method which takes smaller steps when the solutions become close. When moving from frequency ω to $\omega + \Delta\omega$, the solutions obtained using Newton's method at $\omega + \Delta\omega$ are compared to the solutions at ω. If the solutions have moved by too great a distance, they are considered invalid, $\Delta\omega$ is reduced and the procedure is repeated. The guess is also modified by a small random

amount to prevent the algorithm becoming "stuck". If the solutions are valid, then the step is taken and the step-size is increased slightly. The maximum permitted distance for a valid step is half the distance between the solutions, or some user-determined distance — whichever is the smaller. When scanning over a range of laser intensities, I use a similar adaptive method to control the intensity step-size ΔI.

B.4   Propagation model

I wrote a propagation code to simulate the spatial profiles observed in chapter 7. The model was an approximate one in which the only variations considered in the laser pulse are those of intensity and arrival time. Since the single-atom response $E_S(I, q)$ is only affected trivially by the arrival time, it only needs to be calculated over the two-dimensional parameter range (I, q), where q is the harmonic order. The laser beam is defined at its centre frequency, and can be assigned any spatial intensity profile $I_L(r, z)$ and phase profile $\phi_L(r, z)$. Cylindrical symmetry is assumed throughout. The phase of the laser profile is treated as a time delay $\phi_L(r, z)/\omega_L$ on the pulse, resulting in a phase $q\phi_L(r, z)$ in the q-th harmonic. This approximation is equivalent to ignoring variations in the carrier-envelope phase of the pulse, and is valid for many-cycle pulses.

The model includes the refractive index contributions of the gas at both the laser and harmonic wavelengths, and the absorption of the harmonics. The gas density is low enough that its absorption and refractive index contribution are proportional to the local pressure $g(z)$. The harmonics therefore experience complex-valued wavenumber $\beta_q g(z)$, where $\beta_q = (n_q - 1 + i\alpha_q/2)q k_L$ which includes both the refractive index n_q and absorption α_q. It is convenient to define these values at one atmosphere, in which case the units of $g(z)$ are atmospheres. Likewise, the laser has wavenumber $\beta_L g(z)$ where $\beta_L = (n_L - 1)k_L$. Putting this all together, the Hankel transform of the macroscopic response $E_M(k_T, L, \omega)$ is written (neglecting constant factors) as

$$E_M(k_T, L, \omega) = \int_{-\infty}^{\infty} e^{-i\frac{k_T^2}{2qk_L}(L-z') + i\left[\bar{g}(\infty) - \bar{g}(z')\right]\beta_q} \, \mathrm{HT}\left\{E_S\left[\tilde{I}_L(r, z), q\right] e^{iq[\beta_L \bar{g}(z') + \phi_L(r, z)]}\right\} g(z') \, \mathrm{d}z' \quad \text{(B.1)}$$

where $\bar{g}$ is the "effective cumulative pressure-thickness" of the gas $\bar{g}(z) = \int_{-\infty}^{z} g(z') \, \mathrm{d}z'$.

References

[1] Rakov, V. & Uman, M. *Lightning: physics and effects.* Cambridge Univ Pr (2003). 6

[2] Shah, J. *Ultrafast spectroscopy of semiconductors and semiconductor nanostructures.* Springer Verlag (1999). 7

[3] Jimenez, R., Fleming, G. R., Kumar, P. V., & Maroncelli, M. *Nature* **369**, 471 (1994). 7

[4] Zewail, A. H. *J. Phys. Chem. A* **104**, 5660 (2000). 7

[5] Mukamel, S. *Principles of nonlinear optical spectroscopy.* Oxford University Press, New York (1995). 7

[6] Dantus, M. *Annu. Rev. Phys. Chem.* **52**, 639 (2001). 7

[7] Warren, W., Rabitz, H., & Dahleh, M. *Science* **259**, 1581 (1993). 7, 11, 64

[8] Bonadeo, N., Erland, J., Gammon, D., Park, D., Katzer, D., & Steel, D. *Science* **282**, 1473 (1998). 7, 11

[9] Zipfel, W., Williams, R., & Webb, W. *Nat. Biotechnol.* **21**, 1369 (2003). 7

[10] Xu, C., Zipfel, W., Shear, J. B., Williams, R. M., & Webb, W. W. *Proc. Natl. Acad. Sci. U. S. A.* **93**, 10763 (1996). 7

[11] Gattass, R. R. & Mazur, E. *Nat. Photon.* **2**, 219 (2008). 7

[12] Ye, J. & Cundiff, S. *Femtosecond optical frequency comb technology: Principle, operation and application.* Springer Verlag (2005). 8

[13] Udem, T., Holzwarth, R., & Haensch, T. *Nature* **416**, 233 (2002). 8

[14] Holzwarth, R., Udem, T., Hansch, T. W., Knight, J. C., Wadsworth, W. J., & Russell, P. St. J.. *Phys. Rev. Lett.* **85**, 2264 (2000). 8

[15] Bradley, D. & New, G. *Proc. IEEE* **62**, 313 (1974). 8

[16] Korobkin, V. V., Malyutin, A. A., & Shchelev, M. Y. *Sov. Phys. JETP* **11**, 103 (1970). 8

[17] Bradley, D. J., Liddy, B., & Sleat, W. E. *Opt. Commun.* **2**, 391 (1971). 8

[18] Bradley, D., Liddy, B., Roddie, A., Sibbett, W., & Sleat, W. In *Photo-Electronic Image Devices, Proceedings of the Fifth Symposium Held at Imperial College*, vol. 33, Part 2 of *Advances in Electronics and Electron Physics*, pages 1145 – 1156. Academic Press (1972). 9

[19] Ruddock, I. S. & Bradley, D. J. *Appl. Phys. Lett.* **29**, 296 (1976). 9

[20] Armstrong, J. A. *Appl. Phys. Lett.* **10**, 16 (1967). 9, 32

[21] Weber, H. P. *J. Appl. Phys.* **39**, 6041 (1968). 9

[22] Haus, H. *IEEE J. Quantum. Elect.* **11**, 736 (1975). 9

[23] Wise, F., Walmsley, I., & Tang, C. *Opt. Lett.* **13**, 129 (1988). 9

[24] Treacy, E. B. *Phys. Lett. A* **28**, 34 (1968). 9

[25] Treacy, E. B. *Appl. Phys. Lett.* **14**, 112 (1969). 9, 34

[26] Cubeddu, R. & Svelto, O. *IEEE J. Quantum. Elect.* **5**, 495 (1969). 9

[27] Svelto, O. *Appl. Phys. Lett.* **17**, 83 (1970). 9

[28] Ippen, E. P. & Shank, C. V. *Appl. Phys. Lett.* **27**, 488 (1975). 9

[29] Fork, R. L., Greene, B. I., & Shank, C. V. *Appl. Phys. Lett.* **38**, 671 (1981). 9

[30] Fork, R., Shank, C., Yen, R., & Hirlimann, C. *IEEE J. Quantum. Elect.* **19**, 500 (1983). 9

[31] Dietel, W., Döpel, E., Kühlke, D., & Wilhelmi, B. *Opt. Commun.* **43**, 433 (1982). 10

[32] Mourou, G. & II, T. S. *Opt. Commun.* **41**, 47 (1982). 10

[33] Dietel, W., Fontaine, J. J., & Diels, J.-C. *Opt. Lett.* **8**, 4 (1983). 10

[34] Fontaine, J., Dietel, W., & Diels, J.-C. *IEEE J. Quantum. Elect.* **19**, 1467 (1983). 10

[35] Valdmanis, J. A., Fork, R. L., & Gordon, J. P. *Opt. Lett.* **10**, 131 (1985). 10

[36] Valdmanis, J. & Fork, R. *IEEE J. Quantum. Elect.* **22**, 112 (1986). 10

[37] Brabec, T. & Krausz, F. *Rev. Mod. Phys.* **72**, 545 (2000). 10

[38] Asaki, M., Huang, C., Garvey, D., Zhou, J., Kapteyn, H., & Murnane, M. *Opt. Lett.* **18**, 977 (1993). 10

[39] Curley, P., Spielmann, C., Brabec, T., Krausz, F., Wintner, E., & Schmidt, A. *Opt. Lett.* **18**, 54 (1993). 10

[40] Proctor, B. & Wise, F. *Appl. Phys. Lett.* **62**, 470 (1993). 10

[41] Zhou, J., Taft, G., Huang, C., Murnane, M., Kapteyn, H., & Christov, I. *Opt. Lett.* **19**, 1149 (1994). 10

[42] Fork, R. L., Cruz, C. H. B., Becker, P. C., & Shank, C. V. *Opt. Lett.* **12**, 483 (1987). 10, 11

[43] Spielmann, C., Curley, P., Brabec, T., & Krausz, F. *IEEE J. Quantum. Elect.* **30**, 1100 (1994). 10

[44] Stingl, A., Lenzner, M., Spielmann, C., Krausz, F., & Szipöcs, R. *Opt. Lett.* **20**, 602 (1995). 10

[45] Trebino, R. & Kane, D. J. *J. Opt. Soc. Am. A* **10**, 1101 (1993). 10, 38, 40

[46] Iaconis, C. & Walmsley, I. A. *Opt. Lett.* **23**, 792 (1998). 10, 44

[47] Taft, G., Rundquist, A., Murnane, M., Kapteyn, H., DeLong, K., Trebino, R., & Christov, I. *Opt. Lett.* **20**, 743 (1995). 10

