Imperial College London

Few-Femtosecond Deep-UV Pulses for Transient-Absorption Experiments

Malte Christian Brahms

March 2018

Thesis submitted for the degree of Doctor of Philosophy

Laser Consortium
Quantum Optics and Laser Science
Department of Physics
Imperial College London
London SW7 2BW, United Kingdom

Declaration of Originality

Except when otherwise noted or cited, all work contained in this thesis is my own.

Malte Christian Brahms

Copyright

Summary

In this thesis I describe the development, implementation and characterisation of a source of wavelength-tunable few-femtosecond laser pulses in the deep ultraviolet spectral region for use in time-resolved experiments. I also propose and model an extension of this source capable of simultaneously generating a single-cycle driving pulse for extreme nonlinear optics as well as a few-femtosecond ultraviolet pulse.

Building on advances in the field of femtochemistry, ultrafast science is moving towards ever shorter timescales and more complex systems. One of the key building blocks for the next generation of experiments studying ultrafast dynamics in molecules will be the availability of few-femtosecond pulses to directly address electronic resonances whose corresponding photon energy lies in the vacuum and deep ultraviolet spectral regions.

By harnessing the capabilities of soliton self-compression in novel micro-structured waveguides, we have generated pulses in the deep ultraviolet with energies of hundreds of nanojoules. The delivery of these pulses to an experiment as well as the measurement of their temporal profile pose significant challenges due to the dispersive properties of optical materials in the ultraviolet. We have developed an in-vacuum device for ultrafast pulse characterisation, and by directly coupling the waveguide to vacuum we were able to measure distortion-free pulses with durations below 10 fs at a range of different central wavelengths.

Numerical modelling of a scaled-up version of the apparatus shows that the self-compressed driving pulse in the ultraviolet pulse generation process can maintain its shape when delivered directly to vacuum. The single-cycle pulse duration makes it an ideal driver for extreme nonlinear optics and the generation of isolated attosecond pulses in the soft X-ray spectral region.

We do these things not because they are easy, but because we thought they were going to be easy.

- The programmer's credo

Acknowledgments

Despite the fact that only one name is on the front, this thesis was in many ways a group effort. I owe thanks to a huge number of people who helped me along the way in one respect or another.

First and foremost, I would not be able to write anything at all without the tireless effort, support, and patience I received from my parents. Not only did they put me through university in one of the most expensive places I could have chosen, they also stoically endured my studying science and consequently losing most of my ability to relate to normal people.

I started my physics degree, like so many first-years, hoping I'd be the next Stephen Hawking. Luckily, I quite quickly turned to experimental physics instead. This was made possible in large part by Jon (Marangos), who gave me something interesting to do one summer after I knocked on his door and asked for a project. Following on from that, John (Tisch) let me use his laser for final year project, and I thank him for that and many pieces of advice. After that, Jon (Marangos) gave me the opportunity to do my PhD in the Laser Consortium, and he has been a reliable source of sage advice all the way through. I look forward to more productive collaboration in the future.

This project could not have happened without the support from the guys at the Max Planck Institute for the Science of Light. In particular, I'd like to thank Francesco for answering my endless newbie questions about PCF, and John (Travers) for his help with the numerics – and for my new job, of course.

There is no telling how many more hacked-together pieces of kit I would have had to use were it not for the excellent skills of our technicians, Andy and Connor, who I'm pretty sure are wizards. And while I seemed to be the only PhD student without an

Arduino project, Susan's help with an ailing laser's electronics was invaluable, not to mention the fact that my vacuum chamber was one of the many things she acquired in her inimitable magpie-style. Judith supplied me with advice about purchase orders as well as life in general, and I already miss our morning chats.

The Laser Consortium always contains a colourful group of people, and my time here would not have been nearly as fun without them. Teatime on level 8 with a hefty helping of cynicism and rage about politics from Dan and Konstantin will forever remain a cherished memory, as will all the other rambling conversations. In Allan and Paloma, I found not only daily lunch and frequent dinner companions, but something very close to family. I hope we can keep up the biannual visits!

The H012 team has changed members so many times, I'd be certain to forget someone if I tried to list them all. I owe a debt of thanks to all of them for the help in the lab. There are some, however, who deserve a special mention. Christian was there right in the beginning, in the formative months of my PhD. He impressed upon me the importance of doing things properly instead of relying on quick botch jobs, and I will carry this wisdom with me for the rest of my life. Dane helped in so many ways, big and small, managing the lab and teaching me untold multitudes of useful bits of physics – but most importantly, he made me realise that a good understanding of Fourier transforms is all you ever really need in life. Head and shoulders above the rest, however, stands my favourite argumentative Canadian. Allan not only proof-read and critiqued the entirety of this thesis, he also stayed on weekends, hacked equipment, and fought the valiant, grinding, interminable battle against the laser with me for months on end. It is no exaggeration to say that he taught me how to be a scientist.

Lastly, I count myself very lucky to be married to someone who cares so little about physics that I have no choice but to stop thinking about work when I get home. I'm also quite sure I would have driven myself insane at some point without Abi's continued support and encouragement in the face of longer and longer days as the PhD went on. Getting married in the last year of a PhD might not seem like the smartest idea, but it's the most important thing I've ever done. Thank you for everything.

Contents

List of abbreviations

ADK Ammosov-Delone-Krainov

BBO beta barium borate

CaF_2 calcium fluoride

CEP carrier-envelope phase

DFG difference frequency generation

DUV deep ultraviolet

DWE dispersive wave emission

ePIE extended ptychographic iterative engine

FME forward Maxwell equation

FROG frequency-resolved optical gating

FTL Fourier transform limit

FWHM full width at half maximum

FWM four-wave mixing

GDD group delay dispersion

GNLSE generalised nonlinear Schrödinger equation

GVD group velocity dispersion

HHG high-harmonic generation

LBO lithium triborate

MgF_2 magnesium fluoride

NIR near infrared

NLSE nonlinear Schrödinger equation

ODE ordinary differential equation

OPA optical parametric amplifier

PBG photonic band gap

PCF photonic crystal fibre

PCGPA principal components generalised projection algorithm

PID proportional-integral-differential

PIE ptychographic iterative engine

PPT Perelomov-Popov-Terent'ev

RMS root mean square

rPIE regularised ptychographic iterative engine

SFA strong-field approximation

SFG sum frequency generation

SHG second harmonic generation

SPM self-phase modulation

SVD singular value decomposition

SWIR short-wavelength infrared

SXR soft X-ray

THG third harmonic generation

Ti:Sa titanium-doped sapphire

TOD third-order dispersion

UV ultraviolet

XANES X-ray absorption near-edge structure

XFROG cross-correlation frequency-resolved optical gating

XPM cross-phase modulation

1. Introduction

At many points in the history of science, momentous progress was made possible by new methods of measurement. The telescope advanced the field of astronomy by bringing the night sky that much closer, while the microscope has been fundamental to both biology and medicine since its invention. Fast-forward a few centuries, and the use of lasers has revolutionised measurement by providing precise rulers both in space and time.

Arguably, the history of ultrafast time-resolved measurements begins with the efforts of Eadward Muybridge who, in 1878, determined that all four of a horse's feet leave the ground at one point of a galloping gait. This was the first recorded use of the freeze-frame technique to measure the details of a process that is too fast for the human senses to follow. Since then, high-speed photography has been used to slow down a great variety of phenomena in all areas of science and technology.

In all fast measurements, there are two different ways to achieve the desired time resolution. The first, and most commonly used, is to use a very fast shutter on the detector – in the case of high-speed photography, a camera. By only collecting light for a very short amount of time, the scene is frozen in place at that moment. The second method is to keep the shutter open and instead let the illumination be the fast event – that is, to use a flash. If the scene is sufficiently dark outside the illumination time, this also yields a frozen image.

In going beyond the timescales accessible using camera shutters, which are limited to exposure times of hundreds of microseconds ($1\,\mu s = 10^{-6}\,s$), it is the second of these two approaches that has been most successful. The speed record for measurements was first pushed below a microsecond by discharge flash lamps, and then further to the

picosecond ($1\,\mathrm{ps} = 10^{-12}\,\mathrm{s}$) and femtosecond ($1\,\mathrm{fs} = 10^{-15}\,\mathrm{s}$) range by the development of pulsed lasers.

With the advent of ultrafast light sources, the timescales of the atomic and molecular dynamics involved in chemical reactions came within reach, leading to rapid advances in the field of femtochemistry [1]. Besides going well beyond the speed of events usually described as fast – such as horses, speeding bullets, and explosions – this area of physics is concerned with movement on atomic length scales. Processes such as the ejection of an electron from an atom or molecule, the making and breaking of chemical bonds, and the reorganisation of charge following an excitation all occur at extremely small scales in both time and space, and thus in the realm of quantum physics.

The new frontier in ultrafast measurement is the behaviour of electrons inside the molecules and atoms themselves, as part of chemical reactions as well as other phenomena, such as light harvesting. This requires even better time resolution of a few femtoseconds or less.

Ultrafast dynamics and the need for time-resolved measurements

The presence of behaviour on ultrafast timescales in quantum systems is largely a consequence of the superposition of states with different energies. This can be seen by considering the time-dependent Schrödinger equation, which describes general quantum systems at non-relativistic velocities:

$$-i\hbar\partial_t\Psi(\mathbf{R},t) = \left[\frac{1}{2}\hat{\mathbf{p}}^2 + V(\mathbf{R},t)\right]\Psi(\mathbf{R},t)\,, \tag{1.0.1}$$

where $\Psi(\mathbf{R},t)$ is the wavefunction, t is time, $\mathbf{R}$ is a generalised coordinate for the system, $\hat{\mathbf{p}}$ is the momentum operator, and $V(\mathbf{R},t)$ is the potential. If the potential is independent of time or if the dependence on time is slow, the wavefunction can be written as a sum over the energy eigenstates[1]:

$$\Psi(\mathbf{R},t) = \sum_n a_n\psi_n(\mathbf{R})\mathrm{e}^{-iE_nt/\hbar}\,, \tag{1.0.2}$$

[1] This sum becomes an integral when continuum states are considered as well, however the central points can be illustrated without them.

where n is the principal quantum number, a_n is a weighting coefficient, and E_n is the energy corresponding to the n$^{\text{th}}$ eigenstate's spatial wavefunction $\psi_n(\mathbf{R})$, which satisfies the time-independent Schrödinger equation:

$$\left[\frac{1}{2}\hat{\mathbf{p}}^2 + V(\mathbf{R})\right]\psi_n(\mathbf{R},t) = E_n\psi_n(\mathbf{R})\,. \tag{1.0.3}$$

Each state thus evolves with a frequency given by its energy. Therefore, if Ψ is comprised of several eigenstates, that is, if several of the coefficients a_n are non-zero, the overall state will exhibit dynamics at timescales set by the energy separation between its component eigenstates. As with a superposition of electromagnetic waves, more complex dynamics are created when a larger number of individual states are present in the wavefunction.

In many cases, the states in the superposition can be distinguished and their characteristics analysed by static spectroscopic methods. A simple example is the measurement of the excited state lifetime, for instance after excitation with light. For each excited state that is populated, the lifetime is encoded in the width in frequency of the fluorescence emitted when the system decays back to the ground state. When the system is in a superposition, the component states simply appear as different peaks in the fluorescence spectrum, and their lifetimes can be measured in the same way. This means that even for energy separations of a few electron volts, which lead to dynamics on few-femtosecond timescales, a time-resolved measurement is not necessarily required to investigate the behaviour of the system in question.

There are, however, situations in which such a static measurement cannot yield the desired information. The most obvious example is that of dark states. If the system is in such a state, it cannot decay by emitting a single photon, nor can it be promoted to a higher-lying state by absorbing one, since no electric dipole transitions are allowed; the dipole matrix element with all other states is zero [2]. It follows that such states do not appear in static fluorescence or absorption spectroscopy measurements.

Dark states and the dynamics of non-radiative transitions connected with them can play an important role in physical processes. One particularly illuminating example is that of ultraviolet (UV) absorption in DNA and RNA. To avoid damage, the energy

absorbed in these systems is internally converted in a non-radiative way in a very short time [3]. The dynamics of this process were investigated using experiments in which the RNA base uracil was excited with a short pulse of UV light, subsequently ionised by a strong laser pulse, and the relative yield of the resulting molecular fragments recorded as a function of delay between excitation and fragmentation [4]. In this way, the time-dependent population and thus decay rates of several states could be extracted, one of which is a dark state.

On picosecond and femtosecond timescales, time-resolved measurements have been used extensively to go beyond the limitations of static spectroscopy [1]. Beside the invisible presence of dark states, other issues that can be addressed in this fashion are those that make an absorption or emission spectrum difficult to interpret. For instance, the appearance of several spectral lines that overlap can preclude the extraction of line widths. More fundamentally, there are some processes in which the evolution of a particular excited state does not follow a simple exponential decay, so that the connection between lifetime and line width breaks down. One example is Auger decay of an ionised atom or molecule; here, several decay channels are partially mutually coherent, and the resulting interference leads to a more complex structure in the decay [5].

The common feature among these examples of time-resolved measurements is that the absolute value of the absorption or emission spectrum created by a process does not contain sufficient information. To characterise the dynamics, the *phase* of the spectrum has to be measured. In the absence of detectors that can access the relevant timescales directly, the only way to do this is to use two events. One event starts the process, and another investigates the state of the system some time later. By performing a time-integrated (i.e, phase-insensitive) measurement repeatedly while varying the delay between the two events, the dynamics of the process – or equivalently, the spectral phase – can be reconstructed. The event forming the "starting gun" is often termed the *pump* and the second event the *probe*, and consequently this type of measurement is known as a *pump-probe* experiment.

Several different methodologies exist for ultrafast measurements. Collecting the products of ionisation-induced molecular fragmentation, as already mentioned above, is one such method. Another very commonly used technique is that of time-resolved absorption spectroscopy, also known as transient absorption spectroscopy. Here, absorption spectra are taken at different delays after an initial excitation event, building up a picture of the dynamics [6]. Since the processes most commonly investigated using transient absorption spectroscopy result from the vibrational and rotational degrees of freedom of a molecule, the relevant transitions can be addressed using UV, visible and infrared light.

The study of ultrafast dynamics is progressing to ever faster processes, and thus to the behaviour of electrons rather than molecular rotation and vibration. The laser wavelength required to "speak to" the molecule for the excitation becomes more and more important as the relevant absorption resonances spread further apart and shift towards the deep ultraviolet (DUV) (200 nm to 300 nm) or even vacuum ultraviolet (100 nm to 200 nm). A tunable source of sufficiently short pulses in this spectral region is therefore very desirable. At the same time, the utility of transient absorption spectroscopy is greatly enhanced by moving to the X-ray spectral region, where absorption is due to resonances with deeply bound electrons, forming an excellent probe of individual atoms in a molecule. This, too, requires new tools and light sources.

Shorter pulses at shorter wavelengths

The generation of short pulses at the relevant wavelengths is one of the biggest obstacles to applying the ideas of ultrashort measurement to processes involving electronic excitation. As already mentioned, the required time resolution is of the order of a few femtoseconds or even below a femtosecond. No laser sources can currently access this regime directly, so pulse compression techniques form the foundational technology in this field.

Although extreme time resolution can be achieved by using pulse-compression techniques, they still need to overcome a second hurdle: the restriction of laser emission

to the gain bands of available laser media leaves many regions of the electromagnetic spectrum inaccessible. Frequency conversion of laser light by way of nonlinear optical processes has conquered many parts of the visible (400 nm to 700 nm) and near-infrared (700 nm to 2500 nm) spectral regions. However, the deep and vacuum ultraviolet remain difficult to reach, especially when both a short pulse as well as good conversion efficiency are required in addition to the ability to tune the wavelength of the pulse.

The goal of the work for this thesis was to address the challenge of generating and delivering bright, tunable and short pulses across the deep ultraviolet while using the minimum amount of laser pulse energy possible. The approach we used to achieve this goal is based on recent developments in the field of nonlinear optics in gas-filled hollow-core waveguides [7,8]. The unique combination of a high damage threshold and broadband low-loss guidance at small core sizes in hollow-core photonic crystal fibres has allowed many phenomena of nonlinear fibre optics, including the self-compression of laser pulses and the creation of ultraviolet secondary pulses, to be extended to unprecedented intensities. With the gas species and pressure offering a degree of control over the dispersive and nonlinear properties of the waveguide, the potential for broadband deep-UV pulse generation was realised soon after this type of fibre was first created [9,10]. After introducing the relevant concepts as well as the experimental context and requirements in chapters 2 and 3, I describe the design and implementation of an apparatus to harness this capability in chapter 4.

The other main aim of the work presented here was to measure the generated pulses in near-identical conditions to an experiment. This is only possible in vacuum, both because of the dispersive properties of optical materials and air in the ultraviolet and because soft X-ray pulses, the other half of the planned experiment, are strongly absorbed in all gases. Previous measurements of UV pulses generated in this manner were done in air and included other elements that stretched the pulses [11]. While numerical back-propagation showed that short pulses were generated in the waveguide, so far the delivery of a few-femtosecond pulse to an experiment has not been demonstrated. As a result of the vacuum requirement, the UV pulse generation had to take place in

the presence of a pressure gradient in the waveguide, a situation that has only been examined cursorily so far [10]. The pulse characterisation is detailed in chapter 5.