[48] Kasper, A. & Witte, K. J. *J. Opt. Soc. Am. B* **15**, 2490 (1998). 10

[49] Dudley, J. M., Boussen, S. M., Cameron, D. M. J., & Harvey, J. D. *Appl. Opt.* **38**, 3308 (1999). 10, 39

[50] Gallmann, L., Sutter, D. H., Matuschek, N., Steinmeyer, G., Keller, U., Iaconis, C., & Walmsley, I. A. *Opt. Lett.* **24**, 1314 (1999). 10, 44, 45

[51] Sutter, D. H., Steinmeyer, G., Gallmann, L., Matuschek, N., Morier-Genoud, F., Keller, U., Scheuer, V., Angelow, G., & Tschudi, T. *Opt. Lett.* **24**, 631 (1999). 10

[52] Morgner, U., Kärtner, F. X., Cho, S. H., Chen, Y., Haus, H. A., Fujimoto, J. G., Ippen, E. P., Scheuer, V., Angelow, G., & Tschudi, T. *Opt. Lett.* **24**, 411 (1999). 10

[53] Strickland, D. & Mourou, G. *Opt. Commun.* **56**, 219 (1985). 10

[54] Backus, S., Durfee III, C. G., Murnane, M. M., & Kapteyn, H. C. *Rev. Sci. Instrum.* **69**, 1207 (1998). 10

[55] Kohler, B., Yakovlev, V. V., Wilson, K. R., Squier, J., DeLong, K. W., & Trebino, R. *Opt. Lett.* **20**, 483 (1995). 11

[56] Dorrer, C., de Beauvoir, B., Le Blanc, C., Ranc, S., Rousseau, J., Rousseau, P., Chambaret, J., & Salin, F. *Opt. Lett.* **24**, 1644 (1999). 11, 44

[57] Dorrer, C. & Walmsley, I. *Opt. Lett.* **27**, 1947 (2002). 11, 61

[58] Dorrer, C., Kosik, E., & Walmsley, I. *Appl. Phys. B: Lasers Opt.* **74**, 209 (2002). 11, 19, 57, 61, 105, 107, 108

[59] Dorrer, C., Kosik, E. M., & Walmsley, I. A. *Opt. Lett.* **27**, 280 (2002). 11, 57, 61, 105, 107, 108

[60] Nisoli, M., Silvestri, S. D., & Svelto, O. *Appl. Phys. Lett.* **68**, 2793 (1996). 11, 65, 104, 165

[61] Baltuška, A., Pshenichnikov, M. S., & Wiersma, D. A. *Opt. Lett.* **23**, 1474 (1998). 11, 39

[62] Cheng, Z., Fürbach, A., Sartania, S., Lenzner, M., Spielmann, C., & Krausz, F. *Opt. Lett.* **24**,

247 (1999). 11

[63] Stibenz, G., Zhavoronkov, N., & Steinmeyer, G. *Opt. Lett.* **31**, 274 (2006). 11

[64] Gu, X., Xu, L., Kimmel, M., Zeek, E., O'Shea, P., Shreenath, A., Trebino, R., & Windeler, R. *Opt. Lett.* **27**, 1174 (2002). 11, 40, 66

[65] Diddams, S. A., Eaton, H. K., Zozulya, A. A., & Clement, T. S. *Opt. Lett.* **23**, 379 (1998). 11

[66] Zaïr, A., Guandalini, A., Schapper, F., Holler, M., Biegert, J., Gallmann, L., Keller, U., Couairon, A., Franco, M., & Mysyrowicz, A. *Opt. Express* **15**, 5394 (2007). 11, 57

[67] Diddams, S., Eaton, H., Zozulya, A., & Clement, T. In *Lasers and Electro-Optics*, pages 519– (1998). 11, 59

[68] Solli, D., Ropers, C., Koonath, P., & Jalali, B. *Nature* **450**, 1054 (2007). 11

[69] Frei, F., Galler, A., & Feurer, T. *J. Chem. Phys.* **130**, 034302 (2009). 11

[70] Sussman, B. J., Lausten, R., & Stolow, A. *Phys. Rev. A* **77**, 043416 (2008). 11

[71] Tanabe, T., Tanabe, H., Teramura, Y., & Kannari, F. *J. Opt. Soc. Am. B* **19**, 2795 (2002). 11, 30, 60

[72] Feurer, T., Vaughan, J., Koehl, R., & Nelson, K. *Opt. Lett.* **27**, 652 (2002). 11, 57

[73] Cohen, L. *Time-frequency Analysis.* Prentice Hall (1995). 16

[74] Cohen, L. *Proc. IEEE* **77**, 941 (1989). 16

[75] Paye, J. *IEEE J. Quantum. Electron.* **28** (1992). 16

[76] Bor, Z., Gogolak, Z., & Szabo, G. *Opt. Lett.* **14**, 862 (1989). 19, 60

[77] Hebling, J. *Opt. Quantum. Electron.* **28**, 1759 (1996). 19

[78] Akturk, S., Gu, X., Zeek, E., & Trebino, R. *Opt. Express* **12**, 4399 (2004). 19

[79] Cundiff, S. T., Knox, W. H., Ippen, E. P., & Haus, H. A. *Opt. Lett.* **21**, 662 (1996). 20

[80] Jones, D. J., Diddams, S. A., Ranka, J. K., Stentz, A., Windeler, R. S., Hall, J. L., & Cundiff, S. T. *Science* **288**, 635 (2000). 22

[81] Keusters, D., Tan, H., O'Shea, P., Zeek, E., Trebino, R., & Warren, W. *J. Opt. Soc. Am. B* **20**, 2226 (2003). 22, 63, 67, 69, 79, 91, 93, 217

[82] Czerny, M. & Turner, A. *Z. Phys. A: Hadrons Nucl.* **61**, 792 (1930). 25

[83] Reynaud, F., Salin, F., & Barthelemy, A. *Opt. Lett.* **14**, 275 (1989). 26

[84] Froehly, C., Lacourt, A., & Viénot, J. C. *Nouvelle Revue d'Optique* **4**, 183 (1973). 26

[85] Piasecki, J., Colombeau, B., Vampouille, M., Froehly, C., & Arnaud, J. A. *Appl. Opt.* **19**, 3749 (1980). 26

[86] Sáinz, C., Jourdain, P., Escalona, R., & Calatroni, J. *Opt. Commun.* **111**, 632 (1994). 26

[87] Fittinghoff, D., Bowie, J., Sweetser, J., Jennings, R., Krumbügel, M., DeLong, K., Trebino, R., & Walmsley, I. *Opt. Lett.* **21**, 884 (1996). 26

[88] Lepetit, L., Chériaux, G., & Joffre, M. *J. Opt. Soc. Am. B* **12**, 2467 (1995). 26, 31

[89] Dorrer, C. & Joffre, M. *C. R. Acad. Sci. Paris* **2**, 1415 (2001). 27

[90] Knox, W. H., Pearson, N. M., Li, K. D., & Hirlimann, C. A. *Opt. Lett.* **13**, 574 (1988). 27

[91] Naganuma, K. & Sakai, Y. *Opt. Lett.* **19**, 487 (1994). 27

[92] Takeda, M., Ina, H., & Kobayashi, S. *J. Opt. Soc. Am.* **72**, 156 (1982). 28, 55, 60, 158

[93] Vakhtin, A., Peterson, K., Wood, W., & Kane, D. *Opt. Lett.* **28**, 1332 (2003). 30

[94] Kovács, A. P., Osvay, K., Bor, Z., & Szipöcs, R. *Opt. Lett.* **20**, 788 (1995). 30

REFERENCES

[95] Kovács, A., Osvay, K., Kurdi, G., Gã̋rbe, M., Klebniczki, J., & Bor, Z. *Appl. Phys. B: Lasers Opt.* **80**, 165 (2005). 30

[96] Amir, W., Planchon, T. A., Durfee, C. G., & Squier, J. A. *Opt. Lett.* **32**, 939 (2007). 30, 60

[97] Birge, J. R., Ell, R., & Kärtner, F. X. *Opt. Lett.* **31**, 2063 (2006). 31, 48

[98] Weber, H. P. *J. Appl. Phys.* **38**, 2231 (1967). 32

[99] Giordmaine, J. A., Rentzepis, P. M., Shapiro, S. L., & Wecht, K. W. *Appl. Phys. Lett.* **11**, 216 (1967). 32

[100] Shapiro, S. L. *Appl. Phys. Lett.* **13**, 19 (1968). 32

[101] Auston, D. H. *Appl. Phys. Lett.* **18**, 249 (1971). 32

[102] Peatross, J. & Rundquist, A. *J. Opt. Soc. Am. B* **15**, 216 (1998). 32

[103] Chung, J.-H. & Weiner, A. *IEEE J. Sel. Top. Quant.* **7**, 656 (2001). 32

[104] Koumans, R. & Yariv, A. *IEEE J. Quantum. Elect.* **36**, 137 (2000). 33

[105] Diels, J. M., Fontaine, J. J., McMichael, I. C., & Simoni, F. *Applied Optics* **24**, 1270 (1985). 33

[106] Mindl, T. and Hefferle, P. and Schneider, S. and Dörr, F. *Appl. Phys. B: Lasers Opt.* **31**, 201 (1983). 33

[107] Naganuma, K., Mogi, K., & Yamada, H. *IEEE J. Quantum. Elect.* **25**, 1225 (1989). 33

[108] Hache, F., Driscoll, T. J., Cavallari, M., & Gale, G. M. *Appl. Opt.* **35**, 3230 (1996). 33

[109] Feurer, T., Glass, A., & Sauerbrey, R. *Appl. Phys. B: Lasers Opt.* **65**, 295 (1997). 33

[110] Ranka, J. K., Gaeta, A. L., Baltuška, A., Pshenichnikov, M. S., & Wiersma, D. A. *Opt. Lett.* **22**, 1344 (1997). 33

[111] Spielmann, C., Xu, L., & Krausz, F. *Appl. Opt.* **36**, 2523 (1997). 33

[112] Mashiko, H., Suda, A., & Midorikawa, K. *Appl. Phys. B: Lasers Opt.* **87**, 221 (2007). 33

[113] Kozma, I. Z., Baum, P., Schmidhammer, U., Lochbrunner, S., & Riedle, E. *Rev. Sci. Instrum.* **75**, 2323 (2004). 33

[114] Ventalon, C., Fraser, J., Likforman, J., Villeneuve, D., Corkum, P., & Joffre, M. *J. Opt. Soc. Am. B* **23**, 332 (2006). 33

[115] Meshulach, D., Yelin, D., & Silberberg, Y. *J. Opt. Soc. Am. B* **15**, 1615 (1998). 34

[116] Nicholson, J. W., Jasapara, J., Rudolph, W., Omenetto, F. G., & Taylor, A. J. *Opt. Lett.* **24**, 1774 (1999). 34

[117] Nicholson, J. W. & Rudolph, W. *J. Opt. Soc. Am. B* **19**, 330 (2002). 34

[118] Dahlström, L. & Källberg, B. *Opt. Commun.* **4**, 285 (1971). 35

[119] Albrecht, H.-S., Heist, P., Kleinschmidt, J., Lap, D., & Schröder, T. *Appl. Phys. B: Lasers Opt.* **55**, 362 (1992). 35

[120] Luan, S., Hutchinson, M. H. R., Smith, R. A., & Zhou, F. *Meas. Sci. Technol.* **4**, 1426 (1993). 35

[121] Divall, E. J. & Ross, I. N. *Opt. Lett.* **29**, 2273 (2004). 35

[122] Collier, J. L., Hernandez-Gomez, C., Allott, R., Danson, C., & Hall, A. *Laser Part. Beams* **19**, 231 (2001). 35

[123] Dorrer, C., Bromage, J., & Zuegel, J. D. *Opt. Express* **16**, 13534 (2008). 35

[124] Nicholson, J. W., Mero, M., Jasapara, J., & Rudolph, W. *Opt. Lett.* **25**, 1801 (2000). 35

[125] Feurer, T., Niedermeier, S., & Sauerbrey, R. *Appl. Phys. B: Lasers Opt.* **66**, 163 (1998). 35

[126] Janszky, J., Corradi, G., & Gyuzalian, R. N. *Opt. Commun.* **23**, 293 (1977). 35

[127] Gyuzalian, R., Sogomonian, S., & Horvath, Z. *Opt. Commun.* **29**, 239 (1979). 35

[128] Ishida, Y., Yajima, T., & Watanabe, A. *Opt. Commun.* **56**, 57 (1985). 35

[129] Salin, F., Georges, P., Roger, G., & Brun, A. *Appl. Opt.* **26**, 4528 (1987). 35

[130] Simon, P., Gerhardt, H., & Szatmari, S. *Meas. Sci. Technol.* **1**, 637 (1990). 35

[131] Kane, D. J. *IEEE J. Quantum Electron.i* **35**, 421 (1999). 37

[132] Xu, L., Zeek, E., & Trebino, R. *J. Opt. Soc. Am. B* **25**, A70 (2008). 37

[133] Yellampalle, B., Kim, K., & Taylor, A. *Opt. Lett.* **32**, 3558 (2007). 37

[134] Yellampalle, B., Kim, K., & Taylor, A. J. *Opt. Lett.* **33**, 2854 (2008). 37

[135] Yellampalle, B., Kim, K., & Taylor, A. J. *Opt. Lett.* **34**, 2603 (2009). 37

[136] Xu, L., Kane, D. J., & Trebino, R. *Opt. Lett.* **34**, 2602 (2009). 37

[137] Gomes, A., Gouveia-Neto, A., & Taylor, J. *Electron. Lett.* **22**, 41 (1986). 37

[138] Olson, C., Izadpanah, H., & Lin, C. *Electron. Lett.* **25**, 1018 (1989). 37

[139] Mori, K., Morioka, T., & Saruwatari, M. *Electron. Lett.* **29**, 987 (1993). 37

[140] Wong, V. & Walmsley, I. A. *J. Opt. Soc. Am. B* **14**, 944 (1997). 37, 38

[141] Treacy, E. B. *J. Appl. Phys.* **42**, 3848 (1971). 37, 38

[142] Chilla, J. & Martinez, O. *IEEE J. Quantum. Electron.* **27**, 1228 (1991). 37, 38

[143] Reid, D. *IEEE J. Quantum. Elect.* **35**, 1584 (1999). 38

[144] Reid, D. T. & Cormack, I. G. *Opt. Lett.* **27**, 658 (2002). 38

[145] Panasenko, D., Sun, P.-C., Alic, N., & Fainman, Y. *Appl. Opt.* **41**, 5185 (2002). 38

[146] Panasenko, D. & Fainman, Y. *Opt. Lett.* **27**, 1475 (2002). 38

[147] Watanabe, A., Saito, H., Ishida, Y., & Yajima, T. *Opt. Commun.* **63**, 320 (1987). 38

[148] Ishida, Y., Naganuma, K., & Yajima, T. *IEEE J. Quantum. Electron.* **21**, 69 (1985). 38

[149] Watanabe, A., Tanaka, S., Kobayashi, H., Ishida, Y., & Yajima, T. *Rev. Sci. Instrum.* **56**, 2259 (1985). 38

[150] Paye, J., Ramaswamy, M., Fujimoto, J. G., & Ippen, E. P. *Opt. Lett.* **18**, 1946 (1993). 38

[151] Paye, J. *IEEE J. Quantum. Electron.* **30**, 2693 (1994). 38

[152] DeLong, K. W., Trebino, R., Hunter, J., & White, W. E. *J. Opt. Soc. Am. B* **11**, 2206 (1994). 38

[153] Barry, L., Del Burgo, S., Thomsen, B., Watts, R., & Reid, D. *IEEE Photonic. Tech. L.* **14**, 971 (2002). 39

[154] Barry, L., Dudley, J., Bollond, P., Harvey, J., & Leonhardt, R. *Electron. Lett.* **32**, 2339 (1996). 39

[155] Barry, L., Thomsen, B., Dudley, J., & Harvey, J. *IEEE Photonic. Tech. L.* **10**, 935 (1998). 39

[156] Miao, H., Yang, S., Langrock, C., Roussev, R., Fejer, M., & Weiner, A. *J. Opt. Soc. Am. B* **25**, 41 (2008). 39

[157] Baltuška, A., Pshenichnikov, M., & Wiersma, D. *IEEE J. Quantum. Electron.* **35**, 459 (1999). 39

[158] Akturk, S., D'Amico, C., & Mysyrowicz, A. *J. Opt. Soc. Am. B* **25**, A63 (2008). 39

[159] Richman, B. A., Krumbügel, M. A., & Trebino, R. *Opt. Lett.* **22**, 721 (1997). 39

[160] Dmitriev, V., Gurzadyan, G., Nikogosyan, D., & Braun, C. *Handbook of Nonlinear Optical Crystals.* Springer (1997). 39

[161] Fittinghoff, D. N., Squier, J. A., Barty, C., Sweetser, J. N., Trebino, R., & Müller, M. *Opt. Lett.*

REFERENCES

23, 1046 (1998). 39

[162] Fittinghoff, D., Millard, A., Squier, J., & Muller, M. *IEEE J. Quantum. Electron.* **35**, 479 (1999). 39

[163] Gallmann, L., Steinmeyer, G., Sutter, D. H., Matuschek, N., & Keller, U. *Opt. Lett.* **25**, 269 (2000). 39

[164] Stibenz, G. & Steinmeyer, G. *Opt. Express* **13**, 2617 (2005). 39

[165] Stibenz, G. & Steinmeyer, G. *IEEE J. Sel. Topics in Quantum Electron.* **12**, 286 (2006). 39

[166] Zeek, E., Shreenath, A., O'Shea, P., Kimmel, M., & Trebino, R. *Appl. Phys. B: Lasers Opt.* **74**, s265 (2002). 39

[167] Kane, D. & Trebino, R. *IEEE J. Quantum. Electron.* **29**, 571 (1993). 39

[168] Tsang, T., Krumbügel, M. A., DeLong, K. W., Fittinghoff, D. N., & Trebino, R. *Opt. Lett.* **21**, 1381 (1996). 40

[169] Ramos-Ortiz, G., Cha, M., Thayumanavan, S., Mendez, J., Marder, S. R., & Kippelen, B. *Appl. Phys. Lett.* **85**, 3348 (2004). 40

[170] Chadwick, R., Spahr, E., Squier, J. A., Durfee, C. G., Walker, B. C., & Fittinghoff, D. N. *Opt. Lett.* **31**, 3366 (2006). 40