Finally, as a view towards the future of ultrafast light sources, in chapter 6 I propose and model a variation of the UV pulse generation scheme that allows for the simultaneous creation of two ultrashort pulses in different spectral regions at even higher pulse energy. The proposed source is predicted to be capable of improving upon the state of the art for both highly energetic few-femtosecond UV pulses and sub-cycle long-wavelength drivers for the generation of soft X-ray attosecond pulses.

Contributions

All work in this thesis was performed collaboratively. All photonic crystal fibre used in the experiments as well as the gas cells holding that fibre were designed and produced at the Max Planck Institute for the Science of Light in Erlangen, Germany. I developed the pre-compression and driving pulse diagnostics for the deep ultraviolet pulse generation apparatus, and implemented the source in the laboratory with help from Francesco Tani and Alexey Ermolov on the initial setup. The in-vacuum pulse characterisation apparatus was designed by Dane Austin and myself, and I implemented it in the laboratory as well as carrying out the data analysis. The source of soft X-ray attosecond pulses that forms the other half of the planned experiments was developed by Allan Johnson. All numerical modelling was carried out using software I wrote with help from John Travers and Dane Austin.

2. General Background

2.1. Fundamentals of ultrafast optics

The work in this thesis deals primarily with events on timescales only accessible by use of ultrashort laser pulses and with the characteristics of these pulses. To enable the discussion of ultrafast phenomena in later chapters, here I will introduce the formal description of ultrashort laser pulses and their interaction with matter at low intensities. High-intensity phenomena are discussed in sections 2.2 and 2.3.

2.1.1. Formal description of ultrashort laser pulses

The propagation of ultrashort laser pulses is governed by the equations of classical electromagnetism, known as Maxwell's equations. To illustrate the formal description of ultrashort laser pulses, it is useful to initially restrict the discussion to propagation in vacuum or dielectric materials, so that there are no free charges and currents. The magnetic field is then uniquely determined by the electric field, and it is sufficient to discuss only the latter. The electric field in a dielectric is governed by the three-dimensional electromagnetic wave equation [12]:

$$\nabla^2 \boldsymbol{\mathcal{E}} - \frac{1}{c^2}\frac{\partial^2 \boldsymbol{\mathcal{E}}}{\partial t^2} = \mu_0 \frac{\partial^2 \boldsymbol{\mathcal{P}}}{\partial t^2}\,, \tag{2.1.1}$$

where $\boldsymbol{\mathcal{E}}(\mathbf{r}, t)$ is the electric field as a function of spatial coordinate $\mathbf{r}$ and time t, similarly $\boldsymbol{\mathcal{P}}(\mathbf{r}, t)$ is the polarisation induced in the medium, $\nabla^2 = \frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2} + \frac{\partial^2}{\partial z^2}$ is the Laplace operator, μ_0 is the magnetic permeability of vacuum, and c is the speed of light in vacuum. Bold symbols indicate vector quantities. To further simplify the following discussion, we assume that the field is linearly polarised, so that we can use the scalar field $\mathcal{E}$ and polarisation $\mathcal{P}$. Unless the field is very strong, the polarisation is proportional to the field and given by

$$\mathcal{P} = \varepsilon_0 \chi^{(1)} \mathcal{E}\,, \tag{2.1.2}$$

where ε_0 is the electric permittivity of vacuum and $\chi^{(1)}$ is the linear electric susceptibility of the medium [13]. Since the polarisation $\mathcal{P}$ encapsulates the response of bound charges in the medium, this is equivalent to approximating the potential surrounding these

charges as harmonic, so that their displacement is always proportional to the applied force. Using this approximation leads to the simpler wave equation in linear media:

$$\nabla^2 \mathcal{E} - \frac{n^2}{c^2} \frac{\partial^2 \mathcal{E}}{\partial t^2} = 0\,, \tag{2.1.3}$$

where $n = \sqrt{1 + \chi^{(1)}}$ is the refractive index.

Since the electric field $\mathcal{E}(t)$ is a real-valued quantity, it can be written in terms of a complex *analytic* field $E(t)$ as

$$\mathcal{E}(x,y,z,t) = \frac{1}{2}(E(x,y,z,t) + E(x,y,z,t)^*) = \mathrm{Re}\{E(x,y,z,t)\}\,. \tag{2.1.4}$$

Using the analytic field greatly simplifies algebraic manipulation, and can be done without loss of information. In the time-frequency Fourier domain, it is equivalent to discarding negative frequency information, which is redundant because of the Hermitian symmetry of the Fourier transform of a real signal.

The wave equation can be solved by separating variables into space and time coordinates, yielding plane wave solutions for the field E:

$$E(x,y,z,t) = \mathrm{e}^{i(k_x x + k_y + k_z z - \omega t)} = \mathrm{e}^{i(\mathbf{k}\cdot\mathbf{r} - \omega t)}\,, \tag{2.1.5}$$

where k_j is the component of the *wave vector* $\mathbf{k}$ in direction j with $j = x, y, z$. Thus the magnitude of $\mathbf{k}$ is given by

$$k \equiv |\mathbf{k}| = \sqrt{k_x^2 + k_y^2 + k_z^2}\,, \tag{2.1.6}$$

and it is related to the angular frequency ω by $\omega = \frac{c}{n}k$ and to the wavelength λ by $k = 2\pi n/\lambda$. The wave vector $\mathbf{k}$ has an intuitive physical interpretation as the direction in which the plane wave component identified by $\mathbf{k}$ moves in space. In general the electric field can be written as a superposition of such plane waves:

$$E(\mathbf{r}, t) = \int a(\mathbf{k})\, \mathrm{e}^{i(\mathbf{k}\cdot\mathbf{r} - \omega t)}\, \mathrm{d}\mathbf{k}\,, \tag{2.1.7}$$

where $a(\mathbf{k})$ is the complex amplitude of the plane wave with wave vector $\mathbf{k}$ and the integral extends over all of $\mathbf{k}$-space.

By considering the transverse (x and y) spatial and temporal Fourier transform of the field $E(k_x, k_y, z, \omega)$ in a reference plane, usually chosen as $z = 0$, and comparing it to the expression in eq. (2.1.7), it can be shown that it is related to the field in *any* z-plane by

$$E(k_x, k_y, z, \omega) = E(k_x, k_y, z = 0, \omega) \mathrm{e}^{ik_z z} \,. \tag{2.1.8}$$

Therefore the field at any point in space can be obtained from

$$E(x, y, z, \omega) = \mathcal{F}^{-1}_{k_x, k_y} \left[\mathcal{F}_{x,y} \left[E(x, y, 0, \omega) \right] \mathrm{e}^{ik_z z} \right] \,, \tag{2.1.9}$$

where $\mathcal{F}_{r_j}$ and $\mathcal{F}^{-1}_{k_j}$ denote the Fourier transform along coordinate r_j and inverse Fourier transform along reciprocal coordinate k_j, respectively:

$$E(r_j, \ldots) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} E(k_j, \ldots) \, \mathrm{e}^{ik_j r_j} \mathrm{d}k_j \;\; = \mathcal{F}^{-1}_{k_j} \left[E(k_j, \ldots) \right] \tag{2.1.10}$$

$$E(k_j, \ldots) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} E(r_j, \ldots) \, \mathrm{e}^{-ik_j r_j} \mathrm{d}r_j = \mathcal{F}_{r_j} \left[E(r_j, \ldots) \right] . \tag{2.1.11}$$

The Fourier transform propagation relationship eq. (2.1.9) is especially intuitive for laser beams, for which it is easy to define a propagation direction z. To obtain the field of the laser in some plane, all that is required is to decompose the field at $z = 0$ into its transverse plane wave components by a Fourier transform and propagate each component independently according to its value of the longitudinal wave vector k_z given by $k_z = \sqrt{k^2 - k_x^2 - k_y^2}$. We can therefore consider only the transverse field $E(x, y, t)$ in a plane of interest when discussing the characteristics of a laser pulse, safe in the knowledge that we can easily find the field further along the beam path if required. Note that the choice of z as the propagation direction is arbitrary; any orientation of the coordinate system will yield the same result.

The field $E(x, y, t)$ describes an arbitrarily simple or complicated electric field distribution and its evolution in time. However, for many laser pulses, the dependence on time and space can be treated approximately independently, so that the field can be written as a product of spatial and temporal factors:

$$E(x, y, t) = E(x, y) E(t) \,. \tag{2.1.12}$$

As long as this partition is valid, that is, in the absence of *space-time coupling*, or when considering only a single point in space, we can meaningfully discuss the field as a function of time only.

The time-only field $E(t)$ is commonly represented in the frequency domain as $E(\omega)$, which is related to $E(t)$ by the Fourier transform:

$$E(t) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} E(\omega) \mathrm{e}^{-i\omega t} \, \mathrm{d}\omega = \mathcal{F}_{\omega}^{-1} \left[E(\omega) \right] \tag{2.1.13}$$

$$E(\omega) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} E(t) \mathrm{e}^{i\omega t} \, \mathrm{d}t \qquad = \mathcal{F}_{t} \left[E(t) \right] . \tag{2.1.14}$$

The field in the frequency domain is furthermore most often written in polar notation,

$$E(\omega) = A(\omega) \mathrm{e}^{i\phi(\omega)} , \tag{2.1.15}$$

which defines the *spectral amplitude* $A(\omega)$ and the *spectral phase* $\phi(\omega)$, which are both real-valued quantities.

The connection between the time and frequency domains implies that the extent of a laser pulse in time and frequency, that is, its duration Δt and bandwidth $\Delta\omega$, respectively, are related via the bandwidth theorem:

$$\Delta\omega \Delta t \geq K , \tag{2.1.16}$$

where the value of K depends on the shape of the pulse, and also on the definition of Δt and $\Delta\omega$. The most common measure in ultrafast optics is the full width at half maximum (FWHM) of the intensity and spectral energy density, which are proportional to $|E|^2$ and $|A(\omega)|^2$, respectively. For very complicated pulses, e.g. those with many pre- and post-pulses, this definition becomes less useful. Unless otherwise stated, all pulse durations given in this thesis are defined as the FWHM of the intensity.

The pulse in the time domain can also be written in polar notation similar to eq. (2.1.15). However, the phase of $E(t)$ will contain a strong linear component, which corresponds to the central frequency of the pulse. It is therefore useful to partition the pulse into an envelope and carrier:

$$E(t) = E_0(t) \, \mathrm{e}^{-i\phi(t)} \, \mathrm{e}^{-i\omega_0 t} , \tag{2.1.17}$$

where $E_0(t)$ and $\phi(t)$ are both real-valued quantities and vary slowly over the duration of the pulse, and $E_0(t) > 0$. The (negative) time derivative of the total phase $\phi_T = \omega_0 t + \phi(t)$ gives the instantaneous frequency. Therefore the time derivative of $\phi(t)$ is the variation of the carrier frequency around ω_0 over the course of the pulse. By removing the fast-oscillating carrier frequency term in eq. (2.1.17), the *baseband signal* of the pulse is obtained. In the frequency domain, this is centred on $\omega = 0$ instead of ω_0.

The slowly varying parts of the pulse are also often combined into a single complex envelope $A(t)$, which is the central quantity of interest in propagation methods using the slowly evolving envelope approximation (see section 3.2).

2.1.2. Polynomial phases

In most cases, the interaction between light and matter depends on the frequency of the light. It is therefore often very useful to use $E(\omega)$ instead of $E(t)$ when discussing or analysing light-matter interactions. To describe the characteristics of an ultrafast laser pulse in the frequency domain, the spectral phase $\phi(\omega)$ is often expanded in a Taylor series around a central frequency ω_0:

$$\phi(\omega) = \sum_{n=0}^{\infty} \frac{1}{n!} \frac{\partial^n \phi(\omega)}{\partial \omega^n}\bigg|_{\omega_0} (\omega - \omega_0)^n = \sum_{n=0}^{\infty} \frac{1}{n!} \phi_n(\omega_0)\,(\omega - \omega_0)^n\,, \tag{2.1.18}$$

where we have defined the n^{th} order phase $\phi_n(\omega) \equiv \partial_\omega^n \phi(\omega)$.

These phases carry intuitive meaning. The 0^{th}-order phase $\phi_0(\omega_0)$ is usually called the carrier-envelope phase (CEP) and gives the offset between the peak of the electric field and that of the envelope of the pulse. $\phi_1(\omega_0)$ is the *group delay* at the central frequency and simply encodes the arrival time of the whole pulse[1]. If all other $\phi_n(\omega_0)$ are zero, the pulse is said to be at the Fourier transform limit (FTL). Transform-limited pulses have the shortest duration for a given spectral amplitude $A(\omega)$. The 2^{nd}-order phase $\phi_2(\omega_0)$, known as group delay dispersion (GDD), is the lowest-order term that influences the shape of the pulse envelope in the time domain. $\phi_2(\omega)$ is the derivative

[1] With some added group delay ϕ_1, the pulse in the domain becomes $E'(t) = \mathcal{F}_\omega^{-1}\left[E(\omega)e^{i\phi_1(\omega-\omega_0)}\right] = e^{-i\omega_0\phi_1} E(t-\phi_1)$ by the shifting property of the Fourier transform, i.e., it is simply shifted.

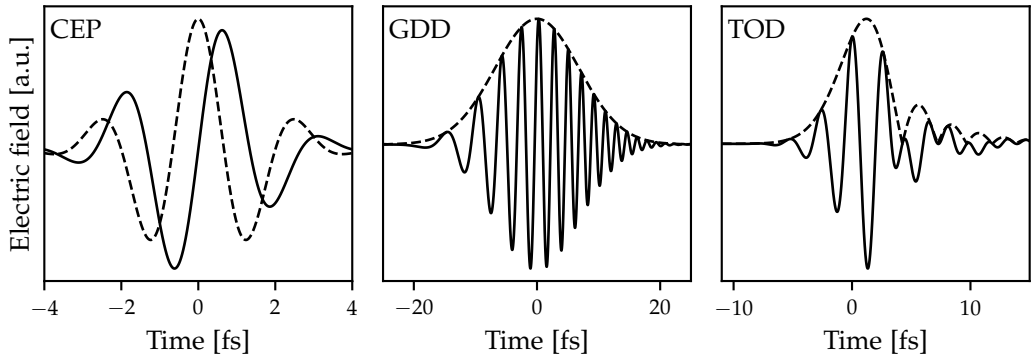

Figure 2.1.: Demonstration of three important orders of polynomial phase. In each case the transform-limited pulse is a Gaussian envelope with a carrier frequency corresponding to a wavelength of 800 nm and a FWHM duration of 2.5 fs. Left: the difference between a CEP of $\pi/2$ (solid) and 0 (dashed). Centre: the pulse with 10 fs^2 of GDD added. Right: the pulse with 10 fs^3 of TOD added.

with respect to frequency of the group delay $\phi_1(\omega)$, so if the GDD is non-zero, the group delay is frequency-dependent and different frequencies arrive at different times. The pulse is therefore longer than its FTL, and furthermore the instantaneous frequency of the field changes over the course of the pulse. This variation of the instantaneous frequency is known as *chirp*. Note that in the case of $\phi_2 > 0$, higher frequency (blue) light arrives after lower frequency (red) light, and vice versa in the case of $\phi_2 < 0$. The next term in the Taylor series is third-order dispersion (TOD). This causes further stretching of the pulse and leads to the formation of pre- or post-pulses depending on the sign of $\phi_3(\omega_0)$. Higher-order phases are often present, especially for complicated pulses. However, their interpretation is much less straightforward.

2.1.3. Dispersion

An important consideration when working with ultrafast laser pulses is how they are affected by propagation through material, for instance optical elements, the air in the

laboratory or other gases. As shown in eqs. (2.1.2) and (2.1.3), at low intensity this is determined by the refractive index n. However, in the interest of illustrating the central points, eq. (2.1.2) omits a crucial fact. The response of the material to incident radiation as encoded in the linear susceptibility is not instantaneous, so $\chi^{(1)}$ is time-dependent: $\chi^{(1)} = \chi^{(1)}(t)$. To take into account the response as induced by the electric field at each point in time, the (analytic) polarisation P is calculated as

$$P(t) = \varepsilon_0 \int_{-\infty}^{t} \chi^{(1)}(t - t')E(t')\,\mathrm{d}t' \,. \tag{2.1.19}$$

Since this is a convolution[2] of the response function $\chi^{(1)}(t)$ with the electric field $E(t)$, the convolution theorem means that it is a multiplication when expressed in the frequency domain:

$$P(\omega) = \varepsilon_0 \chi^{(1)}(\omega)E(\omega) \,. \tag{2.1.20}$$

Consequently, the refractive index $n = n(\omega)$ is also frequency-dependent, a fact that is commonly referred to as *dispersion*. Using the formalism developed in section 2.1.1, we can easily compute the effect of dispersion has on an ultrafast laser pulse as it traverses a medium. For simplicity, we consider only the on-axis component of the laser beam, that is, the one for which $k_x = k_y = 0$ and therefore $k_z = k$. Applied to the propagation through a slab of material of length L with a refractive index $n(\omega)$, eq. (2.1.8) then becomes

$$E(L, \omega) = E(0, \omega)\, \mathrm{e}^{ik_0 L n(\omega)} \tag{2.1.21}$$

where $k_0 = \omega/c$ is the magnitude of the wave vector in vacuum. Comparison to eq. (2.1.15) shows that propagation in this manner simply adds to the spectral phase $\phi(\omega)$ a contribution $k_0 L n(\omega)$. This means that propagation of a pulse through several dispersive elements can easily be described by simply adding their phase contributions one by one, and furthermore that the order in which they are traversed is irrelevant.

It is often useful to employ the same expansion to the on-axis wave vector as to the phase $\phi(\omega)$, so that individual polynomial phase terms can be added directly. The

[2] Due to causality, $\chi^{(1)}(t) = 0$ for $t < 0$. Therefore, the upper limit of the integral in eq. (2.1.19) can be extended to infinity.

on-axis wave vector is commonly given the symbol $\beta(\omega)$, and its expansion reads:

$$\beta(\omega) = \sum_{n=0}^{\infty} \frac{1}{n!} \frac{\partial^n \beta(\omega)}{\partial \omega^n}\bigg|_{\omega_0} (\omega - \omega_0)^n = \sum_{n=0}^{\infty} \frac{1}{n!} \beta_n(\omega_0)\,(\omega - \omega_0)^n\,, \tag{2.1.22}$$

where, similarly to before, $\beta_n(\omega) = \partial_\omega^n \beta(\omega)$. The coefficients β_n are closely connected to those of the spectral phase, ϕ_n. For instance, β_1 is simply the group delay *per unit length*, whose inverse is called the *group velocity* v_g, $\beta_1 = v_\mathrm{g}^{-1}$. Consequently, β_2 is GDD per unit length, and is known as group velocity dispersion (GVD).

The GVD β_2 presents one of the biggest practical challenges in ultrafast optics. This is because for the vast majority of optical glasses, β_2 is positive for frequencies in the visible spectral region and up to wavelengths of around 1500 nm[3] [14]. This is known as *normal* or positive dispersion. Therefore, once a laser pulse has positive ϕ_2, it is very difficult to remove this phase contribution and re-compress the pulse back towards the FTL.

To introduce *anomalous* (negative) dispersion, one option is to use geometrical dispersion, such as in grating and prism compressors [15]. Another method is to use dispersive mirrors, also known as chirped mirrors. These are coated with dielectric layers in such a way that longer wavelengths penetrate deeper into the coating before being reflected, thus delaying them relative to shorter wavelengths and inducing negative GDD. Owing to their complex structure, these optics are very expensive and their surface quality can be poor [15]. The difficulty of compensating GDD is one of the central reasons that ultrafast optics is done nearly exclusively with reflective elements, since the GDD induced by a reflection is negligible compared to transmission through a bulk optic.

It is important to note that the refractive index is not necessarily a purely real quantity, with any imaginary part representing absorption in the medium. Adding an imaginary component $i\,\alpha(\omega)/2$ to $n(\omega)$ in eq. (2.1.21) leads to

$$E(L,\omega) = E(0,\omega)\,\mathrm{e}^{iL\beta(\omega)}\mathrm{e}^{-\frac{1}{2}\alpha(\omega)L}, \tag{2.1.23}$$

[3]The opposite is true for frequencies further in the infrared.

which means that the intensity of the pulse decreases as

$$I(L,\omega) = I(0,\omega)\,\mathrm{e}^{-\alpha(\omega)L}\,. \tag{2.1.24}$$

For transmission of a laser pulse through glass, α is usually small enough to neglect. However, strong variations of α with frequency can lead to significant reshaping of the spectral amplitude of the laser pulse, and therefore change the temporal profile as well.

2.2. Nonlinear optics

In the discussion in section 2.1, we explicitly assumed that the electric field was weak, leading to linear optical propagation. One consequence of the strictly linear polarisation is that no new frequencies can be created, since linear optics only describes phase shifts and attenuation of already existing frequency components. To convert light from one frequency to another, *nonlinear* optics have to employed.

2.2.1. Nonlinear polarisation

If a strong electric field is applied to a dielectric medium, the effect is analogous to that of a spring being pushed or pulled with great force. As bound charges are driven further away from their equilibrium distribution, it becomes less and less accurate to approximate the potential surrounding them as harmonic. Instead of the simple proportionality of eq. (2.1.2), the polarisation is then given by a more general function of the electric field. Since the nonlinear contributions are usually much weaker than the linear one, this function is commonly expressed as a power series in the field $\mathcal{E}$:

$$\mathcal{P}(t) = \varepsilon_0 \int_{-\infty}^{t} \chi^{(1)}(t-t')\mathcal{E}(t')\,\mathrm{d}t' + \chi^{(2)}\mathcal{E}(t)^2 + \chi^{(3)}\mathcal{E}(t)^3 + \cdots , \tag{2.2.1}$$

where $\chi^{(2)}$ and $\chi^{(3)}$ are the second and third-order susceptibilities, respectively. We will limit our discussion to these two nonlinear orders, since higher orders are usually very weak. Furthermore, we will only discuss the electronic nonlinear response which is extremely fast, with response times well below a femtosecond [16], and we can therefore treat $\chi^{(2)}$ and $\chi^{(3)}$ as time-independent. While other parts of the medium response exhibit nonlinearity with longer response times, for instance the vibrational and rotational Raman effects [13], they are not important for the techniques and effects discussed in this thesis.

It is necessary here to use the real field $\mathcal{E}$ once again, since the real part of a product of two complex quantities is not simply the product of the real parts, and so $E^N \neq \mathcal{E}^N$. We can express the nonlinear polarisation as an analytic quantity, however. To see this,

we express the real driving field $\mathcal{E}$ in eq. (2.2.1) with the analytic field E as in eq. (2.1.4):

$$\mathcal{P}_2(t) = \varepsilon_0 \chi^{(2)} \left[\frac{1}{2} (E(t) + E(t)^*) \right]^2 \tag{2.2.2}$$

$$= \frac{\varepsilon_0 \chi^{(2)}}{4} \left(|E(t)|^2 + E(t)^2 \right) + \text{c.c.}\,. \tag{2.2.3}$$

Comparing this expression to eq. (2.1.4), the analytic second-order polarisation P_2 is found as

$$P_2(t) = \frac{\varepsilon_0 \chi^{(2)}}{2} \left(|E(t)|^2 + E(t)^2 \right). \tag{2.2.4}$$

Proceeding similarly for the third-order polarisation $P_3(t)$, we arrive at

$$P_3(t) = \frac{\varepsilon_0 \chi^{(3)}}{4} \left(3|E(t)|^2 E(t) + E(t)^3 \right). \tag{2.2.5}$$

The strength of the nonlinear response varies from material to material, and depends on pressure in the case of gases. This is captured in varying magnitudes of the nonlinear susceptibilities $\chi^{(2)}$ and $\chi^{(3)}$. If $\chi^{(2)}$ is non-zero, this implies that the response of the medium is asymmetric, i.e. $P(-E) \neq -P(E)$. This can only occur in materials that are *non-centrosymmetric* in nature. Most materials are centrosymmetric, including all gases and liquids, and only some crystals exhibit a non-zero value of $\chi^{(2)}$. Third-order effects, on the other hand, can be observed in all materials [13]. It is important to note that the nonlinear susceptibility is generally not a scalar quantity, but a tensor. This is especially critical for the second-order susceptibility $\chi^{(2)}$, since it only arises in necessarily anisotropic crystals. In short, this leads to the *effective nonlinearity* changing as the angle between the field polarisation and the crystal axis (or axes) varies [17].

The effects of the nonlinear response depend on whether the driving field is composed of a single pulse or several. We will first consider the effects of nonlinear interaction on a single laser pulse, and subsequently on combinations of laser pulses.

2.2.2. Harmonic generation and single-pulse Kerr effect

The first term in eq. (2.2.4) represents a constant offset of the charge density in the material for a plane wave, and a slowly varying one for a single laser pulse. This

is known as *optical rectification*. While it has applications in the generation of THz radiation [18], it is not important for the work in this thesis and we will not discuss it further. The effects of the second terms in eqs. (2.2.4) and (2.2.5) can be seen by considering one plane wave component: since

$$\left(\mathrm{e}^{-i\omega t}\right)^N = \mathrm{e}^{-iN\omega t}\,, \tag{2.2.6}$$

the polarisation contains plane wave components at N times the original frequency. This process is referred to as *harmonic generation*, specifically second harmonic generation (SHG) for $N = 2$ and third harmonic generation (THG) for $N = 3$.

For a broadband pulse, something very similar occurs. Considering the pulse in the time domain as an envelope and a carrier as shown in eq. (2.1.17), the harmonic generation terms lead to

$$P_{\mathrm{N}}(t) = \left(E_0(t)\,\mathrm{e}^{-i\phi(t)}\,\mathrm{e}^{-i\omega_0 t}\right)^N = E_0(t)^N \mathrm{e}^{-iN\phi(t)}\mathrm{e}^{-iN\omega_0 t}\,. \tag{2.2.7}$$

This is simply a pulse at N times the original frequency. However, since the amplitude envelope $E_0(t)$ is also raised to the power N, the pulse of the harmonic is shorter than that of the driving field. Correspondingly, the frequency bandwidth increases[4]. Harmonic generation was the first nonlinear optical frequency conversion mechanism to be discovered [19], and it drives some of the most prominent applications of nonlinear optics.

The nonlinear refractive index

The effects of the first term in the third-order response in eq. (2.2.5) are somewhat less straightforward. The crucial fact is that while it involves the third power of the field overall, it is not shifted in frequency. This part of the third-order polarisation oscillates at the original frequency, and can therefore influence the driving pulse itself. This is known as the *Kerr effect*.

[4] This can be understood as the result of the convolution of the pulse in frequency space with itself $N-1$ times, each time broadening the spectrum and shifting it by ω_0.

Going back to the electromagnetic wave equation (in its analytic form),

$$\nabla^2 E - \frac{1}{c^2}\frac{\partial^2 E}{\partial t^2} = \mu_0 \frac{\partial^2 P}{\partial t^2}\,, \tag{2.2.8}$$

and introducing the polarisation P consisting only of the linear part and the Kerr effect,

$$P = P_{\text{L}} + P_{\text{Kerr}} = \varepsilon_0 \chi^{(1)} E + \frac{3}{4}\varepsilon_0 \chi^{(3)} |E|^2 E\,, \tag{2.2.9}$$

we find that there is an additional contribution to the refractive index n as defined in eq. (2.1.3), since there is another term present that multiplies E. Note that we have ignored dispersion (that is, the time dependence of $\chi^{(1)}$) to keep the expressions simple, however the result is the same when we include it. The *effective* refractive index is given by

$$n_{\text{eff}} = \sqrt{1 + \chi^{(1)} + \frac{3}{4}\chi^{(3)}|E|^2} \tag{2.2.10}$$

$$= \sqrt{n^2 + \frac{3}{4}\frac{2}{c\varepsilon_0}\chi^{(3)} I}\,, \tag{2.2.11}$$

where we have used the fact that the intensity of the field is given by $I = c\varepsilon_0 |E|^2/2$ [12]. Since the nonlinear response of the medium will generally be much weaker than the linear one, we can approximate the square root to first order:

$$n_{\text{eff}} \approx n + \frac{3}{4c\varepsilon_0 n^2}\chi^{(3)} I = n + n_2 I\,, \tag{2.2.12}$$

which defines the *nonlinear refractive index* n_2. The Kerr effect thus leads to an increase of the refractive index in proportion to the intensity of the field. Since the intensity of a laser pulse varies over time as well as across the laser beam profile, this can significantly reshape a pulse both in time and in space.

Self-focusing

As demonstrated in section 2.1.1, a laser pulse accumulates phase as it propagates in z, given by $\phi = k\,z$, with $k = k_0 n$ where $k_0 = \omega/c$. For an intense laser pulse, the refractive index varies across the (non-uniform) beam profile, so that the phase ϕ has a spatial dependence as

$$\phi(x,y) = k_0 z\, n_{\text{eff}}(x,y) = k_0 z[n + n_2 I(x,y)]\,. \tag{2.2.13}$$

For the vast majority of materials, n_2 is positive. Furthermore, for most laser beams the intensity is highest in the centre and drops off towards the edges. The effective refractive index is therefore *larger* in the centre of the beam. Contrasting this with propagation in vacuum, it is equivalent to a situation in which the centre of the beam passes through a slab of material (with $n > 1$) which is thicker in the centre than at the edges. Such an object is usually referred to as a converging lens. The Kerr effect leads to the laser beam *self-focusing* inside the nonlinear material. Moreover, the intensity of the laser beam increases as the beam size decreases. Nonlinear effects therefore become even stronger, leading to more rapid self-focusing in a runaway process, which can lead to beam breakup and material damage.

This has wide-ranging implications for the design of lasers and other optical systems. The design principle of chirped-pulse amplification systems, for instance, is based on the need to avoid catastrophic damage to the laser optics through nonlinear self-focusing when amplifying a laser pulse. The Kerr effect also means that a short high-energy laser pulse cannot be transmitted through thick optics such as lenses, and reflective optics have to be used instead.

Self-phase modulation

In the time domain, the presence of the Kerr effect means that the effective refractive index varies over the course of the laser pulse. As with self-focusing, the effects of this can be illustrated by calculating the accumulated phase. Considering only the time-dependent field at different planes along z, the pulse after transmission of some distance L through a nonlinear material is given by

$$E(z,t) = E(0,t)\,\mathrm{e}^{ik_0 n_{\mathrm{eff}} L} = E_0(t)\,\mathrm{e}^{-i[\omega_0 t+\phi(t)]}\,\mathrm{e}^{ik_0 L[n+n_2 I(t)]}\,. \tag{2.2.14}$$

As described in section 2.1.1, the negative time derivative of the total phase gives the instantaneous frequency:

$$\omega(t) = -\frac{\partial \phi_{\mathrm{T}}}{\partial t} = \omega_0 + \frac{\partial \phi}{\partial t} - k_0 L n_2 \frac{\partial I(t)}{\partial t}\,, \tag{2.2.15}$$

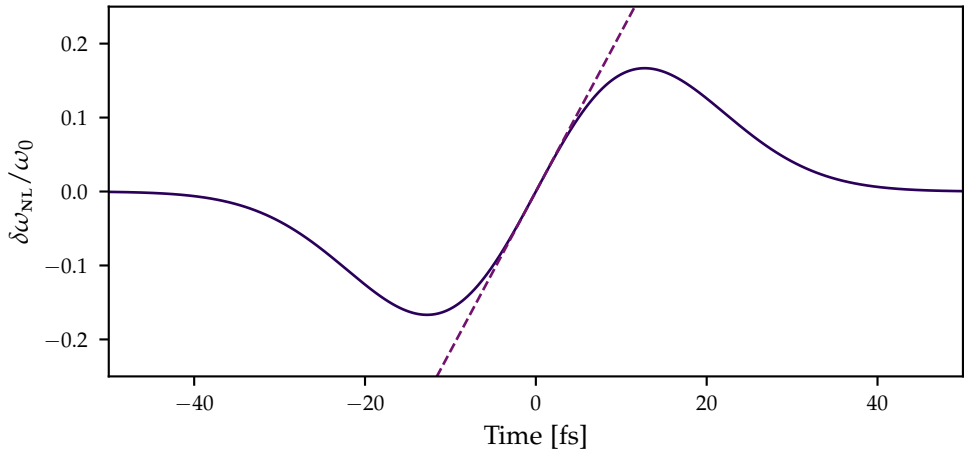

Figure 2.2.: Frequency shift $\delta\omega_{\mathrm{NL}}$ (in units of the fundamental frequency ω_0) induced by self-phase modulation (SPM) for a Gaussian pulse of 30 fs FWHM duration centred at 790 nm and a peak intensity of $3.5 \times 10^{12}\,\mathrm{W\,cm^{-2}}$ after propagating through 1 mm of material with a nonlinear refractive index n_2 of $3 \times 10^{-16}\,\mathrm{cm^2\,W^{-1}}$, a typical value for optical glass. The dashed line shows the linear approximation according to eq. (2.2.18).

which shows that the nonlinear interaction changes the instantaneous frequency by an amount

$$\delta\omega_{\mathrm{NL}}(t) = -k_0 L n_2 \frac{\partial I(t)}{\partial t}\,. \tag{2.2.16}$$

The rising (or leading) edge of a laser pulse is thus red-shifted, whereas the falling (or trailing) edge is blue-shifted. This effect is known as *self-phase modulation* (SPM). For an unchirped pulse, this means that the bandwidth of the pulse expands while the temporal profile remains the same. Figure 2.2 shows the frequency shift $\delta\omega_{\mathrm{NL}}(t)$ for a pulse with a Gaussian profile in time. While the overall shift is not linear – which would correspond to purely quadratic spectral phase, and is easiest to compensate – the shift around the peak of the pulse can be approximated as such. This can be seen by considering times close to the peak of the pulse, where the Gaussian profile can be

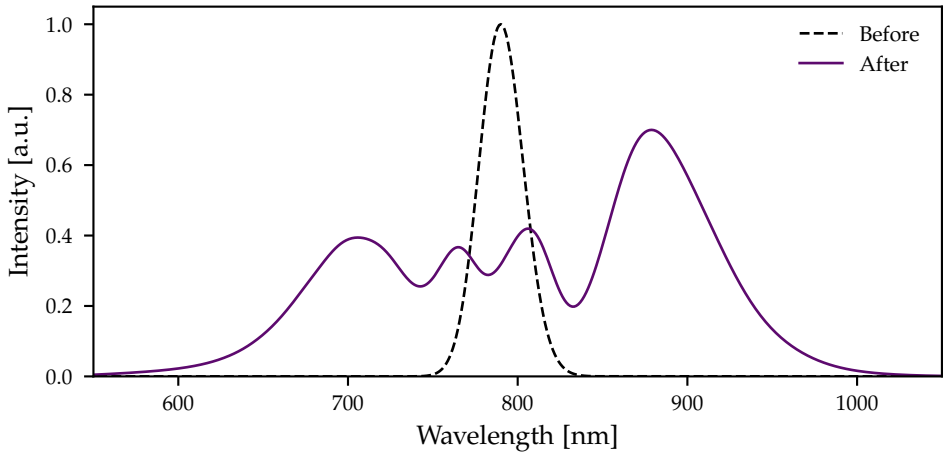

Figure 2.3.: Spectrum of a Gaussian pulse after undergoing SPM in the same conditions as in fig. 2.2. The dashed line shows the spectrum of the initial pulse (scaled down for visibility). The effect of self-steepening is visible in the asymmetry of the spectrum.

expressed as a Taylor expansion to first order:

$$I(t) = I_0 \, \mathrm{e}^{-(t/\tau)^2} \approx I_0 \left[1 - \left(\frac{t}{\tau} \right)^2 \right], \tag{2.2.17}$$

where τ is the Gaussian half-width pulse duration, so that the frequency shift is given by

$$\delta\omega_{\mathrm{NL}}(t) \approx 2 I_0 k_0 L n_2 \frac{t}{\tau^2}, \tag{2.2.18}$$

which is linear in time.