[171] Sweetser, J. N., Fittinghoff, D. N., & Trebino, R. *Opt. Lett.* **22**, 519 (1997). 40

[172] Lacourt, P.-A., Dudley, J. M., Merolla, J.-M., Porte, H., Goedgebuer, J.-P., & Rhodes, W. T. *Opt. Lett.* **27**, 863 (2002). 40

[173] Kang, I., Dorrer, C., Zhang, L., Dinu, M., Rasras, M., Buhl, L., Cabot, S., Bhardwaj, A., Liu, X., Cappuzzo, M., Gomez, L., Wong-Foy, A., Chen, Y., Dutta, N., Patel, S., Neilson, D., Giles, C., Piccirilli, A., & Jaques, J. *IEEE J. Sel. Top. Quant.* **14**, 758 (2008). 40

[174] Linden, S., Kuhl, J., & Giessen, H. *Opt. Lett.* **24**, 569 (1999). 40

[175] Linden, S., Giessen, H., & Kuhl, J. *Phys. Status Solidi B* **206**, 119 (1998). 40, 65

[176] Reid, D. T., Loza-Alvarez, P., Brown, C. T. A., Beddard, T., & Sibbett, W. *Opt. Lett.* **25**, 1478 (2000). 40

[177] Zhang, J., Shreenath, A., Kimmel, M., Zeek, E., Trebino, R., & Link, S. *Opt. Express* **11**, 601 (2003). 40

[178] Zhang, J.-Y., Lee, C.-K., Huang, J., & Pan, C.-L. *Opt. Express* **12**, 574 (2004). 40

[179] Dudley, J. M., Gu, X., Xu, L., Kimmel, M., Zeek, E., O'shea, P., Trebino, R., Coen, S., & Windeler, R. S. *Opt. Express* **10**, 1215 (2002). 40

[180] Efimov, A. & Taylor, A. J. *Opt. Express* **16**, 5942 (2008). 40

[181] Kane, D. J., Taylor, A. J., Trebino, R., & DeLong, K. W. *Opt. Lett.* **19**, 1061 (1994). 41

[182] O'Shea, P., Kimmel, M., Gu, X., & Trebino, R. *Opt. Lett.* **26**, 932 (2001). 41

[183] Radzewicz, C., Wasylczyk, P., & Krasinski, J. S. *Opt. Commun.* **186**, 329 (2000). 41

[184] Shannon, C. *Proceedings of the IRE* **37**, 10 (1949). 42

[185] Anderson, M., Monmayrant, A., Gorza, S.-P., Wasylczyk, P., & Walmsley, I. *Laser Physics Letters* **5**, 259 (2008). 44

[186] Walmsley, I. A. & Dorrer, C. *Adv. Opt. Photon.* **1**, 308 (2009). 44, 63, 67

[187] Iaconis, C. & Walmsley, I. *IEEE J. Quantum. Electron.* **35**, 501 (1999). 44

[188] Shuman, T., Walmsley, I. A., Waxer, L., Anderson, M., Iaconis, C., & Bromage, J. *Opt. Express* **5**, 134 (1999). 44

[189] Kornelis, W., Biegert, J., Tisch, J. W. G., Nisoli, M., Sansone, G., Vozzi, C., Silvestri, S. D., & Keller, U. *Opt. Lett.* **28**, 281 (2003). 44

[190] Gallmann, L., Sutter, D., Matuschek, N., Steinmeyer, G., & Keller, U. *Appl. Phys. B: Lasers Opt.* **70**, 67 (2000). 45

[191] Stibenz, G. & Steinmeyer, G. *Rev. Sci. Instrum.* **77**, 073105 (2006). 45

[192] Dorrer, C. *Opt. Lett.* **24**, 1532 (1999). 45

[193] Müller, A. & Laubscher, M. *Opt. Lett.* **26**, 1915 (2001). 45

[194] Baum, P., Lochbrunner, S., & Riedle, E. *Opt. Lett.* **29**, 210 (2004). 46

[195] Baum, P. & Riedle, E. *J. Opt. Soc. Am. B* **22**, 1875 (2005). 46

[196] Kosik, E., Radunsky, A., Walmsley, I., & Dorrer, C. *Opt. Lett.* **30**, 326 (2005). 46, 57, 91, 93

[197] Wyatt, A., Walmsley, I., Stibenz, G., & Steinmeyer, G. *Opt. Lett.* **31**, 1914 (2006). 46, 57, 104

[198] Dorrer, C., Londero, P., & Walmsley, I. A. *Opt. Lett.* **26**, 1510 (2001). 46

[199] Radunsky, A., Walmsley, I., Gorza, S.-P., & Wasylczyk, P. *Opt. Lett.* **32**, 181 (2007). 47

[200] Gorza, S.-P., Wasylczyk, P., & Walmsley, I. A. *Opt. Express* **15**, 15168 (2007). 47, 82, 91

[201] Miao, H., Leaird, D. E., Langrock, C., Fejer, M. M., & Weiner, A. M. *Opt. Express* **17**, 3381 (2009). 48

[202] Messager, V., Louradour, F., Froehly, C., & Barthelemy, A. *Opt. Lett.* **28**, 743 (2003). 48

[203] Lelek, M., Louradour, F., Barthélémy, A., & Froehly, C. *Opt. Commun.* **261**, 124 (2006). 48

[204] Lelek, M., Louradour, F., Barthélémy, A., Froehly, C., Mansourian, T., Mouradian, L., Chambaret, J.-P., Chériaux, G., & Mercier, B. *J. Opt. Soc. Am. B* **25**, A17 (2008). 48

[205] Oksenhendler, T., Coudreau, S., Forget, N., Crozatier, V., Grabielle, S., Herzog, R., Gobert, O., & Kaplan, D. *Appl. Phys. B: Lasers Opt.* **99**, 7 (2010). 49

[206] Kolner, B. H. & Nazarathy, M. *Opt. Lett.* **14**, 630 (1989). 50

[207] Christov, I. P. *Opt. Quantum. Electron.* **22**, 473 (1990). 50

[208] Bennett, C. & Kolner, B. *Opt. Lett.* **24**, 783 (1999). 50, 52

[209] Bennett, C. V., Scott, R. P., & Kolner, B. H. *Appl. Phys. Lett.* **65**, 2513 (1994). 50

[210] Vampouille, M., Marty, J., & Froehly, C. *IEEE J. Quantum. Elect.* **22**, 192 (1986). 50, 52

[211] Kauffman, M. T., Banyai, W. C., Godil, A. A., & Bloom, D. M. *Appl. Phys. Lett.* **64**, 270 (1994). 50

[212] Azaña, J., Berger, N. K., Levit, B., & Fischer, B. *Appl. Opt.* **43**, 483 (2004). 50

[213] Mouradian, L., Louradour, F., Messager, V., Barthelemy, A., & Froehly, C. *IEEE J. Quantum. Elect.* **36**, 795 (2000). 50, 52

[214] Foster, M., Salem, R., Geraghty, D., Turner-Foster, A., Lipson, M., & Gaeta, A. *Nature* **456**, 81 (2008). 50, 52

[215] Gerchberg, R. & Saxton, W. *Optik* **35**, 237 (1972). 50, 61

[216] Saxton, W. *Computer Techniques for Image Processing in Electron Microscopy*. Academic Press (1978). 50, 61

[217] Cohn, P. *Further algebra and applications*. Springer Verlag (2003). 51

[218] Kak, A. & Slaney, M. *Principles of computerized tomographic imaging*. Society for Industrial Mathematics (1999). 51

[219] Webb, S. *From the watching of shadows: the origins of radiological tomography*. Taylor &

REFERENCES

Francis (1990). 51

[220] Beck, M., Raymer, M. G., Walmsley, I. A., & Wong, V. *Opt. Lett.* **18**, 2041 (1993). 51

[221] Alieva, T., Bastiaans, M., & Stankovic, L. *IEEE T. Signal. Proces.* **51**, 112 (2003). 51

[222] Dorrer, C. *Opt. Lett.* **30**, 3237 (2005). 51

[223] Dorrer, C. & Kang, I. *Opt. Lett.* **28**, 1481 (2003). 51

[224] Ng, T., Parmigiani, F., Ibsen, M., Zhang, Z., Petropoulos, P., & Richardson, D. *IEEE Photonic. Tech. L.* **20**, 1097 (2008). 52

[225] Nibbering, E. T. J., Franco, M. A., Prade, B. S., Grillon, G., Chambaret, J.-P., & Mysyrowicz, A. *J. Opt. Soc. Am. B* **13**, 317 (1996). 52

[226] Ferreiro, J. J., de la Fuente, R., & López-Lago, E. *Opt. Lett.* **26**, 1025 (2001). 52

[227] Shack, R. & Platt, B. *J. Opt. Soc. Am* **61**, 656 (1971). 54

[228] Hartmann, J. *Z. Instrumentenkd* **20**, 47 (1900). 54

[229] Platt, B. *J. Refract. Surg.* **17**, 573 (2001). 54

[230] Bonaretti, F., Faccio, D., Clerici, M., Biegert, J., & Trapani, P. D. *Opt. Express* **17**, 9804 (2009). 54