The spectral broadening effect of SPM is the foundation of many pulse compression schemes, since a larger bandwidth allows for a shorter minimum pulse duration. Importantly, after undergoing SPM the pulse is positively chirped, so that anomalous dispersion has to be added to compress the newly broadened spectrum into a shorter pulse. As mentioned in section 2.1.3, this is often challenging, requiring the use of special optics.

A typical spectrum of a Gaussian pulse after undergoing SPM is shown in fig. 2.3. The oscillatory structure is due to the fact that the same value of $\delta\omega_{\mathrm{NL}}(t)$ occurs at two

different times t as can be seen in in fig. 2.2. There are thus two components of the pulse at the same frequency but different phases, and the interference between these is seen as fringes in the frequency domain.

Given the symmetric spectral intensity distribution of the initial Gaussian pulse, it would be natural to expect the spectrum to remain symmetric after SPM. The evident asymmetry is caused by an effect usually called *self-steepening*. The intensity dependent refractive index shifts not only the instantaneous frequency, but also the group velocity of the pulse. The group velocity, being inversely proportional to the refractive index, is lower for higher intensities since $n_2 > 0$ for most materials. The peak of the pulse in the time domain therefore moves at a slower speed than the wings, and the trailing edge of the pulse becomes steeper as a consequence. This asymmetry is reflected in the spectral domain. Since SPM blue-shifts the trailing edge of the pulse, self-steepening leads to stronger spectral broadening in this part of the spectrum, so that the total spectral energy density extends further towards higher frequencies but is larger at lower frequencies. A more thorough treatment of self-steepening in waveguides can be found in section 3.2.

2.2.3. Frequency mixing and cross-pulse Kerr effect

If the driving field consists of several pulses, the nonlinear polarisation has dramatic effects. Whereas in linear optics each frequency component propagates independently, nonlinear optics leads to the mixing of frequencies. This is already evident in the temporal compression of a laser pulse upon harmonic generation. However, considering multiple pulses in the driving field explicitly illustrates the power of nonlinear optics in creating light at a desired frequency.

Sum and difference frequency generation

A field consisting of two pulses with different frequencies ω_1 and ω_2, with $\omega_1 > \omega_2$, can be expressed in the time domain as

$$E(t) = A_1(t)\,\mathrm{e}^{-i\omega_1 t} + A_2(t)\,\mathrm{e}^{-i\omega_2 t}\,, \tag{2.2.19}$$

where $A_1(t)$ and $A_2(t)$ are the slowly varying complex envelopes of the two pulses. The second-order polarisation according to eq. (2.2.4) (but omitting the common factor of $\varepsilon_0\chi^{(2)}/2$) is then given by

$$\begin{aligned} P_2(t) =& |A_1|^2 + |A_2|^2 + A_1\,\mathrm{e}^{-i2\omega_1 t} + A_2\,\mathrm{e}^{-i2\omega_2 t} \\ &+ 2A_1A_2^*\,\mathrm{e}^{-i(\omega_1-\omega_2)t} + 2A_1A_2\,\mathrm{e}^{-i(\omega_1+\omega_2)t}\,, \end{aligned} \tag{2.2.20}$$

where we have taken the complex conjugate of one of the cross-terms to combine the expressions[5]. The first four terms are the same as for the single-pulse case, representing optical rectification and second harmonic generation, respectively. The cross-terms, on the other hand, represent pulses at the difference and sum of the two original frequencies. Appropriately, these effects are known as difference frequency generation (DFG) and sum frequency generation (SFG).

Cross-phase modulation and four-wave mixing

Calculating the third-order polarisation induced by two pulses according to eq. (2.2.5) yields 12 terms. These can be categorised depending on their central oscillation frequencies. Firstly, all terms from the single-pulse case are reproduced, once for each pulse. Secondly, there are two terms that mix the amplitudes, but not frequencies, of the two fields:

$$P_{1\to2} = 6\,|A_1|^2 A_2 \mathrm{e}^{-i\omega_2 t} \tag{2.2.21}$$

$$P_{2\to1} = 6\,|A_2|^2 A_1 \mathrm{e}^{-i\omega_1 t}\,. \tag{2.2.22}$$

These terms describe cross-phase modulation (XPM), which is the result of the nonlinear shift to the refractive index induced by the respective other pulse. Note that the coefficient before these terms is twice that of SPM. Therefore, for a given intensity, the effect of a second pulse is stronger than that of a pulse on itself. This is an important consideration in ultrafast measurement.

The remaining terms in the nonlinear polarisation contain oscillations at frequencies $2\omega_1 \pm \omega_2$ and $2\omega_2 \pm \omega_1$. Note that since the interaction is of third order, the most general

[5] This does not change the real field $\mathcal{P}$ since $\mathrm{Re}[A] = \mathrm{Re}[A^*]$ and $\mathrm{Re}[A+B] = \mathrm{Re}[A] + \mathrm{Re}[B]$.

case involves the mixing of *three* different waves at frequencies ω_1, ω_2 and ω_3, and the resulting polarisation contains terms at all combinations of $\omega_4 = \pm\omega_1 \pm \omega_2 \pm \omega_3$. The effect described by these terms is thus known as four-wave mixing (FWM), since four fields are involved – three inputs and one output.

2.2.4. Phase-matching

The discussion of nonlinear frequency mixing effects so far has been based purely on the terms present in the nonlinear polarisation. This gives an overview of which frequency conversion processes are possible given a set of interacting pulses and a nonlinearity of a given order. However, we have only treated one point in space. To illustrate how propagation affects the situation, we can examine the process of difference frequency generation as an example.

Considering the total electric field as a superposition of several collinear discrete frequency components for simplicity, each component E_j at frequency ω_j is given by

$$E_j(t) = A_j \, \mathrm{e}^{i\left(k_j z - \omega_j t\right)} \,, \tag{2.2.23}$$

with $k_j = n(\omega_j)\omega_j/c$. Together, all of these terms have to obey the nonlinear wave equation eq. (2.2.9). Since we are only considering discrete frequency components, the wave equation can be split into several equations, one for each field at frequency ω_j with only the parts of the polarisation P which also oscillate at ω_j on the right hand side. For DFG between two fields at ω_1 and ω_2, such that $\omega_3 = \omega_1 - \omega_2$, the corresponding term of P is

$$P_{\mathrm{DFG}} = \frac{1}{2}\varepsilon_0 \chi^{(2)} \, 2A_1 A_2^* \mathrm{e}^{i[(k_1 - k_2)z - (\omega_1 - \omega_2)t]} \,. \tag{2.2.24}$$

There are corresponding terms for the inverse process, as well as for SFG between the two fundamental fields E_1 and E_2 and other combinations of all fields present. Considering solely on-axis propagation along z, only the derivatives with respect to z and t remain in the wave equation. Furthermore, upon inserting E_3 into the wave equation, we can neglect the term proportional to $\partial_z^2 A_3$, making the approximation that the amplitude of E_3 varies slowly with z. The result is that the evolution of A_3 with z is

given by

$$\frac{\partial A_3}{\partial z} = i\frac{\mu_0\varepsilon_0}{2}\frac{\omega_3^2}{k_3}\chi^{(2)}A_1A_2^*\,\mathrm{e}^{i(k_1-k_2-k_3)z} \tag{2.2.25}$$

$$= i\frac{\chi^{(2)}}{2c^2}\frac{\omega_3^2}{k_3}A_1A_2^*\,\mathrm{e}^{i\Delta kz}\,, \tag{2.2.26}$$

where we have defined the *wave-vector mismatch* $\Delta k \equiv k_1 - k_2 - k_3$. This is known as a *coupled wave equation*, because it expresses the coupling between the various fields present. It can be integrated along z directly, assuming that A_1 and A_2 stay approximately constant. This is known as the *no-depletion approximation*, and is valid as long as the total amount of energy converted to the field E_3 remains small [13]. If we consider the evolution of the DFG field from $z = 0$ to some length L (e.g. the thickness of a nonlinear medium), we obtain

$$A_3(L) = i\frac{\chi^{(2)}}{2c^2}\frac{\omega_3^2}{k_3}A_1A_2^*\int_0^L \mathrm{e}^{i\Delta kz}\mathrm{d}z \tag{2.2.27}$$

$$= i\frac{\chi^{(2)}}{2c^2}\frac{\omega_3^2}{k_3}A_1A_2^*\,L\,\mathrm{e}^{i\Delta kL}\,\mathrm{sinc}\left(\frac{\Delta kL}{2}\right), \tag{2.2.28}$$

where $\mathrm{sinc}(x) = \sin(x)/x$. The intensity of the component at ω_3 is therefore given by

$$I_3 \propto |A_3|^2 \propto I_1I_2L^2\,\mathrm{sinc}^2\left(\frac{\Delta kL}{2}\right). \tag{2.2.29}$$

The presence of the factor L^2 signifies that this expression describes coherent build-up of the DFG signal. The two fundamental intensities I_1 and I_2 play equal roles, which can be an important advantage in pulse measurement – if one of the two pulses is very weak, the nonlinear signal can be increased by raising the intensity of the other. The last factor is known as the *phase-matching factor* or the *phase-matching efficiency* η and is caused by the dispersion of the medium.

As can be seen in fig. 2.4, the phase-matching term is maximum at $\Delta k = 0$ and lower everywhere else. For efficient frequency conversion, the phase-matching condition is therefore given by

$$k_3 = k_1 - k_2\,. \tag{2.2.30}$$

There are two ways of understanding this criterion. In terms of the propagating electromagnetic waves, Δk can be seen as the difference in phase velocity between the

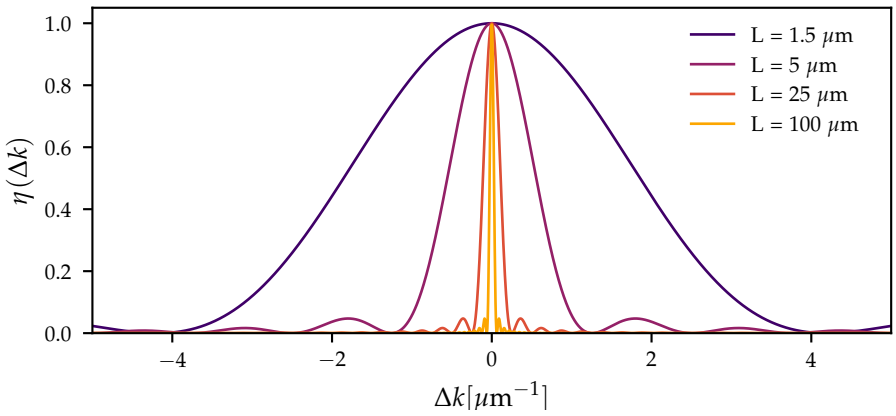

Figure 2.4.: The phase-matching efficiency $\eta = \mathrm{sinc}^2\left(\frac{\Delta kL}{2}\right)$ as a function of Δk for different medium lengths L.

polarisation P_3 that is induced at each plane in z and the already existing field E_3 from previous planes. A non-zero Δk thus implies that these two do not propagate in phase, so that they do not always interfere constructively. Instead of a monotonic build-up of intensity, the intensity I_3 of the DFG signal oscillates between 0 and some fixed value with a period dependent on the magnitude of the wave-vector mismatch.

The second way of understanding the phase-matching condition is through the photon picture of electromagnetic fields. Viewing the DFG process as the "subtraction" of a photon of energy $\hbar\omega_2$ from one energy $\hbar\omega_1$ to yield a photon of energy $\hbar\omega_3$, momentum conservation demands that the total momentum before and after the process is the same. Since the momentum of a photon is $\hbar k$, the phase-matching condition arises from this consideration in the same form.

Non-collinear phase-matching

In practice, it is often difficult and even unwanted to arrange the beams involved in a frequency mixing process in a collinear geometry, so beams are often crossed at an angle in the nonlinear medium. In this case, the propagation of all fields in three dimensions must be considered, and thus vector phase-matching equations must be used. For example, the phase-matching condition for any process mixing two fundamental fields

then reads:

$$\mathbf{k_3} = \pm\mathbf{k_1} \pm \mathbf{k_2}\,, \tag{2.2.31}$$

where the signs are determined by the specific process under consideration. This implies that in addition to the longitudinal phase-matching discussed above, transverse phase-matching is also required. In the transverse plane, there is an additional freedom in the angle of the third beam relative to the two fundamentals. The transverse phase-matching condition therefore determines the angle at which this beam propagates. This is most easily understood in the photon picture, in which the vector momentum of a photon is correspondingly $\hbar\mathbf{k}$ – the total momentum must be conserved, so that the propagation direction of the mixing signal is determined simply the sum (or difference) of the transverse components of the fundamentals. Note that the longitudinal wave vectors are also changed when beams cross at an angle. This changes the longitudinal phase-matching condition significantly if the crossing angle is large.

Birefringent phase-matching

Given that for most optically transparent materials the refractive index in the optical part of the electromagnetic spectrum monotonically increases with frequency [13], it is generally not possible to achieve phase-matching for any desired combination of wavelengths. This is easiest to see in the case of second harmonic generation, where $\omega_2 = 2\omega_1$ and the phase-matching condition (in the collinear case for simplicity) is $k_2 = 2k_1$. This reduces to $n(\omega_2) = n(\omega_1)$, which is impossible to satisfy in a dispersive material except at very particular wavelengths. This problem can be overcome by exploiting birefringence, however.

In a birefringent material, the refractive index varies with the polarisation direction of the electric field due to an anisotropy in the linear susceptibility. The anisotropy can be caused by the structure of the material itself [12], an applied external field [13], or mechanical stress on the bulk material [20]. In general, this is formally expressed by using the susceptibility tensor, but a simpler way is to instead use principal refractive indices.

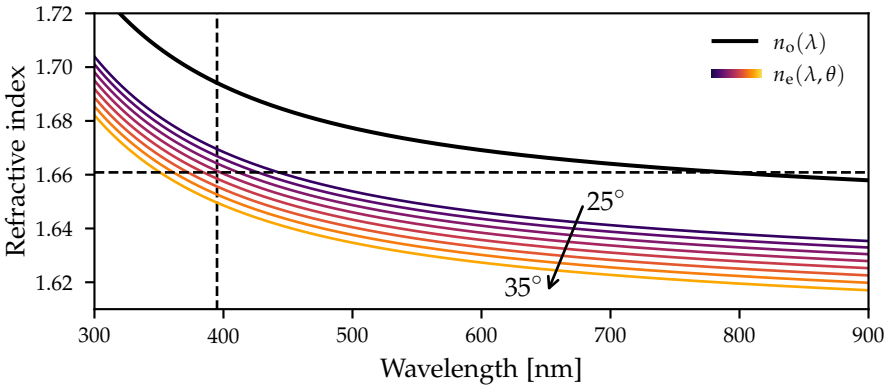

Figure 2.5.: Birefringent type I phase-matching for SHG at 790 nm in a beta barium borate (BBO) crystal. The black solid line shows the ordinary refractive index $\bar{n}_o(\lambda)$ and the coloured lines the extraordinary refractive index $n_e(\lambda,\theta)$ for 8 different angles between 25° and 35°. The dashed horizontal line shows the refractive index at the fundamental wavelength of 790 nm, demonstrating that by choosing the correct crystal angle, it can be matched to the refractive index at the second harmonic wavelength of 395 nm (marked with the vertical dashed line).