[231] Dutton, D., Cornejo, A., & Latta, M. *Appl. Opt.* **7**, 125 (1968). 55

[232] Smartt, R. & Steel, W. *Jpn. J. Appl. Phys* **14**, 351 (1975). 56

[233] Bates, W. J. *Proc. Phys. Soc. London* **59**, 940 (1947). 56

[234] Murty, M. V. R. K. *Appl. Opt.* **3**, 531 (1964). 56

[235] Drew, R. L. *Proc. Phys. Soc. London, Sect. B* **64**, 1005 (1951). 56

[236] Brown, D. *Proc. Phys. Soc. London, Sect. B* **67**, 232 (1954). 56

[237] Saunders, J. *Precision Meas. and Calibration* **10**, 165 (1964). 56

[238] Paez, G., Strojnik, M., & Torales, G. G. *Appl. Opt.* **39**, 5172 (2000). 56

[239] Ronchi, V. *Appl. Opt.* **3**, 437 (1964). 56

[240] Walmsley, I. A. & Bigelow, N. P. *Phys. Rev. A* **57**, R713 (1998). 56

[241] Murty, M. V. R. K. *Appl. Opt.* **3**, 853 (1964). 56

[242] Brown, D. *Her Majesty's Stationery Office, London* pages 253–256 (1959). 56

[243] Hariharan, P. & Sen, D. *J. Sci. Instrum.* **38**, 428 (1961). 56

[244] Murty, M. V. R. K. & Hagerott, E. C. *Appl. Opt.* **5**, 615 (1966). 57

[245] Vaughan, J., Feurer, T., & Nelson, K. *J. Opt. Soc. Am. B* **19**, 2489 (2002). 57

[246] Bowlan, P., Gabolde, P., Shreenath, A., McGresham, K., Trebino, R., & Akturk, S. *Opt. Express* **14**, 11892 (2006). 57

[247] Bowlan, P., Gabolde, P., Coughlan, M., Trebino, R., & Levis, R. *J. Opt. Soc. Am. B* **25**, 81 (2008). 57

[248] Bowlan, P., Gabolde, P., & Trebino, R. *Opt. Express* **15**, 10219 (2007). 57

[249] Szatmári, S. & Kühnle, G. *Opt. Commun.* **69**, 60 (1988). 58

[250] Sacks, Z., Mourou, G., & Danielius, R. *Opt. Lett.* **26**, 462 (2001). 58

[251] Akturk, S., Kimmel, M., O'Shea, P., & Trebino, R. *Opt. Express* **11**, 68 (2003). 58

[252] Gabolde, P. & Trebino, R. *Opt. Express* **12**, 4423 (2004). 59

[253] Meshulach, D., Yelin, D., & Silberberg, Y. *J. Opt. Soc. Am. B* **14**, 2095 (1997). 60

[254] Geindre, J. P., Audebert, P., Rousse, A., Falliès, F., Gauthier, J. C., Mysyrowicz, A., Santos, A. D., Hamoniaux, G., & Antonetti, A. *Opt. Lett.* **19**, 1997 (1994). 60

[255] Jasapara, J. & Rudolph, W. *Opt. Lett.* **24**, 777 (1999). 60

[256] Amir, W., Planchon, T., Durfee, C., Squier, J., Gabolde, P., Trebino, R., & Müller, M. *Opt. Lett.* **31**, 2927 (2006). 60

[257] Bacon, R., Adam, G., Baranne, A., Courtes, D., Dubet, D., Dubois, J., Emsellem, E., Ferruit, P., Georgelin, Y., Monnet, G., *et al. Astronomy and Astrophysics-Supplement Series* **113**, 347 (1995). 60

[258] Tecza, M., Thatte, N., Clarke, F., Goodsall, T., & Symeonidis, M. *New Astron. Rev.* **49**, 647 (2006). 60

[259] Gabolde, P. & Trebino, R. *Opt. Express* **14**, 11460 (2006). 60

[260] Gabolde, P. & Trebino, R. *J. Opt. Soc. Am. B* **25**, A25 (2008). 60

[261] Radzewicz, C., la Grone, M. J., & Krasinski, J. S. *Opt. Commun.* **126**, 185 (1996). 60

[262] Rubino, E., Faccio, D., Tartara, L., Bates, P. K., Chalus, O., Clerici, M., Bonaretti, F., Biegert, J., & Trapani, P. D. *Opt. Lett.* **34**, 3854 (2009). 61

[263] Bragheri, F., Faccio, D., Bonaretti, F., Lotti, A., Clerici, M., Jedrkiewicz, O., Liberale, C., Henin, S., Tartara, L., Degiorgio, V., & Trapani, P. D. *Opt. Lett.* **33**, 2952 (2008). 61

[264] Rouyer, C., Blanchot, N., Neauport, J., & Sauteret, C. *Opt. Express* **15**, 2019 (2007). 61

[265] Kosik, E. M. *Characterization of Ultrashort Optical Pulses.* Ph.D. thesis, University of Rochester (2004). 61, 107, 108

[266] Wyatt, A. *Spectral interferometry for the complete characterisation of near infrared femtosecond and extreme ultraviolet attosecond pulses.* Ph.D. thesis, University of Oxford (2007). 63, 155

[267] Austin, D. R., Witting, T., & Walmsley, I. A. *J. Opt. Soc. Am. B* **26**, 1818 (2009). 63, 93, 98

[268] Dudley, J. M., Genty, G., & Coen, S. *Rev. Mod. Phys.* **78**, 1135 (2006). 64, 65

[269] La Fontaine, B., Vidal, F., Jiang, Z., Chien, C. Y., Comtois, D., Desparois, A., Johnston, T. W., Kieffer, J.-C., Pépin, H., & Mercure, H. P. *Phys. Plasmas* **6**, 1615 (1999). 64

[270] Weiner, A., Heritage, J., & Salehi, J. *Opt. Lett.* **13**, 300 (1988). 64

[271] Chowdhury, I. H., Xu, X., & Weiner, A. M. *Appl. Phys. Lett.* **86**, 151110 (2005). 64

[272] Stratmann, M., Pagel, T., & Mitschke, F. *Phys. Rev. Lett.* **95**, 143902 (2005). 65

[273] Sohn, J. *Appl. Phys. Lett.* **81**, 13 (2002). 65

[274] Rector, K. D. & Fayer, M. D. *Int. Rev. Phys. Chem.* **17**, 261 (1998). 65

[275] de Boeij, W. P., Pshenichnikov, M. S., & Wiersma, D. A. *Chem. Phys. Lett.* **247**, 264 (1995). 65

[276] de Boeij, W. P., Pshenichnikov, M. S., & Wiersma, D. A. *Chem. Phys.* **233**, 287 (1998). 65

[277] Sundermann, K. & de Vivie-Riedle, R. *J. Chem. Phys.* **110**, 1896 (1999). 65

[278] Lee, D., Gabolde, P., & Trebino, R. *J. Opt. Soc. Am. B* **25**, A34 (2008). 66

[279] Anderson, M., de Araujo, L., Kosik, E., & Walmsley, I. *Appl. Phys. B* **70**, S85 (2000). 67

[280] Wang, Z., Zeek, E., Trebino, R., & Kvam, P. *J. Opt. Soc. Am. B* **20**, 2400 (2003). 67

[281] Wang, Z., Zeek, E., Trebino, R., & Kvam, P. *Opt. Express* **11**, 3518 (2003). 67, 68

[282] Witting, T., Austin, D. R., & Walmsley, I. A. *Opt. Express* **17**, 18983 (2009). 68, 82, 91, 98

[283] Dorrer, C. & Walmsley, I. *J. Opt. Soc. Am. B* **19**, 1019 (2002). 68, 88, 224

REFERENCES

[284] Dorrer, C. & Walmsley, I. *J. Opt. Soc. Am. B* **19**, 1030 (2002). 68, 88, 224

[285] Hirasawa, M., Nakagawa, N., Yamamoto, K., Morita, R., Shigekawa, H., & Yamashita, M. *Appl. Phys. B* **74**, S225 (2002). 69, 93

[286] Zavelani-Rossi, M., Polli, D., Cerullo, G., De Silvestri, S., Gallmann, L., Steinmeyer, G., & Keller, U. *Appl. Phys. B* **74**, S245 (2002). 69, 94

[287] Jones, G. A. & Jones, J. M. *Elementary Number Theory*. Springer (1998). 71

[288] Witting, T., Austin, D. R., & Walmsley, I. A. *Opt. Lett.* **34**, 881 (2009). 91

[289] Dorrer, C. & Kang, I. *J. Opt. Soc. Am. B* **25**, A1 (2008). 91

[290] Rimmer, M. P. *Appl. Opt.* **13**, 623 (1974). 92

[291] Velghe, S., Primot, J., Guérineau, N., Cohen, M., & Wattellier, B. *Opt. Lett.* **30**, 245 (2005). 92

[292] Fried, D. L. *Opt. Commun.* **200**, 43 (2001). 92

[293] Ghai, D. P., Senthilkumaran, P., & Sirohi, R. *Opt. Commun.* **281**, 1315 (2008). 92

[294] Austin, D. R., Witting, T., & Walmsley, I. A. *Appl. Opt.* **48**, 3846 (2009). 94

[295] Ghiglia, D. & Romero, L. *J. Opt. Soc. Am. A* **11**, 107 (1994). 106, 117, 119

[296] Wyatt, A., Witting, T., Monmayrant, A., Walmsley, I., Haworth, C., Robinson, J., Tisch, J., & Marangos, J. In *European Conference on Lasers and Electro-Optics and the European Quantum Electronics Conference (CLEO/EQEC)* (2009). 110