Most crystals used for nonlinear optics are *uniaxial*[6], meaning that two elements of the diagonalised susceptibility tensor are equal [12]. Such a crystal has two principal refractive indices, the *ordinary* index $\bar{n}_o(\omega)$ and the *extraordinary* index $\bar{n}_e(\omega)$. A field polarised along the ordinary or extraordinary axis of the crystal experiences the respective refractive index. If the field polarisation lies between the two axes, the refractive index is instead given by [20]

$$n_e(\omega,\theta) = \left[\frac{\cos^2(\theta)}{\bar{n}_o^2(\omega)} + \frac{\sin^2(\theta)}{\bar{n}_e^2(\omega)}\right]^{-\frac{1}{2}}, \tag{2.2.32}$$

where θ is the angle between the propagation direction and the single symmetry axis of the crystal, so that the ordinary refractive index $\bar{n}_o(\omega)$ is recovered at $\theta = 0$. Figure 2.5 shows the effect of this angle-dependent refractive index in a common nonlinear process;

[6]One exception is lithium triborate (LBO), which is biaxial and also commonly used.

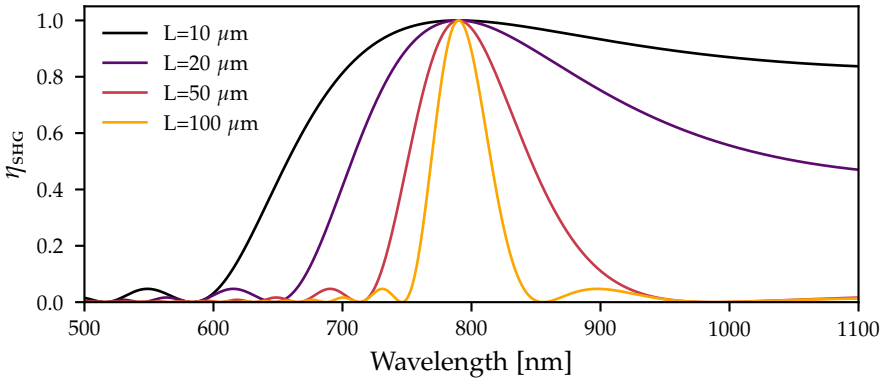

Figure 2.6.: Phase-matching efficiency η_{SHG} according to eq. (2.2.29) for SHG at 790 nm when using a BBO crystal of four different lengths at an angle of $\theta = 29.55°$, which is the type I phase-matching angle.

SHG at the central wavelength of a titanium-doped sapphire (Ti:Sa) laser in a beta barium borate (BBO) crystal. The extra degree of freedom that birefringence offers allows for index-matching between the fundamental and the second harmonic, so the nonlinear process is phase-matched. In practice, the correct phase-matching angle is usually achieved by cutting the crystal such that the phase-matching condition is fulfilled at normal incidence to the medium, however other angles can be obtained by rotating the crystal.

The phase-matching scheme shown here is commonly referred to as *type I* or *o-o-e* phase-matching, since the fields at the longer wavelength are in the *ordinary wave* (polarised along the ordinary axis) [21]. In *type II* or *o-e-e* phase-matching, the polarisations of the two fundamental fields are different; this is usually achieved by simply rotating the crystal around the beam axis [13].

Broadband phase-matching

The phase-matching efficiency curves in fig. 2.6 also make an important point about the mixing of broadband fields. While the coupled wave equation eq. (2.2.25) was given for monochromatic fields, the process of mixing ultrafast pulses can only be accurately

described using the full continuous-frequency nonlinear wave equation. Then, the nonlinear polarisation in the frequency domain is proportional to the convolution (for a $\chi^{(2)}$ nonlinearity) or double convolution (for $\chi^{(3)}$) of the driving field with itself. Therefore each frequency component mixes with all other frequency components of the pulse.

This can be understood qualitatively by imagining a continuum of coupled wave equations, one for each combination of frequencies from the fundamental and mixed fields. The phase-matching condition eq. (2.2.31) then has to hold for each combination. Given that each field will contain many frequencies and therefore wave-vectors, simultaneous phase-matching at all of them is impossible. In this case, the width of the phase-matching efficiency curve gives a measure of the bandwidth of the process. Since this width decreases for a longer medium (in fact, it is inversely proportional to L), there is a trade-off between peak efficiency of the frequency mixing process (proportional to L^2) and the bandwidth that can be converted.

Figure 2.6 shows the phase-matching efficiency for same process shown in fig. 2.5, using the type I phase-matching angle for SHG at 790 nm, which is 29.55°. As the crystal length increases, the bandwidth decreases significantly, demonstrating that very thin ($< 20\,\mu$m) crystals have to be used to efficiently up-convert the entire bandwidth of short pulses. Note that the curve is not symmetric around the central wavelength. This is because BBO (and most other materials) are less dispersive for longer wavelengths, as seen in fig. 2.5 by the lower gradient of the refractive index.

2.2.5. Photoionisation by strong fields

At very high field intensities, it is no longer only the response of bound electrons that affects the propagation of laser pulses. A very strong field can liberate electrons from their binding potential entirely, even when the frequency of the laser pulse is far away from any ionisation resonance. Within the realm of photoionisation by strong fields, two limiting cases can be distinguished: that of multi-photon ionisation, and that of tunnel-ionisation. The difference between the two lies mainly in just how intense the

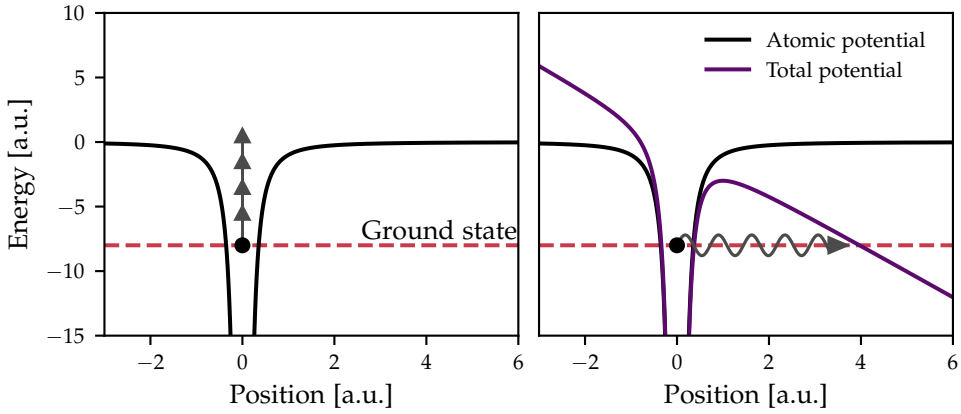

Figure 2.7.: Illustration of the two limiting regimes of non-resonant photoionisation by strong fields. In multi-photon ionisation (left), the process can be considered as the simultaneous absorption of several photons from the laser field, promoting an electron from the ground state (with its energy shown in red) to the continuum. In tunnel-ionisation (right), on the other hand, the field is strong enough to bend the Coulomb potential into a barrier, through which an electron can tunnel to freedom.

laser field is.

Multi-photon ionisation occurs when the field is strong enough that there is a significant probability of an atom absorbing multiple light quanta simultaneously. If the sum of the energies of these photons is equal to or above the ionisation potential of the atom, an electron is freed from the system and can escape (see fig. 2.7). The likelihood of this process decreases dramatically for long-wavelength driving fields, since each individual photon has much less energy. For these cases, tunnel-ionisation dominates.

Tunnel-ionisation requires an even more intense field, such that the field strength is comparable to that of the field between an electron and the nucleus of the atom; a typical value for this field strength is $1\,\mathrm{V\,\mathring{A}^{-1}}$. In this case, the coulombic potential binding the electron is distorted significantly by the addition of a linear component (see

fig. 2.7)[7]. As a consequence, the binding potential energy surface turns into a barrier on one side of the atom, and there is a finite probability of an electron tunnelling through this barrier and leaving the atom.

The importance of photoionisation for ultrashort pulse propagation in particular stems from two facts, namely its influence on the refractive index of the medium and the strong sub-cycle confinement of this change. The presence of free electrons changes the refractive index of the medium in an intensity-dependent manner similar to the Kerr effect, though with the opposite sign. This can be seen by considering the classical equation of motion of a free electron in an electric field, arbitrarily chosen to be polarised in the x direction [20]:

$$m_\mathrm{e}\partial_t^2 x = -\mathrm{e}E(t) \tag{2.2.33}$$

where x is the displacement of the electron and m_e and e are its mass and charge, respectively. Note that since we are considering the effect on femtosecond laser pulses, we can ignore any relaxation or damping, both of which occur on much longer timescales [22]. For a single plane wave, so that $E(t) = A\,\mathrm{e}^{-i\omega t}$, the displacement is given by

$$x(t) = -\frac{\mathrm{e}}{m_\mathrm{e}\omega^2}E(t)\,. \tag{2.2.34}$$

The resulting polarisation is proportional to the displacement x, the charge e, and the electron density ρ_e:

$$P_\mathrm{e}(t) = \mathrm{e}\rho_\mathrm{e}x(t) = -\frac{\rho_\mathrm{e}\mathrm{e}^2}{m_\mathrm{e}\omega^2}E(t) = -\varepsilon_0\frac{\omega_\mathrm{p}^2}{\omega^2}E(t)\,, \tag{2.2.35}$$

where ω_p is the *plasma frequency*:

$$\omega_\mathrm{p} = \sqrt{\frac{e^2\rho_\mathrm{e}}{\varepsilon_0 m_\mathrm{e}}}\,. \tag{2.2.36}$$

[7] The approximation of the laser field as constant across the atom, known as the dipole approximation, is acceptable for wavelengths longer than the scale of the atom; this is the case for photon energies up to the soft X-ray regime.

Since the polarisation is proportional to the field, it contributes to the effective refractive index in a similar way to the Kerr effect. The effective index in this case is given by

$$n_{\text{eff}} = \sqrt{1 + \chi^{(1)} - \frac{\omega_{\text{p}}^2}{\omega^2}}\,. \tag{2.2.37}$$

As long as only a small fraction of the medium is ionised, the plasma frequency is much lower than optical frequencies; for example, an ionisation fraction of 1 %[8] at a gas pressure of 10 atmospheres leads to a plasma frequency of $\omega_{\text{p}} = 0.09\,\text{rad}\,\text{fs}^{-1}$, whereas the frequency of a laser at 790 nm wavelength is $\omega = 2.4\,\text{rad}\,\text{fs}^{-1}$. The contribution of the plasma to the refractive index is therefore small, and it can be approximated as

$$n_{\text{eff}} \approx n - \frac{\omega_{\text{p}}^2}{2\omega^2}\,. \tag{2.2.38}$$

Importantly, since the ionisation rate and thus the electron density ρ_{e} increases with intensity, the refractive index is lower for higher intensities, counteracting the self-focusing caused by the Kerr effect. This is known as *plasma defocusing*, and is an important effect in extreme nonlinear optics [24].

In the time domain, the effect of the plasma cannot be described as simply opposite to that of the third-order nonlinearity. A femtosecond laser pulse is so short that no electrons have time to recombine with their parent atoms. Therefore, the electron density only ever rises over the course of the pulse. Instead of the symmetric red- and blue-shifting of SPM, the pulse is thus purely blue-shifted.

The rate at which a strong field ionises a medium is the subject of decades of research, particularly in the tunnel-ionisation regime [25–28]. In general, the only way to calculate the time-dependent electron density for an arbitrary laser pulse is to solve the time-dependent Schrödinger equation numerically. However, the computational complexity makes this infeasible for all but the simplest systems, and even then it is far too time-consuming to be useful for the modelling of macroscopic phenomena such as ultrashort pulse propagation. Several approximate models for the tunnel-ionisation rate

[8] An ionisation fraction of 1 %, while low, is already substantially higher than what is achieved in most experiments involving gas-based nonlinear optics [23].

have been developed, based on an early quasi-static approach known as the Keldysh model [25]. The Ammosov-Delone-Krainov (ADK) model extends the Keldysh model by considering the structure of hydrogen-like atoms, and it is commonly used because of its relative simplicity [26]. The ADK ionisation rate is given by

$$w(t) = A_{n^*,l^*}\, B_{l,|m|}\, \frac{I_\mathrm{P}}{\hbar} \left(\frac{2\sqrt{m_\mathrm{e}}(2I_\mathrm{P})^{3/2}}{\mathrm{e}\hbar E(t)} \right)^{C} \exp\left[-\frac{2\sqrt{m_\mathrm{e}}(2I_\mathrm{P})^{3/2}}{3\mathrm{e}\hbar E(t)} \right], \tag{2.2.39}$$

where A_{n^*,l^*} and $B_{l,|m|}$ are dimensionless functions of the effective orbital quantum number n^*, the effective angular momentum quantum number l^*, the angular momentum quantum number l and the magnetic quantum number m, $C = 2n^* - |m| - 1$, I_P is the ionisation potential of the medium, and $\hbar$ is the reduced Planck's constant. Further extensions of this model are the Perelomov-Popov-Terent'ev (PPT) model, which takes into account the long-range Coulomb potential of the atom [27], and the Yudin-Ivanov model, which considers sub-cycle non-adiabatic effects [28]. Although these are more accurate approximations, they are also significantly more difficult to implement and computationally expensive.

2.3. High-harmonic generation

The fundamental limit for the generation of sub-femtosecond pulses is that this timescale is below the cycle time of light at the majority of common laser wavelengths; at 800 nm, the most common wavelength used in ultrafast experiments, the oscillation period of the electric field is 2.7 fs. Sub-femtosecond pulses with a spectrum spanning the ultraviolet to infrared spectral regions have been created by extreme nonlinear spectral broadening and subsequent multi-channel compression [16]. However, the most successful technique to generate such short pulses is high-harmonic generation (HHG). Contrary to what the name suggests, this is not simply an extension of the processes of second and third harmonic generation discussed in section 2.2.2 to higher energies. Instead it is more closely related to, and in fact dependent upon, strong-field ionisation.

As described in section 2.2.5, electrons can be liberated from an atom by tunnel ionisation if the field strength is sufficient. This forms the first step in the semi-classical *three-step model* of HHG [29]. The second step was also already alluded to in the discussion of the free-electron polarisation: a free electron exhibits oscillatory motion in the presence of the laser field. This means that after ionisation, the newly freed electron is accelerated away from the ion. When the field reverses polarity, the effect is inverted and the electron instead moves back towards its parent ion. There is then a small probability that the electron recombines with the ion to form a neutral atom once again, emitting a photon in the process. The energy of this photon is larger than that of the driving photons by at least the ionisation potential. Due to the acceleration in the field, however, the electron impacts the atom with significant additional kinetic energy, so a photon at an even higher frequency is created.

Importantly, simultaneous HHG processes in many atoms in an ensemble are phase-locked by their synchronisation to the laser field. In combination with the fact that the process at each atom is coherent, this means that the light source formed by the ensemble of radiating atoms is coherent as well. Harmonic radiation can therefore build up coherently in the same way as described for low-order processes in section 2.2.4, provided that phase-matching between the driving field and the harmonics is achieved

[30]. The synchronisation to the laser field has the additional effect of creating a spatially as well as temporally coherent beam.

The chances of the whole process - ionisation, acceleration, recombination – occurring are very low indeed, leading to very small amounts of energy being converted to high harmonics. One reason for this is that the liberated electron is only properly treated as a quantum particle, and thus as a wavepacket. This wavepacket broadens as the electron propagates away from and back to the atom due to the dispersion inherent to the Schrödinger equation – electrons with more energy travel faster. The transverse spreading of the wavepacket means that the overlap between the returning electron wavepacket and the state it came from is reduced [31]. This, in combination with other factors, leads to typical conversion efficiencies in HHG of below 10^{-6} and as low as 10^{-9} [32].

The kinetic energy of the electron recombining with its parent ion is determined by the intensity of the laser field and how much time it spends being accelerated; this time, in turn, depends on when exactly it is born into the continuum. As a result, the photon energy spectrum created by the HHG process in a single atom is a continuous distribution extending all the way from low-order harmonics to the highest possible energy. This maximum energy is given by

$$E_{\max} = I_{\mathrm{P}} + 3.17\, U_{\mathrm{P}}\,, \tag{2.3.1}$$

where I_{P} is once again the ionisation potential of the atom and U_{P} is the *ponderomotive energy*, which is the average kinetic energy of an electron in a laser field [29]. This is given by

$$U_{\mathrm{P}} = \frac{e^2 A^2}{4 m_{\mathrm{e}} \omega_0^2} \tag{2.3.2}$$

where A is the field amplitude and ω_0 its frequency.

However, if the pulse driving the process is longer than a few cycles of the driving field, so that several cycles are strong enough to cause tunnel ionisation and HHG, a continuous spectrum is not observed. HHG processes occur once for every half-cycle of the field and the resulting bursts of light interfere with each other. Because harmonics

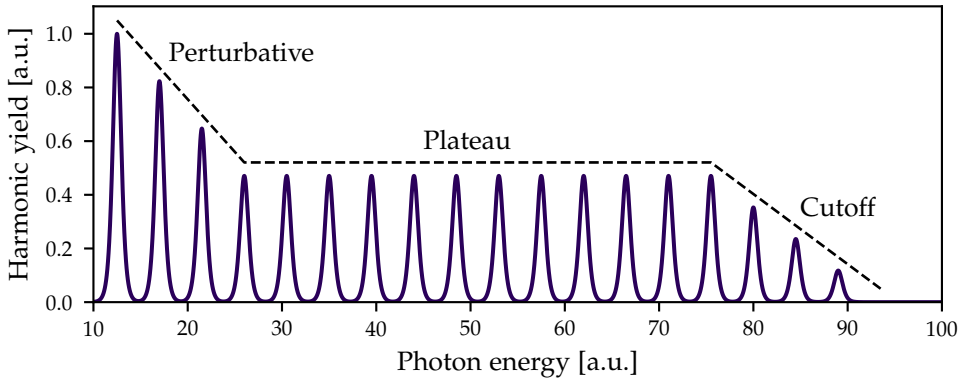

Figure 2.8.: Schematic representation of a typical multi-cycle HHG spectrum. The perturbative regime describes the response of bound electrons and thus low-order harmonic generation. The plateau is an extended region over which the intensity of harmonic emission changes relatively little, before the harmonic yield drops off rapidly near the cutoff.

emitted with the field at one polarity have the opposite phase to those at the other polarity, destructive interference occurs whenever the harmonic frequency is an even multiple of the driving frequency; the harmonic fields at these frequencies accumulate multiples of 2π in phase between half-cycles, so that the additional phase flip from the polarity causes them to cancel. Thus only odd-integer multiples of the driving field frequency appear in a typical harmonic spectrum, a schematic representation of which is shown in fig. 2.8.

In the time domain, the harmonic radiation generated by a multi-cycle pulse is a train of pulses spaced by half the period of the driving field. Because tunnel-ionisation is limited to times when the field strength is near its peak, each pulse is very short, commonly below one femtosecond [33]. A train of such pulses, however, is of limited use as an attosecond probe, since there is no way to distinguish the contributions to the signal from different individual pulses; the time resolution is thus set by the envelope of the driving pulse.

There are several ways to instead generate isolated attosecond pulses using HHG

[34–37]. The method used in our laboratory is known as *amplitude gating* [38]; it consists of simply using a very short pulse to drive HHG. In such a *few-cycle* pulse, the pulse duration is comparable to the cycle time of the field. The field amplitude thus varies rapidly over the course of the field, and only a single half-cycle is strong enough to generate harmonics at the highest energies [38]. While the weaker half-cycles still generate a train of pulses, spectral filtering of the region around the harmonic cutoff will result in a single attosecond pulse.

Soft X-ray attosecond pulses

Attosecond pulses generated with driving lasers operating at 800 nm have been used in a variety of experiments both on atomic and molecular systems [39–41]. However, the maximum cutoff energy achievable in this way is limited, with most 800 nm HHG sources capable of reaching photon energies of around 150 eV but no higher [42]. Although the contribution of the ponderomotive energy in eq. (2.3.1) suggests that the cutoff can be moved to higher energies by increasing the intensity and thus the field amplitude, this cannot be scaled arbitrarily.

As the intensity and thus field amplitude is increased, the total ionisation fraction of the generation medium rises at the same time as the harmonic cutoff energy, because a stronger field leads to more rapid field ionisation. The plasma contribution to the refractive index then leads to an insurmountable phase mismatch between the driving field and the harmonics, precluding efficient build-up of the signal [43]. At the extreme intensities required to push the cutoff much beyond 150 eV, the ground state of the atom is furthermore depleted significantly. As a consequence, a returning electron finds that very little remains of the state it was ionised from; simply put, there is no hole for it to fall back into. Harmonic generation is therefore switched off almost entirely [44].

The need to avoid this regime of poor phase-matching and ground-state depletion is the chief reason that HHG with 800 nm driving pulses cannot reach very high photon energies while generating enough flux for an experiment. For example, at a driving wavelength of 790 nm the intensity required for a harmonic cutoff at 150 eV in a neon

target is ca. $7 \times 10^{14}\,\mathrm{W\,cm^{-2}}$. This intensity is sufficient to cause significant ionisation in a neon atom [43], and thus the cutoff cannot easily be pushed to higher energies. The problem is reduced when HHG is driven with a few-cycle pulse, since cycles preceding the main peak do not contribute much to the overall ionisation, however this does not fundamentally change the limitations.

This energy restriction has consequences for the applicability of conventional high-harmonic sources. Since radiation at these photon energies is very strongly absorbed in almost all materials, only gas-phase samples and the surfaces of condensed-phase materials can be investigated [39]. For the study of biological or bio-chemical processes, this is a significant disadvantage, since the relevant samples are often most easily prepared in solution or in solid form. The region of interest especially for transient absorption spectroscopy of these kinds of samples is the *water window* (see fig. 2.9a). This is the photon energy range between the K-shell absorption edges of carbon, at 284 eV, and that of oxygen, at 540 eV, both of which lie in the soft X-ray (SXR) part of the electromagnetic spectrum. In the water window, water and other carbon-free solvents are nearly transparent, allowing the study of samples in solution as well as thin films of solid materials.

The only remaining parameter in eq. (2.3.1) to increase the harmonic cutoff photon energy into the water window is the driving wavelength. By increasing this (and thus decreasing ω_0), the cutoff can be extended without overly ionising the medium [42, 46, 47]. Since the ponderomotive energy scales with the square of the wavelength, using a driving pulse at a wavelength in the short-wavelength infrared (SWIR) spectral region, for instance 1800 nm, can lead to a maximum photon energy of several hundred electron volts. This can be understood intuitively by considering that the electron can spend more time being accelerated away from and then back towards the parent ion, raising its maximum kinetic energy. The cost of this approach, however, is a very severe drop in conversion efficiency from the already weak pulses obtained at 800 nm.

There are several reasons for the poor scaling of the conversion efficiency. One of the most important ones is the same as the reason for the increased cutoff: the time

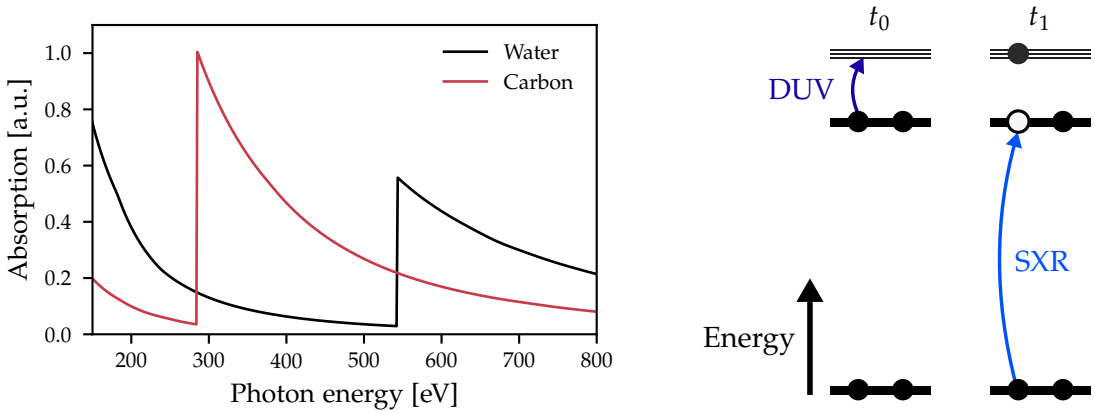

Figure 2.9.: (a) Absorption cross sections (on an arbitrary but common scale) of water and carbon in the SXR spectral region, calculated from the atomic scattering factors [45]. The oxygen K edge is visible in the absorption of water at 540 eV.
(b) Schematic of a possible DUV-SXR pump-probe experiment. At some time t_0, the DUV pulse excites a valence electron, and the SXR pulse probes the state population and coherence at some later time t_1.

spent outside the ion. The longer this is, the more time the electron wavepacket has to disperse. Since the excursion time is proportional to the period of the field and thus the wavelength λ_0, the wavepacket spread in each dimension is proportional to λ_0 as well, reducing the overlap of the returning electron with its ground state by a factor of λ_0^{-3} [31].

When other effects, such as the scaling of the recombination matrix element, are included, the reduction in flux is in fact dramatically worse – as strong as λ_0^{-9} [48]. As a consequence, it is vital that the phase-matching conditions are optimised to achieve maximum overall conversion efficiency, and furthermore that the highest possible pulse energy is used to drive the process. Phase-matching in this regime is achieved by using high-pressure gas targets and very high intensities, and the resulting reshaping of the

laser pulse by plasma effects plays an important role [49].

Time-resolved XANES

The choice of the water window as the photon energy range for the probe pulse has so far only been justified by its applicability to condensed-phase targets. What remains is the question of what information is contained in the absorption spectrum.

Thankfully, the study of X-ray absorption spectra is a well-developed field. The concept most pertinent to the study of few-femtosecond electronic dynamics using soft X-ray attosecond pulses is X-ray absorption near-edge structure (XANES). In this technique, the features in the absorption spectrum around an X-ray absorption edge are investigated as a signature of both the state occupancy and shape of the potential in the atom creating the edge [50]. In a time-resolved experiment, tracking the change in these absorption features after photoexcitation can yield extensive information about, for instance, the formation and disappearance of electron-hole pairs [51,52].

Since X-ray absorption edges such as the carbon K edge arise from the promotion of a deeply bound core electron to a high-lying state or the continuum, a time-resolved XANES experiment offers spatial as well as temporal resolution; core electrons are strongly localised to their parent atom and their corresponding X-ray absorption features therefore encode the dynamics local to that atom. In a molecule with, for instance, a single carbon atom, the carbon K edge XANES signal therefore probes a very spatially localised part of the molecule.

Figure 2.9b shows a simplified schematic of a time-resolved XANES experiment. A valence electron in a molecule is promoted to a superposition of higher-lying states by a pump pulse at time t_0. The hole it leaves behind creates a new feature in the X-ray absorption spectrum, but only for as long as the hole survives. At some time t_1 later, a soft X-ray pulse probes the state by promoting a core electron to the hole if possible. Taking the XANES spectra shown in fig. 2.10 as an example, a valence excitation may create an additional absorption feature below the sulphur L edge, so around 160 eV, or the carbon K edge, so around 280 eV, depending on how the hole is

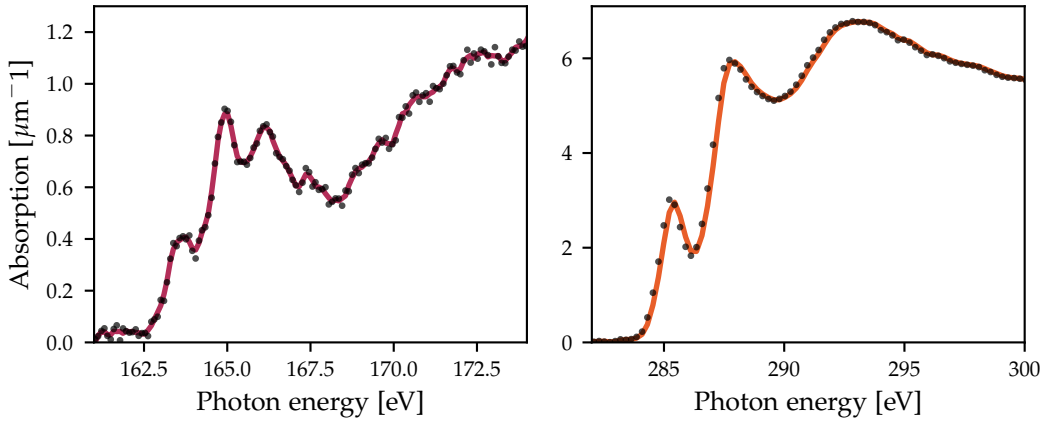

Figure 2.10.: Near-edge X-ray absorption spectra at the sulphur $L_{2,3}$ edge (left) and carbon K edge (right) in a polythiophene film acquired with the high-harmonic supercontinuum source in our laboratory, which was developed by Allan Johnson [53]. The dots show the individual data points and the solid lines show the data after nearest-neighbour smoothing. Different peaks in the absorption spectrum near the edge correspond to the transition of a core electron to different states, and the slower features above the edge contain information about the environment of the atom.

localised. By measuring the strength of the absorption feature as a function of delay between the pump and probe pulses, the dynamics of the excited electron and the hole can be tracked. Electron correlations can furthermore shift other absorption features, yielding another signal channel, however such shifts are more difficult to detect.

Light source requirements

Time-resolved XANES is well-placed to be a very powerful technique in ultrafast measurement, however it places great demands on the light sources for the experiment. Since attosecond pulse generation via HHG is extremely inefficient, especially when using the long driving wavelengths required to reach the X-ray regime, achieving the maximum driving pulse energy possible is absolutely critical.

The system used to generate SXR attosecond pulses in the Laser Consortium at Imperial College is based on a heavily modified commercial laser system (Red Dragon, KM Labs). The laser system generates pulses at 790 nm central wavelength with a FWHM duration of 30 fs and up to 8.3 mJ of pulse energy at a repetition rate of 1 kHz. To create the long-wavelength driving pulse required for SXR generation, around 8 mJ of this laser output is used as the input for a commercial tunable optical parametric amplifier (OPA) (HE-TOPAS, Light Conversion). In the OPA, white-light generation by very strong SPM creates a broadband seed pulse which can be spectrally selected and amplified, resulting in a tunable pulse. The amplification creates two pulses further in the infrared, with the pulse at the shorter wavelength conventionally called the signal, and the one at the longer wavelength the idler.

The SXR generation apparatus uses the idler, which is generated with a central wavelength of around 1750 nm. The pulse created in the OPA is somewhat longer than the laser pulses, around 35 fs to 40 fs. It is passed through a pulse compressor consisting of a gas-filled hollow capillary to expand the spectrum via SPM and a pair of glass wedges to induce the necessary anomalous GDD[9], after which its duration is around 10 fs. At this wavelength, this corresponds to around 1.6 cycles, short enough for the generation of isolated attosecond pulses.

Owing to its operating principle, the conversion efficiency to the idler in the OPA is not very high. Commonly, the idler pulse contains only about 1.4 mJ of energy, less than 20 % of the input. This is further reduced by about 50 % in the pulse compressor. As a result, the pulse driving SXR generation usually has an energy of 600 µJ to 700 µJ. While this is sufficient to generate harmonics across the water window in a suitably optimised target [49, 53], achieving this energy is crucial; a reduction of even 5 % in laser output from its nominal performance precludes operation of the SXR source. It is this inefficiency and the resulting restriction that make an efficient source of ultrashort DUV pulses so critical for a successful experiment involving SXR attosecond pulses.

One important question to be addressed when designing an experiment is the

[9] Recall that most optical glass is anomalously dispersive at this wavelength.

necessary energy for both pulses involved. In the type of experiment discussed so far, the lower limit on the SXR flux is set by the need to detect it after transmission through a sample. Our SXR source has already been used to perform static XANES measurements at the sulphur L and carbon K edges in a target consisting of a thin film of photoactive polymer (see fig. 2.10) [53] and of the oxygen K edge in a Mylar film [49], demonstrating that is bright enough in principle for time-resolved measurements.

The energy requirements on the DUV pulse are harder to estimate, since they are extremely variable depending on the target. The absorption cross-section of valence-electron resonances in small molecules the deep and vacuum UV can range from $10^{-19}\,\mathrm{cm}^2$ to $10^{-16}\,\mathrm{cm}^2$ per molecule [54,55]. Even for otherwise identical experimental conditions, this results in a change of over three orders of magnitude in the required pulse energy. The situation is further complicated by the fact that only those molecules in the focus of the SXR pulse contribute to the transient absorption signal. Given the very short wavelength, this focus is very likely to be smaller than that of the DUV pulse, increasing the total amount of DUV energy required.

For a rough estimate, expertise from the femtochemistry community is very useful. Typical values for the pulse energy in a transient absorption experiment at visible and infrared wavelengths lie in the range of 10 nJ to 1 µJ, which is sufficient to excite around 10 % of molecules in most experiments [6]. Given that valence-electronic resonances are often stronger (i.e., cause more absorption), this can serve as a conservative estimate. It should be noted, however, that these are values for the energy in the target, rather than at generation. Should significant losses be incurred between the two, more energy has to be generated.