[297] Fried, D. L. *J. Opt. Soc. Am.* **67**, 370 (1977). 117

[298] Hudgin, R. H. *J. Opt. Soc. Am.* **67**, 375 (1977). 117

[299] Noll, R. J. *J. Opt. Soc. Am.* **68**, 139 (1978). 117

[300] Hunt, B. R. *J. Opt. Soc. Am.* **69**, 393 (1979). 117

[301] Sandberg, R. L., Paul, A., Raymondson, D. A., Hädrich, S., Gaudiosi, D. M., Holtsnider, J., Tobey, R. I., Cohen, O., Murnane, M. M., Kapteyn, H. C., Song, C., Miao, J., Liu, Y., & Salmassi, F. *Phys. Rev. Lett.* **99**, 098103 (2007). 125

[302] Sarukura, N., Hata, K., Adachi, T., Nodomi, R., Watanabe, M., & Watanabe, S. *Phys. Rev. A* **43**, 1669 (1991). 126

[303] Crane, J., Perry, M., Herman, S., & Falcone, R. *Opt. Lett.* **17**, 1256 (1992). 126

[304] Faldon, M. E., Hutchinson, M. H. R., Marangos, J. P., Muffett, J. E., Smith, R. A., Tisch, J. W. G., & Wahlström, C. G. *J. Opt. Soc. Am. B* **9**, 2094 (1992). 126

[305] Kondo, K., Sarukura, N., Sajiki, K., & Watanabe, S. *Phys. Rev. A* **47**, R2480 (1993). 126

[306] Macklin, J. J., Kmetec, J. D., & Gordon, C. L. *Phys. Rev. Lett.* **70**, 766 (1993). 126

[307] Perry, M. D. & Crane, J. K. *Phys. Rev. A* **48**, R4051 (1993). 126

[308] Wahlström, C.-G., Larsson, J., Persson, A., Starczewski, T., Svanberg, S., Salières, P., Balcou, P., & L'Huillier, A. *Phys. Rev. A* **48**, 4709 (1993). 126

[309] Tisch, J. W. G., Smith, R. A., Muffett, J. E., Ciarrocca, M., Marangos, J. P., & Hutchinson, M. H. R. *Phys. Rev. A* **49**, R28 (1994). 126

[310] Miyazaki, K. & Takada, H. *Phys. Rev. A* **52**, 3007 (1995). 126

[311] McPherson, A., Gibson, G., Jara, H., Johann, U., Luk, T., McIntyre, I. A., Boyer, K., & Rhodes, C. K. *J. Opt. Soc. Am. B* **4**, 595 (1987). 126

[312] Li, X. F., L'Huillier, A., Ferray, M., Lompré, L. A., & Mainfray, G. *Phys. Rev. A* **39**, 5751 (1989). 126

[313] L'Huillier, A. & Balcou, P. *Phys. Rev. Lett.* **70**, 774 (1993). 126

[314] Haight, R. & Seidler, P. F. *Appl. Phys. Lett.* **65**, 517 (1994). 126

[315] Theobald, W., Häßner, R., Wülker, C., & Sauerbrey, R. *Phys. Rev. Lett.* **77**, 298 (1996). 126

[316] Schnürer, M., Streli, C., Wobrauschek, P., Hentschel, M., Kienberger, R., Spielmann, C., & Krausz, F. *Phys. Rev. Lett.* **85**, 3392 (2000). 126

[317] Farkas, G. & Tóth, C. *Phys. Lett. A* **168**, 447 (1992). 126

[318] Harris, S. E., Macklin, J. J., & Hänsch, T. W. *Opt. Commun.* **100**, 487 (1993). 126

[319] Hänsch, T. W. *Opt. Commun.* **80**, 71 (1990). 126

[320] Paul, P. M., Toma, E. S., Breger, P., Mullot, G., Auge, F., Balcou, P., Muller, H. G., & Agostini, P. *Science* **292**, 1689 (2001). 126, 152

[321] Mairesse, Y., de Bohan, A., Frasinski, L. J., Merdji, H., Dinu, L. C., Monchicourt, P., Breger, P., Kovačev, M., Taïeb, R., Carre, B., Muller, H. G., Agostini, P., & Salieres, P. *Science* **302**, 1540 (2003). 126

[322] Corkum, P. B., Burnett, N. H., & Ivanov, M. Y. *Opt. Lett.* **19**, 1870 (1994). 126

[323] Chang, Z. *Phys. Rev. A* **70**, 043802 (2004). 126

[324] Ivanov, M., Corkum, P. B., Zuo, T., & Bandrauk, A. *Phys. Rev. Lett.* **74**, 2933 (1995). 126

[325] Platonenko, V. T. & Strelkov, V. V. *J. Opt. Soc. Am. B* **16**, 435 (1999). 126

[326] Tcherbakoff, O., Mével, E., Descamps, D., Plumridge, J., & Constant, E. *Phys. Rev. A* **68**, 043804 (2003). 126

[327] Altucci, C., Delfin, C., Roos, L., Gaarde, M. B., L'Huillier, A., Mercer, I., Starczewski, T., & Wahlström, C.-G. *Phys. Rev. A* **58**, 3934 (1998). 126

[328] Kovačev, M., Mairesse, Y., Priori, E., Merdji, H., Tcherbakoff, O., Monchicourt, P., Breger, P., Mével, E., Constant, E., Salières, P., Carré, B., & Agostini, P. *Eur. Phys. J. D* **26**, 79 (2003). 126

[329] Sola, I. J., Mével, E., Elouga, L., Constant, E., Strelkov, V., Poletto, L., Villoresi, P., Benedetti, E., Caumes, J.-P., Stagira, S., Vozzi, C., Sansone, G., & Nisoli, M. *Nat Phys* **2**, 319 (2006). 127

[330] Chang, Z. *Phys. Rev. A* **76**, 051403 (2007). 127

[331] Mashiko, H., Gilbertson, S., Li, C., Khan, S. D., Shakya, M. M., Moon, E., & Chang, Z. *Phys. Rev. Lett.* **100**, 103906 (2008). 127

[332] Tzallas, P., Skantzakis, E., Kalpouzos, C., Benis, E. P., Tsakiris, G. D., & Charalambidis, D. *Nat Phys* **3**, 846 (2007). 127

[333] Charalambidis, D., Tzallas, P., Benis, E. P., Skantzakis, E., Maravelias, G., Nikolopoulos, L. A. A., Conde, A. P., & Tsakiris, G. D. *New J. Phys.* **10**, 025018 (2008). 127

[334] Zhou, J., Peatross, J., Murnane, M. M., Kapteyn, H. C., & Christov, I. P. *Phys. Rev. Lett.* **76**, 752 (1996). 127

[335] Schafer, K. J. & Kulander, K. C. *Phys. Rev. Lett.* **78**, 638 (1997). 127

[336] Chang, Z., Rundquist, A., Wang, H., Murnane, M. M., & Kapteyn, H. C. *Phys. Rev. Lett.* **79**, 2967 (1997). 127

[337] Spielmann, C., Burnett, N., Sartania, S., Koppitsch, R., Schnurer, M., Kan, C., Lenzner, M., Wobrauschek, P., & Krausz, F. *Science* **278**, 661 (1997). 127

[338] Schnürer, M., Spielmann, C., Wobrauschek, P., Streli, C., Burnett, N. H., Kan, C., Ferencz, K., Koppitsch, R., Cheng, Z., Brabec, T., & Krausz, F. *Phys. Rev. Lett.* **80**, 3236 (1998). 127

[339] Christov, I. P., Murnane, M. M., & Kapteyn, H. C. *Phys. Rev. Lett.* **78**, 1251 (1997). 127

REFERENCES

[340] Drescher, M., Hentschel, M., Kienberger, R., Tempea, G., Spielmann, C., Reider, G., Corkum, P., & Krausz, F. *Science* **291**, 1923 (2001). 127

[341] Hentschel, M., Kienberger, R., Spielmann, C., Reider, G. A., Milosevic, N., Brabec, T., Corkum, P., Heinzmann, U., Drescher, M., & Krausz, F. *Nature* **414**, 509 (2001). 127

[342] Kienberger, R., Goulielmakis, E., Uiberacker, M., Baltuska, A., Yakovlev, V., Bammer, F., Scrinzi, A., Westerwalbesloh, T., Kleineberg, U., Heinzmann, U., *et al. Nature* **427**, 817 (2004). 127, 152

[343] López-Martens, R., Varjú, K., Johnsson, P., Mauritsson, J., Mairesse, Y., Salières, P., Gaarde, M., Schafer, K., Persson, A., Svanberg, S., *et al. Phys. Rev. Lett.* **94**, 33001 (2005). 127

[344] Stockman, M. I., Kling, M. F., Kleineberg, U., & Krausz, F. *Nat. Photon.* **1**, 539 (2007). 127

[345] Remacle, F. & Levine, R. D. *Proc. Natl. Acad. Sci. U. S. A.* **103**, 6793 (2006). 127

[346] Uiberacker, M., Uphues, T., Schultze, M., Verhoef, A. J., Yakovlev, V., Kling, M. F., Rauschenberger, J., Kabachnik, N. M., Schroder, H., Lezius, M., Kompa, K. L., Muller, H.-G., Vrakking, M. J. J., Hendel, S., Kleineberg, U., Heinzmann, U., Drescher, M., & Krausz, F. *Nature* **446**, 627 (2007). 127

[347] Cavalieri, A. L., Muller, N., Uphues, T., Yakovlev, V. S., Baltuska, A., Horvath, B., Schmidt, B., Blumel, L., Holzwarth, R., Hendel, S., Drescher, M., Kleineberg, U., Echenique, P. M., Kienberger, R., Krausz, F., & Heinzmann, U. *Nature* **449**, 1029 (2007). 127

[348] Lein, M., Hay, N., Velotta, R., Marangos, J. P., & Knight, P. L. *Phys. Rev. A* **66**, 023805 (2002). 128

[349] Lein, M., Hay, N., Velotta, R., Marangos, J. P., & Knight, P. L. *Phys. Rev. Lett.* **88**, 183903 (2002). 128

[350] Torres, R., Kajumba, N., Underwood, J., Robinson, J., Baker, S., Tisch, J., de Nalda, R., Bryan, W., Velotta, R., Altucci, C., *et al. Phys. Rev. Lett.* **98**, 203007 (2007). 128

[351] Velotta, R., Hay, N., Mason, M. B., Castillejo, M., & Marangos, J. P. *Phys. Rev. Lett.* **87**, 183901 (2001). 128

[352] Nalda, R. d., Heesel, E., Lein, M., Hay, N., Velotta, R., Springate, E., Castillejo, M., & Marangos, J. P. *Phys. Rev. A* **69**, 031804 (2004). 128

[353] Vozzi, C., Calegari, F., Benedetti, E., Caumes, J.-P, Sansone, G., Stagira, S., Nisoli, M., Torres, R., Heesel, E., Kajumba, N., Marangos, J. P, Altucci, C., & Velotta, R. *Phys. Rev. Lett.* **95**, 153902 (2005). 128

[354] Kanai, T., Minemoto, S., & Sakai, H. *Nature* **435**, 470 (2005). 128

[355] Itatani, J., Levesque, J., Zeidler, D., Niikura, H., Pépin, H., Kieffer, J., & Corkum, P. *Nature* **432**, 867 (2004). 128

[356] Schwarz, W. H. E. *Angew. Chem., Int. Ed.* **45**, 1508 (2006). 128, 157

[357] Patchkovskii, S., Zhao, Z., Brabec, T., & Villeneuve, D. M. *Phys. Rev. Lett.* **97**, 123003 (2006). 128

[358] Lein, M. *J. Phys. B: At., Mol. Opt. Phys.* **40**, R135 (2007). 128

[359] Niikura, H., Legare, F., Hasbani, R., Bandrauk, A. D., Ivanov, M. Y., Villeneuve, D. M., & Corkum, P. B. *Nature* **417**, 917 (2002). 128

[360] Niikura, H., Legare, F., Hasbani, R., Ivanov, M. Y., Villeneuve, D. M., & Corkum, P. B. *Nature* **421**, 826 (2003). 128

[361] Alnaser, A. S., Osipov, T., Benis, E. P., Wech, A., Shan, B., Cocke, C. L., Tong, X. M., & Lin, C. D.

Phys. Rev. Lett. **91**, 163002 (2003). 128

[362] Lein, M. *Phys. Rev. Lett.* **94**, 053004 (2005). 129

[363] Baker, S., Robinson, J., Haworth, C., Teng, H., Smith, R., Chirila, C., Lein, M., Tisch, J., & Marangos, J. *Science* **312**, 424 (2006). 129, 219

[364] Baker, S., Robinson, J. S., Lein, M., Chirilă, C. C., Torres, R., Bandulet, H. C., Comtois, D., Kieffer, J. C., Villeneuve, D. M., Tisch, J. W. G., & Marangos, J. P. *Phys. Rev. Lett.* **101**, 053901 (2008). 129

[365] Niikura, H., Villeneuve, D. M., & Corkum, P. B. *Phys. Rev. Lett.* **94**, 083003 (2005). 129

[366] Corkum, P. B. *Phys. Rev. Lett.* **71**, 1994 (1993). 129, 133

[367] Kitzler, M. & Lezius, M. *Phys. Rev. Lett.* **95**, 253001 (2005). 131, 188

[368] Lewenstein, M., Balcou, P., Ivanov, M. Y., L'Huillier, A., & Corkum, P. B. *Phys. Rev. A* **49**, 2117 (1994). 132, 135, 138, 140, 187

[369] Gavrila, M. *San Diego* (1992). 134

[370] Ammosov, M., Delone, N., & Krainov, V. *Sov. Phys. JETP* **64**, 1191 (1986). 137, 140

[371] Kohler, M. C., Ott, C., Raith, P., Heck, R., Schlegel, I., Keitel, C. H., & Pfeifer, T. *Phys. Rev. Lett.* **105**, 203902 (2010). 138

[372] Bleistein, N. & Handelsman, R. *Asymptotic Expansions of Integrals.* Dover Publications (1986). 139

[373] Milošević, D. & Becker, W. *Phys. Rev. A* **66**, 063417 (2002). 140, 142, 143

[374] Popruzhenko, S. V., Korneev, P. A., Goreslavski, S. P., & Becker, W. *Phys. Rev. Lett.* **89**, 023001 (2002). 140

[375] Chipperfield, L. E. *High Harmonic Generation with Few-Cycle Pulses.* Ph.D. thesis, Imperial College London (2007). 140, 142, 225

[376] Figueira de Morisson Faria, C., Schomerus, H., & Becker, W. *Phys. Rev. A* **66**, 43413 (2002). 142, 143

[377] Gaarde, M. B., Tate, J. L., & Schafer, K. J. *J. Phys. B: At., Mol. Opt. Phys.* **41**, 132001 (2008). 144

[378] Gaarde, M. B., Salin, F., Constant, E., Balcou, P., Schafer, K. J., Kulander, K. C., & L'Huillier, A. *Phys. Rev. A* **59**, 1367 (1999). 144

[379] Schapper, F., Holler, M., Auguste, T., Zaïr, A., Weger, M., Salières, P., Gallmann, L., & Keller, U. *Opt. Express* **18**, 2987 (2010). 145

[380] Balcou, P., Sali'eres, P., L'Huillier, A., & Lewenstein, M. *Phys. Rev. A* **55**, 3204 (1997). 146

[381] Balcou, P., Dederichs, A. S., Gaarde, M. B., & L'Huillier, A. *J. Phys. B: At., Mol. Opt. Phys.* **32**, 2973 (1999). 146, 188

[382] Lindner, F., Paulus, G. G., Walther, H., Baltuška, A., Goulielmakis, E., Lezius, M., & Krausz, F. *Phys. Rev. Lett.* **92**, 113001 (2004). 147

[383] Lipson, S., Lipson, H., & Tannhauser, D. *Optical physics.* Cambridge Univ Pr (1995). 147

[384] Geissler, M., Tempea, G., & Brabec, T. *Phys. Rev. A* **62**, 033817 (2000). 147, 148

[385] L'Huillier, A., Auguste, T., Balcou, P., Carré, B., Monot, P., Salieres, P., Altucci, C., Gaarde, M., Larsson, J., Mevel, E., *et al. Journal of Nonlinear Optical Physics and Materials* **4**, 647 (1995). 147

[386] Rae, S. C., Burnett, K., & Cooper, J. *Phys. Rev. A* **50**, 3438 (1994). 148

[387] Wood, W. M., Focht, G., & Downer, M. C. *Opt. Lett.* **13**, 984 (1988). 148

REFERENCES

[388] Wood, W. M., Siders, C. W., & Downer, M. C. *Phys. Rev. Lett.* **67**, 3523 (1991). 148

[389] Kan, C., Burnett, N. H., Capjack, C. E., & Rankin, R. g. *Phys. Rev. Lett.* **79**, 2971 (1997). 148

[390] Ruchon, T., Hauri, C. P., Varjú, K., Mansten, E., Swoboda, M., López-Martens, R., & L'Huillier, A. *New J. Phys.* **10**, 025027 (2008). 148

[391] Constant, E., Garzella, D., Breger, P., Mével, E., Dorrer, C., Le Blanc, C., Salin, F., & Agostini, P. *Phys. Rev. Lett.* **82**, 1668 (1999). 148

[392] Hergott, J.-F., Kovačev, M., Merdji, H., Hubert, C., Mairesse, Y., Jean, E., Breger, P., Agostini, P., Carré, B., & Salières, P. *Phys. Rev. A* **66**, 021801 (2002). 148

[393] Lorin, E., Chelkowski, S., & Bandrauk, A. D. *New J. Phys.* **10**, 025033 (2008). 148

[394] Kita, T., Harada, T., Nakano, N., & Kuroda, H. *Appl. Opt* **22**, 12 (1983). 150

[395] Nakano, N., Kuroda, H., Kita, T., & Harada, T. *Appl. Opt.* **23**, 2386 (1984). 150

[396] White, M. G., Rosenberg, R. A., Gabor, G., Poliakoff, E. D., Thornton, G., Southworth, S. H., & Shirley, D. A. *Rev. Sci. Instrum.* **50**, 1268 (1979). 150

[397] Hemmers, O., Whitfield, S. B., Glans, P., Wang, H., Lindle, D. W., Wehlitz, R., & Sellin, I. A. *Rev. Sci. Instrum.* **69**, 3809 (1998). 150

[398] Sekikawa, T., Kosuge, A., Kanai, T., & Watanabe, S. *Nature* **432**, 605 (2004). 151, 157