2.4. Frequency-resolved optical gating

The vast majority of time-resolved measurements in science are performed using some kind of gate, that is, an event with a timescale significantly faster than that of the process under study. The most common example of this is a flash as used in photography. In electronics, this is commonly the sampling clock in a digital oscilloscope, which causes voltage readings to be carried out at a rapid pace. The limit on the time resolution in this case is set by the response time of electronic circuitry as well as available photodetectors, and lies in the range of tens of picoseconds. To measure ultrashort laser pulses, different methods are therefore required. For laser pulses as short as several hundred femtoseconds, a *streak camera* can be used to measure the intensity envelope directly [56]. In this case, the gate is provided by the fast movement of a mirror or electro-optic equivalent, which causes different time-slices of the pulse to be directed to spatially separated parts of a detector.

When faced with the task of measuring short femtosecond or even attosecond pulses, such a strategy cannot be used, because a sufficiently fast gate does not exist: ultrashort laser pulses are the shortest controlled events ever created in a laboratory [57]. No faster event is therefore available to sample the waveform of an ultrashort pulse, and other methods are required. All available techniques used to measure ultrashort laser pulses make use of two or more[10] such pulses and some nonlinear interaction between them to gain information about the temporal shape of a pulse.

The use of a nonlinearity is not merely a convenient way to implement ultrashort pulse measurement, it is a fundamental requirement. The effect of linear optical interaction can be expressed as a product of terms in the frequency domain, as shown in section 2.1.3 for the example of propagation through material. Since the order of linear operations is irrelevant, the linear response of any optical system can be encapsulated in its *response function* $H(\omega)$, the formal properties of which have been investigated in detail in the field of signal processing theory [59]. The electric field of

[10] The exception to this is the dispersion scan method, which only requires the test pulse. However, it can only be used when the GDD of the pulse can be scanned through 0. [58]

the laser pulse at the end of the system is then obtained by simply multiplying the input by the response function. Given the lack of sufficiently fast photodetectors, any linear optical measurement can at most produce the time-integrated (or equivalently, frequency-integrated) intensity of the field, so any such measurement can be described as

$$S = \int_{-\infty}^{\infty} |H(\omega)E(\omega)|^2 \mathrm{d}\omega \,, \tag{2.4.1}$$

where S is the signal acquired in the measurement. This kind of measurement is therefore insensitive to the phase of $E(\omega)$, yielding, for instance, the power spectrum $|A(\omega)|^2$ (if $H(\omega) = \delta(\omega - \omega_0)$ with ω_0 variable) or the total energy (if $H(\omega) = 1$). The spectral phase can only be measured if the response is nonlinear.

This section introduces the most common ultrashort pulse characterisation scheme, frequency-resolved optical gating (FROG), two variants of which were implemented in the course of the work for this thesis. Detailed descriptions of these can be found in section 4.3.3 and chapter 5.

2.4.1. General description

In its simplest form, FROG is best explained as a frequency-resolved measurement of the autocorrelation of the unknown pulse (the *test pulse*). The autocorrelation, in turn, is arguably the simplest measurement of the profile of an ultrashort pulse.

In the context of ultrashort pulses, the term autocorrelation most often refers to the second-order intensity autocorrelation $S(\tau)$, which is essentially the convolution of the pulse intensity with itself [60]:

$$S(\tau) = \int_{-\infty}^{\infty} |E_\mathrm{t}(t)E_\mathrm{t}(t-\tau)|^2 \, \mathrm{d}t = \int_{-\infty}^{\infty} I_\mathrm{t}(t)I_\mathrm{t}(t-\tau) \, \mathrm{d}t \,, \tag{2.4.2}$$

where $E_\mathrm{t}(t)$ is the electric field of the test pulse and $I_\mathrm{t}(t)$ its intensity (neglecting constants). $S(\tau)$ is particularly easy to obtain using nonlinear optics by measuring the total power in the sum-frequency signal of a pulse with a delayed replica of itself as a function of that delay. Writing the driving field in the nonlinear medium for such an

arrangement as a function of both time t and delay τ,

$$E(t,\tau) = E_t(t) + E_t(t-\tau)\,, \tag{2.4.3}$$

the second-order nonlinear polarisation (omitting the optical rectification term) is then given by

$$P_2(t,\tau) = \frac{1}{4}\varepsilon_0\chi^{(2)}\left[E_t(t)^2 + E_t(t-\tau)^2 + 2E_t(t)E_t(t-\tau)\right]. \tag{2.4.4}$$

As shown in section 2.2.3, the first two terms in the nonlinear polarisation represent the second harmonic of the two delayed replicas. The third term represents the sum-frequency generation process between them, and its intensity, proportional to $|E_t(t)E_t(t-\tau)|^2$, has the same form as the integrand in eq. (2.4.2). It is therefore sufficient to measure total power in the sum-frequency signal between two replicas of the test pulse as a function of the delay between them to obtain the intensity autocorrelation, as long as the phase-matching conditions are appropriate (see section 2.4.2).

The autocorrelation contains information about the pulse. If it is known that the pulse has a particular shape, then the autocorrelation yields its duration. Furthermore, the root mean square (RMS) widths of $S(\tau)$ and $I_t(t)$ are related unambiguously, since the definition of the RMS width is independent of pulse shape. However, it is not possible to fully reconstruct the field as a function of time including the phase [60]. There are an infinite number of pulse profiles $E_t(t)$ which will result in the same autocorrelation trace $S(\tau)$, because the problem of gaining phase information from the intensity-only autocorrelation (known as the *one-dimensional phase retrieval problem*) is under-constrained. Thankfully, there is a simple way of adding enough information to the measurement for the phase retrieval to succeed, and that is to measure not the time-integrated energy of the autocorrelation, but instead its spectral energy density. This is the concept of FROG [61].

To spectrally resolve the sum-frequency part of the second-order polarisation eq. (2.4.4) is to Fourier transform it before measuring the intensity, and yields a signal which is dependent on both delay and frequency:

$$S(\omega,\tau) = \left|\int_{-\infty}^{\infty} E_t(t)E_t(t-\tau)e^{i\omega t}dt\right|^2. \tag{2.4.5}$$

The fundamental difference between this and the autocorrelation is that by adding the second dimension to the data, the problem of finding the underlying field $E_t(t)$ turns into the *two-dimensional* phase retrieval problem, which can be solved as long as there is some additional constrain on the solution [62]. In the case of FROG, the additional constraint is given by knowledge of the process that creates the signal $S(\omega, \tau)$. Here, it is the form of the second-order polarisation eq. (2.4.4). It should be noted, however, that sum-frequency generation is far from the only process capable of producing FROG traces. A frequency-resolved third-order autocorrelation, for example, is also a FROG measurement. In fact, the first demonstrated FROG measurement was carried out using polarisation gating, a consequence of the Kerr effect when two pulses with different polarisations interact [61]. Any instantaneous nonlinearity can be used, and it is only the particular form of the integrand in eq. (2.4.5) which changes [63]. In order for the measurement to be meaningful, however, it is important for several conditions to be met, and furthermore a practical solution to the phase retrieval problem posed by the FROG trace has to be found.

2.4.2. Practical limitations and conditions

The simplest yet most important condition on any attempted FROG measurement is that the nonlinear interactions need to be phase-matched [64]. As detailed in section 2.2.4, it is impossible to achieve perfect phase-matching across a large bandwidth, with the conversion efficiency dropping as the wave vector mismatch Δk increases. If the phase-matching bandwidth is insufficient, the phase-matching efficiency reshapes the FROG trace. It is possible to treat the effect of phase-matching as a simple spectral filter and correct for it [65]. This is an approximation, though one that is very often used, since including the full two-dimensional phase-matching function can substantially slow down the phase retrieval. Another way of circumventing this problem is to use a sufficiently thin nonlinear medium, so that phase-matching effects are negligible. For very short pulses this becomes challenging since the total signal generated scales with the square of the medium length, however it is often the only available option. A thick

medium may not yield any detectable signal in parts of the pulse spectrum due to the narrow phase-matching window, so that a simple correction factor cannot be used.

Spectral reshaping of the FROG signal occurs because of other reasons, too. The overall ω^2-dependence of the harmonic signal as seen in eq. (2.2.25) similarly becomes more important as the test pulse bandwidth increases. The frequency-dependent transmission or reflection of the optics in the apparatus also plays a role, as does the spectral sensitivity of the spectrometer. These are much simpler to account for, however, as they depend purely on the wavelength of the signal, and not that of the interacting pulses.

In most beam geometries used for FROG, the two (or three) copies of the test pulse cross at an angle in the nonlinear medium [60]. In addition to being easier to construct, this makes it possible to spatially separate the signal beam created in the medium from the incident beams, since transverse phase-matching will cause it to travel in a different direction. This removes a potential source of background from the signal. In the case of SHG FROG, the second harmonic of the two pulse replicas has the same central frequency as the signal and would thus overlap with it in a spectrometer; a similar problem is present for variants of FROG which do not change the central frequency. The non-collinear geometry creates its own problems, however. Since the pulses impinge on the nonlinear medium at an angle, one edge of each beam will reach the medium before the other, and which side hits first or last is reversed between the two pulses. The delay between the pulses is therefore not constant across the face of the medium. The amount of time by which the extreme edges of the beam are delayed relative to each other (at nominal zero delay, i.e., when the centres of the beams overlap in time) is easily found from the geometry of the situation, and given by

$$\Delta\tau = \frac{2d}{c}\tan\left(\frac{\theta}{2}\right), \tag{2.4.6}$$

where d is the beam diameter in the medium, θ is the crossing angle between the beams, and c is the speed of light in vacuum [60]. For a beam focused to a spot of 50 µm diameter, a crossing angle of 1° is sufficient to cause a spread of around 2 fs in delay across the beam profile. While this can be significant compared to the pulse duration,

the impact on the measured pulse is not as dramatic: the blurring and the real pulse duration add in quadrature, so that 2 fs of blurring is very small for all but the shortest pulses. This can be shown analytically for a pulse that has a Gaussian shape in both time and space [64].

If required, there are two ways of alleviating this problem. The first is to simply choose as small a crossing angle as possible, though this can only reduce the amount of delay blurring to around one optical cycle of the field [60]. The second is to re-image the interaction region in the nonlinear medium onto the entrance slit of the spectrometer with a magnification such that only the centre of the signal beam enters the detector. While this reduces the amount of signal collected, it very easily removes temporal blurring.

Instead of reducing the problem of temporal blurring, another option is to use it. Since the smearing leads to a mapping from space to delay in the medium, imaging this plane onto an imaging spectrograph can yield a FROG trace in a single shot. This is a commonly used pulse characterisation technique, however it is significantly more complicated to set up in practice than a multi-shot, scanning FROG device [60].

2.4.3. FROG trace reconstruction

The final task in any FROG measurement is to solve the phase retrieval problem, a process often called *reconstruction*. This is impossible to do in a single step, given that eq. (2.4.5) cannot be solved analytically to yield the test pulse field $E_t(t)$. Instead, an iterative algorithm has to be employed which makes use of the two constraints of the measurement; the measured trace on the one hand, and the underlying mathematical form of the interaction on the other [60]. Several different algorithms exist, which mostly differ in how they apply the mathematical-form constraint at every iteration, whereas the data constraint is used in the same way. All FROG algorithms start with an initial guess for the test pulse field, update it using the constraints, and then generate a FROG trace from it. These steps are repeated until the difference between this reconstructed and the measured trace converges to a minimum value. This difference

is commonly measured using the RMS error between the two traces, also known as the *FROG error*:

$$\Delta = \sqrt{\frac{1}{N}\sum_{i,j}\left[\left(S_{\mathrm{rec}}(\omega_i,\tau_j) - \mu\, S_{\mathrm{meas}}(\omega_i,\tau_j)\right]^2}\,, \tag{2.4.7}$$

where N is the total number of points in the FROG trace, ω_i are the frequency samples, τ_j the delay samples, and $S_{\mathrm{meas}}(\omega_i,\tau_j)$ and $S_{\mathrm{rec}}(\omega_i,\tau_j)$ are the measured and reconstructed FROG traces, respectively [60]. The factor μ is chosen to minimise the error and is given by [58]

$$\mu = \frac{\sum_{i,j} S_{\mathrm{meas}}(\omega_i,\tau_j) S_{\mathrm{rec}}(\omega_i,\tau_j)}{\sum_{i,j} S^2_{\mathrm{meas}}(\omega_i,\tau_j)}\,. \tag{2.4.8}$$

Applying the data constraint is quite simple. From the current guess for $E_{\mathrm{t}}(t)$, the complex-valued FROG trace *amplitude* $\sigma(\omega,\tau)$ is generated, which is the FROG trace before taking the absolute value:

$$\sigma(\omega,\tau) = \int_{-\infty}^{\infty} E_{\mathrm{t}}(t) E_{\mathrm{t}}(t-\tau) \mathrm{e}^{i\omega t}\mathrm{d}t\,. \tag{2.4.9}$$

The magnitude of this quantity is then replaced by the square root of the measured trace $S(\omega,\tau)$:

$$\sigma'(\omega,\tau) = \frac{\sigma(\omega,\tau)}{|\sigma(\omega,\tau)|}\sqrt{S(\omega,\tau)}\,. \tag{2.4.10}$$

In the simplest FROG reconstruction algorithm, the form constraint is then applied by Fourier transforming $\sigma'(\omega,\tau)$ along the frequency axis to obtain $\sigma'(t,\tau)$ and integrating this over the delay τ; since the test pulse $E_{\mathrm{t}}(t)$ can be factored out of this integral, the result will be proportional to it. This then forms the updated guess, and the next iteration starts. Unfortunately, this basic algorithm often fails in the presence of noise, and is generally slow to converge.

A more reliable method is given by the method of generalised projections [66]. The main difference to the basic algorithm is that instead of generating the updated guess for $E_{\mathrm{t}}(t)$ by simply integrating, a generalised projection onto the set of all possible $\sigma'(\omega,\tau)$ that satisfy the form constraint is carried out. This yields an updated guess which fulfils the form constraint exactly. It can also be shown that the process of replacing the amplitude as shown in eq. (2.4.10) is a generalised projection onto the set of all waveforms which satisfy the data constraint [67].

While the method of generalised projections is much more robust to noise and works in a greater variety of circumstances, it is relatively slow, since it involves the computation of gradients in as many dimensions as there are frequency samples in the FROG trace [66]. A faster implementation of a similar idea is the principal components generalised projection algorithm (PCGPA) [68]. By noting that a FROG trace can be obtained by the outer product of two vectors representing the two pulses interacting in the medium, the phase retrieval problem becomes tractable using singular value decomposition (SVD) techniques. Further simplification by replacing the accurate, but computationally expensive full SVD with the power method results in an algorithm that is fast enough to allow for real-time measurements of ultrashort pulses [69,70].

Due to its relatively simple operating principle and a well-developed range of variations, FROG has become the most commonly used technique for the measurement of ultrashort laser pulses. Two different FROG set-ups were used in the work for this thesis; one to measure the pulse driving the frequency conversion to the DUV, and one to measure its output pulses.

3. Ultrashort Pulses in Waveguides

The multitude of nonlinear optical effects introduced in section 2.2 have found applications in many areas of science and technology. As shown in section 2.2.4, the refractive index profile of the nonlinear material can be as important as the order and strength of the nonlinearity in achieving efficient frequency conversion. In order to fully realise the potential of nonlinear optics, it is therefore important to be able to use the whole spectrum of materials as the medium, from dilute gases to liquids and solids. Using a gas as the nonlinear medium additionally offers the ability to adjust the strength of the nonlinearity by changing the gas pressure.

One of the most successful strategies in accessing this range of possible materials is the use of guided light instead of free-space propagation. Optical waveguides can be made out of solid materials, or they can have hollow cores which can then be filled with different liquids or gases to serve as the nonlinear medium. Fibres furthermore allow very long interaction lengths in the range of metres, enabling the use of materials with relatively weak higher-order susceptibilities while still achieving strong nonlinear effects. This is in contrast to the most common geometry in nonlinear optics experiments of focusing a laser beam into a nonlinear medium, where the interaction length is often shorter than the Rayleigh length of the focus, commonly in the range of tens to hundreds of micrometres.

I will review the mechanism by which both solid and hollow optical waveguides confine light, with particular focus on the two types of hollow-core waveguide used in the work for this thesis. I will also develop the theoretical background of nonlinear pulse propagation in these types of waveguides, including both the full (single-mode) propagation equation and an approximate version which is useful for qualitative

understanding and analytical insight. This enables the discussion of the theory of soliton self-compression, which is the basic mechanism underlying the DUV generation scheme implemented in the work for this thesis.

3.1. Propagation in Fibres and Hollow Capillaries

The first and by far the most common type of optical waveguide, conventional optical fibre consists of a solid core with circular cross-section surrounded by a solid cladding with a different refractive index. Its main application is in telecommunications, however it is ubiquitous in both scientific and commercial contexts.

The guiding mechanism in this so-called step-index fibre is appropriately known as *index guiding* and is easily understood from a ray-optical point of view in terms of total internal reflection. Total internal reflection occurs whenever a ray impinges on an interface between materials at an angle such that its refracted angle according to Snell's law becomes larger than 90°. This means that the ray cannot propagate in the other material and it is reflected instead [71]. One can imagine the guiding of light inside a step-index fibre as the repeated total internal reflection of rays on the interface between the core and the cladding. Since total internal reflection is only possible when a ray approaches the interface from inside the material with larger refractive index, the refractive index of the cladding must be smaller than that of the core.