[399] Tzallas, P., Charalambidis, D., Papadogiannis, N. A., Witte, K., & Tsakiris, G. D. *Nature* **426**, 267 (2003). 151, 157

[400] Muller, H. *Appl. Phys. B* **74**, 17 (2002). 152

[401] Itatani, J., Quéré, F., Yudin, G. L., Ivanov, M. Y., Krausz, F., & Corkum, P. B. *Phys. Rev. Lett.* **88**, 173903 (2002). 152

[402] Kitzler, M., Milosevic, N., Scrinzi, A., Krausz, F., & Brabec, T. *Phys. Rev. Lett.* **88**, 173904 (2002). 152

[403] Kosik, E. M., Corner, L., Wyatt, A. S., Cormier, E., Walmsley, I. A., & Dimauro, L. F. *J. Mod. Opt.* **52**, 361 (2005). 152, 153

[404] Mairesse, Y. & Quéré, F. *Phys. Rev. A* **71**, 011401 (2005). 152

[405] Goulielmakis, E., Schultze, M., Hofstetter, M., Yakovlev, V. S., Gagnon, J., Uiberacker, M., Aquila, A. L., Gullikson, E. M., Attwood, D. T., Kienberger, R., Krausz, F., & Kleineberg, U. *Science* **320**, 1614 (2008). 152

[406] Quéré, F., Itatani, J., Yudin, G. L., & Corkum, P. B. *Phys. Rev. Lett.* **90**, 073902 (2003). 152

[407] Mauritsson, J., Lpez-Martens, R., L'Huillier, A., & Schafer, K. *Opt. Lett.* **28**, 2393 (2003). 153

[408] Cormier, E., Walmsley, I. A., Kosik, E. M., Wyatt, A. S., Corner, L., & DiMauro, L. F. *Phys. Rev. Lett.* **94**, 033905 (2005). 153

[409] Mairesse, Y., Gobert, O., Breger, P., Merdji, H., Meynadier, P., Monchicourt, P., Perdrix, M., Salières, P., & Carré, B. *Phys. Rev. Lett.* **94**, 173903 (2005). 153

[410] Lee, D. G., Park, J. J., Sung, J. H., & Nam, C. H. *Opt. Lett.* **28**, 480 (2003). 153

[411] Valentin, C., Gautier, J., Goddet, J.-P, Hauri, C., Marchenko, T., Papalazarou, E., Rey, G., Sebban, S., Scrick, O., Zeitoun, P., Dovillaire, G., Levecq, X., Bucourt, S., & Fajardo, M. *J. Opt. Soc. Am. B* **25**, B161 (2008). 154

[412] Frumker, E., Paulus, G. G., Niikura, H., Villeneuve, D. M., & Corkum, P. B. *Opt. Lett.* **34**, 3026 (2009). 154

[413] Witting, T., Wyatt, A., Monmayrant, A., Walmsley, I., Haworth, C., Robinson, J., Tisch, J., &

Marangos, J. In *Conference on Lasers and Electro-Optics/Quantum Electronics and Laser Science Conference and Photonic Applications Systems Technologies*, pages 1–1 (2007). 155

[414] Wyatt, A., Witting, T., Monmayrant, A., Walmsley, I., Haworth, C., Robinson, J., Tisch, J., & Marangos, J. In *Conference on Lasers and Electro-Optics/Quantum Electronics and Laser Science Conference and Photonic Applications Systems Technologies* (2007). 155

[415] Gordon, A., Santra, R., & Kärtner, F. X. *Phys. Rev. A* **72**, 063411 (2005). 157

[416] Gordon, A. & Kärtner, F. X. *Phys. Rev. Lett.* **95**, 223901 (2005). 157

[417] Kästner, A., Grossmann, F., Schmidt, R., & Rost, J.-M. *Phys. Rev. A* **81**, 023414 (2010). 157

[418] McNeil, B. W. J., Clarke, J. A., Dunning, D. J., Hirst, G. J., Owen, H. L., Thompson, N. R., Sheehy, B., & Williams, P. H. *New J. Phys.* **9**, 82 (2007). 157

[419] Lambert, G., Hara, T., Garzella, D., Tanikawa, T., Labat, M., Carre, B., Kitamura, H., Shintake, T., Bougeard, M., Inoue, S., Tanaka, Y., Salieres, P., Merdji, H., Chubar, O., Gobert, O., Tahara, K., & Couprie, M.-E. *Nat Phys* **4**, 296 (2008). 157

[420] Press, W. H., Teukolsky, S. A., Vetterling, W. T., & Flannery, B. P. *Numerical Recipes in C, The Art of Scientific Computing.* Cambridge University Press (1992). 165

[421] Haworth, C. A. *Development and Application of an Attosecond Light Source Based on High Harmonic Generation.* Ph.D. thesis, Imperial College London (2007). 165, 166

[422] Robinson, J. S. *The generation and application of intense, few-cycle laser pulses.* Ph.D. thesis, Imperial College London (2006). 165

[423] Chipperfield, L. E., Knight, P. L., Tisch, J. W. G., & Marangos, J. P. *Opt. Commun.* **264**, 494 (2006). 188

[424] Zaïr, A., Holler, M., Guandalini, A., Schapper, F., Biegert, J., Gallmann, L., Keller, U., Wyatt, A. S., Monmayrant, A., Walmsley, I. A., Cormier, E., Auguste, T., Caumes, J. P., & Salières, P. *Phys. Rev. Lett.* **100**, 143902 (2008). 188

[425] Auguste, T., Salières, P., Wyatt, A. S., Monmayrant, A., Walmsley, I. A., Cormier, E., Zaïr, A., Holler, M., Guandalini, A., Schapper, F., Biegert, J., Gallmann, L., & Keller, U. *Phys. Rev. A* **80**, 033817 (2009). 188

[426] Kitzler, M., O'Keeffe, K., & Lezius, M. *J. Mod. Opt.* **53**, 57 (2006). 188

[427] Pfeifer, T., Gallmann, L., Abel, M. J., Neumark, D. M., & Leone, S. R. *Opt. Lett.* **31**, 975 (2006). 188

[428] Taranukhin, V. D. *J. Opt. Soc. Am. B* **21**, 419 (2004). 188

[429] Zeng, Z., Cheng, Y., Song, X., Li, R., & Xu, Z. *Phys. Rev. Lett.* **98**, 203901 (2007). 188

[430] Kim, B., An, J., Yu, Y., Cheng, Y., Xu, Z., & Kim, D. E. *Opt. Express* **16**, 10331 (2008). 188

[431] Dudovich, N., Smirnova, O., Levesque, J., Mairesse, Y., Ivanov, M. Y., Villeneuve, D. M., & Corkum, P. B. *Nat. Phys.* **2**, 781 (2006). 188, 190

[432] Pfeifer, T., Gallmann, L., Abel, M. J., Nagel, P. M., Neumark, D. M., & Leone, S. R. *Phys. Rev. Lett.* **97**, 163901 (2006). 188, 190

[433] Joosen, W., Bakker, H. J., Noordam, L. D., Muller, H. G., & van Linden van den Heuvell, H. B. *J. Opt. Soc. Am. B* **8**, 2537 (1991). 215

[434] Joosen, W., Agostini, P., Petite, G., Chambaret, J. P., & Antonetti, A. *Opt. Lett.* **17**, 133 (1992). 215

[435] Banfi, G. P., Trapani, P. D., Danielius, R., Piskarskast, A., Righini, R., & Sa'nta, I. *Opt. Lett.* **18**,

1547 (1993). 215

[436] Danielius, R., Piskarskas, A., Stabinis, A., Banfi, G. P., Trapani, P. D., & Righini, R. *J. Opt. Soc. Am. B* **10**, 2222 (1993). 215

[437] Petrov, V., Seifert, F., & Noack, F. *Appl. Phys. Lett.* **65**, 268 (1994). 215

[438] Seifert, F., Petrov, V., & Noack, F. *Opt. Lett.* **19**, 837 (1994). 215

[439] Nisoli, M., Silvestri, S. D., Magni, V., Svelto, O., Danielius, R., Piskarskas, A., Valiulis, G., & Varanavicius, A. *Opt. Lett.* **19**, 1973 (1994). 215

[440] Yakovlev, V. V., Kohler, B., & Wilson, K. R. *Opt. Lett.* **19**, 2000 (1994). 215

[441] Wilson, K. R. & Yakovlev, V. V. *J. Opt. Soc. Am. B* **14**, 444 (1997). 215

[442] Piel, J., Beutter, M., & Riedle, E. *Opt. Lett.* **25**, 180 (2000). 215

[443] Cerullo, G. & Silvestri, S. D. *Rev. Sci. Instrum.* **74**, 1 (2003). 215

[444] Toma, E. S., Antoine, P., de Bohan, A., & Muller, H. G. *J. Phys. B: At., Mol. Opt. Phys.* **32**, 5843 (1999). 215

[445] Strelkov, V. V., Mével, E., & Constant, E. *New J. Phys.* **10**, 083040 (18pp) (2008). 215

[446] Birge, J. R., Crespo, H. M., & Kärtner, F. X. *J. Opt. Soc. Am. B* **27**, 1165 (2010). 217

[447] Peebles, P. *Probability, random variables, and random signal principles*. McGraw-Hill Companies (1980). 222

[448] Haykin, S. *Communication systems*. Wiley-India (2008). 222