To facilitate the comparison with other guiding mechanisms, it is also useful to analyse a step-index fibre using wave optics. Here, the first important fact is that the largest value that any component k_i of the wave vector $\mathbf{k}$ can take is $k_i = |\mathbf{k}| = n\omega/c$. The second is that at an interface in the transverse plane of the waveguide, the longitudinal component of wave vector k_z cannot change due to the symmetry of the situation – there is no refractive index change in the longitudinal direction [72]. Any plane wave component of the field for which $k_z > n_{\text{clad}}\omega/c$ can therefore not propagate in the cladding and is confined to the core. As in the ray-optical picture this necessitates $n_{\text{clad}} < n_{\text{core}}$ for any light to be confined.

Hollow capillaries

The requirement that the cladding index be smaller than the core index rules out total internal reflection as a guiding mechanism for any structure with a hollow core, since

the refractive index of glass and other solid materials is much larger than that of a gas, even at high pressures. However, it is still possible to guide light in a hollow capillary, that is, a waveguide consisting of a circular hollow core surrounded by a dielectric such as glass [73]. Again it is useful to consider this situation from different perspectives.

To understand confinement in a hollow core in ray optics, the critical piece is a mechanism to replace total internal reflection. This is provided by glancing reflection, i.e. the fact that at very shallow angles of incidence, a large proportion of light is reflected at the interface between any two media with a large difference in refractive index, irrespective of the sign of the difference. While this provides a route towards light confinement, it also points to one of the drawbacks of capillary waveguides – since only most, rather than all, of the light is reflected at the interface, transmission through capillaries is very lossy compared to step-index fibres.

A wave-optical treatment of a capillary waveguide begins very similarly to the general derivation outlined in section 2.1, which is to say with Maxwell's equations. Solving Maxwell's equations taking into account a capillary's refractive index profile yields the modes of the waveguide. In the case of free space propagation discussed in section 2.1.1, the modes were simply plane waves. The geometry of a capillary makes cylindrical coordinates (radial coordinate r, the azimuthal angle θ and the longitudinal direction z) the natural choice. The natural basis functions in this coordinate system are Bessel functions, and in free space any combination of these is a valid solution to the wave equation.

A mode of the capillary is a particular linear combination of basis functions which satisfies the boundary conditions arising from the refractive index profile and results in a single propagation constant, which is to say that the spatial profile of a mode does not change as it propagates. The (linear) evolution of the field in a specific waveguide mode, commonly labelled by an integer m, along the z direction can be expressed as

$$E_m(r,\theta,z,\omega) = a_m(\omega)\, e_m(r,\theta)\, \mathrm{e}^{i\beta_m(\omega)z}, \tag{3.1.1}$$

where $a_m(\omega)$ is the (complex) spectral amplitude at frequency ω, $e_m(r,\theta)$ is the transverse mode profile in mode m and $\beta_m(\omega)$ is the propagation constant for that mode.

The frequency dependence of β_m is known as *modal dispersion*. The total field in the waveguide is then a sum over modes:

$$E(r,\theta,z,\omega) = \sum_m E_m(r,\theta,z,\omega). \tag{3.1.2}$$

The form of the modes is given by the waveguide geometry, i.e., its refractive index profile, which enters the mathematical treatment as the boundary conditions on the magnetic and electric fields.

In a capillary, many different modes exist which differ in their transverse intensity profile as well as their polarisation direction, in addition to their propagation constant. However, a large proportion of possible modes are excluded in real-world applications by the laser beam that is coupled into the waveguide. Which modes are excited is determined by decomposing the field at the entrance face of the capillary into the modes of the waveguide. For the idealised case of a perfectly linearly polarised Gaussian laser beam, the field at the entrance face of the waveguide is cylindrically symmetric and linearly polarised. Therefore it only couples to modes which share those characteristics.

The class of modes that satisfy this condition is that of the *hybrid modes*, often labelled HE_n [73]. They are known as hybrid modes since the longitudinal field E_z is non-zero, in contrast to the *transverse electric* or *transverse magnetic* modes – there are no modes in which both the electric and magnetic field are transverse as they are in free space. The transverse scalar field of the hybrid modes (dropping the polarisation vector) takes a particularly simple form:

$$e_m(r,\theta) = J_0\left(k_{\perp,m} r\right), \tag{3.1.3}$$

where $J_0(x)$ is the 0^{th} order Bessel function of the first kind, $0 \leq r \leq a$ with a the waveguide core radius, and $k_{\perp,m}$ is a constant which, in analogy to the case of plane waves, is often called the *transverse* wave vector[1].

[1] This can be understood by noting that a higher value of $k_{\perp,m}$ leads to a higher number of subsidiary peaks in the field distribution within the waveguide core, which corresponds to a higher spatial frequency. This is analogous to a higher value of k for plane waves, which directly represents the (angular) spatial frequency.

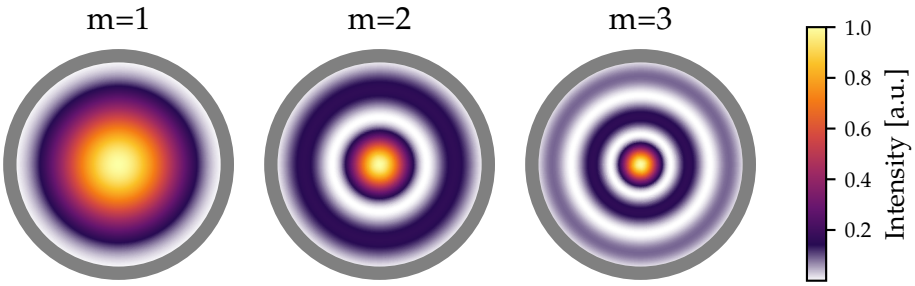

Figure 3.1.: The intensity profile $|e(r,\theta)|^2$ of the first three hybrid modes in a hollow capillary on a false colour scale. The grey circles signify the boundary of the core. The increasing number of subsidiary peaks is clearly visible, which corresponds to a higher spatial frequency and larger magnitude of $k_\perp$.

It is important to note that the modes as shown in eq. (3.1.3) are derived using several approximations, which are often collectively referred to as the *Marcatili model* or the *capillary model* [73]. One of the approximations is that the field vanishes at the edge of the waveguide core, where $r = a$. This restricts the possible values of $k_{\perp,m}$:

$$\begin{aligned} e_m(r=a,\theta) &= J_0\left(k_{\perp,m}a\right) = 0 \\ \implies k_{\perp,m} &= \frac{u_m}{a}, \end{aligned} \tag{3.1.4}$$

where u_m is the m^{th} root of $J_0(x)$. The modes of a capillary are thus *truncated* Bessel functions. Figure 3.1 shows the transverse intensity profile of the first three hybrid modes in a capillary, corresponding to the first three roots of $J_0(x)$. With this simple expression for the transverse wave vector, we can derive the propagation constant β for each mode of the capillary: in the core of the waveguide, the magnitude of the wave vector, $k(\omega)$, is simply $\omega n(\omega)/c$, where $n(\omega)$ is the frequency-dependent refractive index in the core. Expanding k^2 in terms of the transverse wave vector $k_\perp$ and the

propagation constant $\beta(\omega) = k_z(\omega)$ yields the propagation constant:

$$\begin{aligned} k^2 = \frac{\omega^2}{c^2} n(\omega)^2 &= k_\perp^2 + \beta(\omega)^2 \\ \implies \beta_m(\omega) &= \frac{\omega}{c}\sqrt{n(\omega)^2 - \frac{c^2 u_m^2}{a^2\omega^2}}\,. \end{aligned} \tag{3.1.5}$$

In an evacuated capillary, $n(\omega) = 1$, whereas with a gas fill it is simply the refractive index of the gas. This means that it is not only the nonlinearity that is tunable via the gas pressure and species, but also the dispersion of the waveguide.

Importantly, the GVD of an evacuated capillary is always anomalous. This can be seen by expanding the square root in eq. (3.1.5) to first order:

$$\beta_m(\omega) \approx \frac{\omega}{c}\left(1 - \frac{c^2 u_m^2}{2a^2\omega^2}\right), \tag{3.1.6}$$

so that the GVD is given by

$$\beta_{2,m} = \frac{\partial^2 \beta_m}{\partial \omega^2} \approx -\frac{c u_m^2}{a^2 \omega^3}\,. \tag{3.1.7}$$

For large core diameters, this effect is quite weak, and the normal dispersion arising from any gas fill can easily be larger than the anomalous contribution from the waveguide. If the core is only a few tens of µm in diameter, however, the waveguide dispersion can be very strong and the total dispersion remain anomalous even with high-pressure gas fills.

A more thorough derivation of the capillary modes and $\beta(\omega)$ also yields the imaginary part of the propagation constant, i.e. the loss, which is given by[2]

$$\alpha_m(\omega) = \frac{c^2 u_m^2}{a^3\omega^2}\frac{\nu^2 + 1}{\sqrt{\nu^2 - 1}}, \tag{3.1.8}$$

where ν is the ratio of the refractive index of the cladding to that of the core [73]. Importantly, the loss scales as λ^2/a^3, so it increases very rapidly with decreasing core diameter, but smaller cores can be used for shorter wavelengths. If significant losses

[2]Note that in this expression, α is a factor of 2 larger than given in the original work by Marcatili and Schmeltzer [73]. This is because in that work, α is defined to be the attenuation coefficient for the *field* as $E(z) \propto \exp\{-\alpha z\}$.

are to be avoided, this limits the maximum intensity inside the core for a given pulse energy and duration.

The loss is much higher for larger modal indices m, so that higher-order modes are attenuated much more rapidly than the lowest one. The guiding mechanism of glancing reflection explains why this is the case: as detailed above, the transverse wave vector $k_\perp$ is larger for higher-order modes. In the ray picture, this means that the light approaches the core-cladding interface at a steeper angle, which leads to larger losses through refraction. The effect of this mode-dependent loss is that a capillary acts as a *spatial filter*. If the spatial profile of the beam is distorted at the input of the capillary, this will translate into the presence of higher-order modes. This is because the coupling efficiency between the incoming beam and each mode of the fibre is determined by an overlap integral [74]:

$$\eta_m = \frac{\left[\int_0^\infty \int_0^{2\pi} E(r,\theta) e_m^*(r,\theta)\, r\mathrm{d}r\mathrm{d}\theta\right]^2}{\int_0^\infty \int_0^{2\pi} |e_m(r,\theta)|^2\, r\mathrm{d}r\mathrm{d}\theta \int_0^\infty \int_0^{2\pi} |E(r,\theta)|^2\, r\mathrm{d}r\mathrm{d}\theta}\,, \tag{3.1.9}$$

where $e_m(r,\theta)$ is the field profile of the m^{th} mode and $E(r,\theta)$ is the field profile of the incoming beam at the entrance face of the fibre. Perfect coupling into the fundamental mode is therefore only possible if the beam has exactly the correct shape. Since higher-order modes are attenuated, however, the spatial profile of the output will consist mostly of the fundamental mode, effectively removing the distortions from the beam. Both this and another implementation of spatial filtering are essential elements of the work done for this thesis (see section 4.3).

Apart from the adjustable nonlinearity and dispersion that comes with guidance in a hollow core, capillary waveguides offer other important advantages over conventional optical fibre. The first is that their dispersion is significantly reduced, even when compared to dispersion-compensated fibre [75]. Furthermore, the guiding bandwidth for all modes is many times larger. While multi-mode optical fibres can guide very broadband light, the presence of many different modes leads to strong inter-modal dispersion and interference. In contrast, every mode in a hollow capillary can contain effectively the entirety of the optical part of the electromagnetic spectrum, though

subject to the frequency-dependent loss. This is especially important when using very short pulses, whose bandwidth can span over an octave in frequency. Lastly, it is essentially impossible to guide pulses with high peak powers in any fibre with a solid core. Because of the strong nonlinearity of glass as compared to gas or vacuum [76], self-focusing quickly leads to catastrophic damage to the solid core of a standard fibre. For high-power applications, a hollow core is therefore essential.

The main drawback of a hollow capillary waveguide is its relatively large loss, even assuming a perfectly straight capillary. Whereas solid-core fibre can be bent significantly while maintaining very good guidance, capillaries suffer twofold from even a small deviation: the loss increases rapidly with decreasing bend radius [73], and the waveguide modes become coupled, meaning that energy is transferred between them. Therefore great care must be taken to keep capillaries straight.

Hollow-core photonic crystal fibres

The limitations of both traditional solid-core fibre as well as hollow capillaries have inspired research into a great range of novel waveguide geometries. The first demonstration of light guidance in a micro-structured fibre, commonly known as photonic crystal fibre (PCF), was in 1996 [77]. While the guidance mechanism in this first type of PCF is very similar to that in a step-index fibre, the novelty lies in the underlying reason for the refractive index difference between the core and the cladding. Instead of employing a lower-index material, the cladding is made by creating holes in a regular pattern around the (solid) core. Thus the average, or effective, refractive index of the cladding region is lowered, and the fibre can guide light via index guiding.

A more radical departure from the index-guiding mechanism is offered by photonic band gap (PBG) fibre [78]. This type of fibre guides light in a hollow core, but does not rely on a difference in average refractive index, as all other types of fibres discussed so far do. Light is confined to the core by the presence of a photonic band gap, which is analogous to the concept of a band gap in solid-state physics: whereas electronic band gaps are formed by the Bragg reflection of the electron wavefunction off of atomic

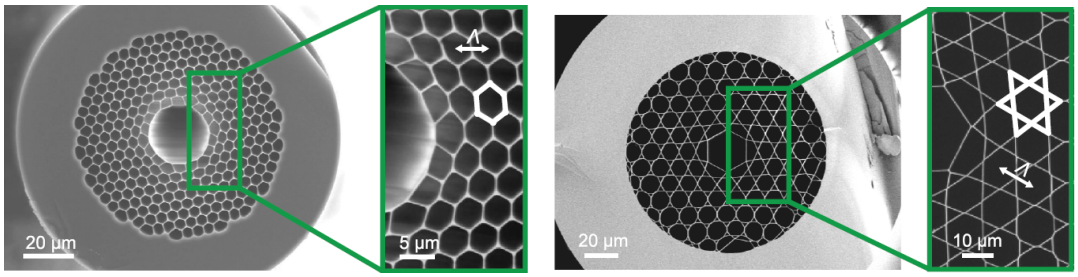

(a) (b)

Figure 3.2.: Scanning electron micrographs of the cross-section of a typical PBG fibre (a) and a kagomé-style anti-resonant PCF (b). Adapted with permission from reference [7], Optical Society.

centres in a crystalline material, a photonic band gap is due to Bragg reflection of the electric field by a periodic arrangement of refractive index changes [79] (see fig. 3.2a). In the cladding region, light at certain frequencies can not propagate, and therefore it remains in the core.

The particular pattern forming the photonic crystal in the cladding can vary substantially, with different structures optimising different parameters [80]. This type of waveguide offers true guidance of light in a hollow core, in contrast to the weak grazing-incidence confinement in a capillary. This means that bending losses are much less extreme [81]. However, the guiding bandwidth is primarily determined by the width of the band gaps, which limits PBG fibres to a maximum bandwidth of a few hundred nanometres [82].

For applications requiring transmission over very large bandwidths in a hollow core, which includes the generation of deep-UV pulses, the use of *anti-resonant* PCF is very advantageous. These fibres typically consist of a hollow core surrounded by a lattice of thin silica struts (see fig. 3.2b). The term refers to the guiding mechanism, which is based on anti-resonant reflection of light by the cladding structure. In contrast to PBG fibres, where the confinement is due to the overall structure of the cladding region, in anti-resonant fibres it is the geometry of the individual struts that determine the

location and width of confinement spectral regions, as well as the propagation loss [83].

The confinement mechanism is analogous to reflection by a Fabry-Perot cavity. There, waves that have been reflected by the cavity boundaries different numbers of times interfere with each other. This leads to near-perfect transmission through the cavity whenever the optical path difference between reflections is an integer multiple of the wavelength. Conversely, other wavelengths are reflected efficiently, and this is known as anti-resonant reflection.

The width of the transmission resonances as well as the reflectivity at anti-resonant wavelengths is determined by the *finesse* of the cavity, which in the absence of loss is purely a function of the reflectivity of the cavity ends: the higher the reflectivity, the narrower the resonances and the higher the total reflectivity of the cavity away from the resonances. The wavelengths at which resonance occurs are set by the round-trip phase, which in turn depends on the refractive index inside the cavity as well as its length [12].

Fabry-Perot cavities exist in two main forms. They can either consist of two end-mirrors with free space between them, or be contained entirely inside a slab of dielectric material, often called an *etalon*, with the Fresnel reflection at the two interfaces forming the cavity ends. While often considered as a transmissive frequency filter, an etalon can also be seen as a mirror with a reflectivity dependent on the finesse of the cavity, and zero reflectivity at the resonances. In the case of anti-resonant PCF, it is this second incarnation of a Fabry-Perot cavity that is very informative. The thin glass struts around the waveguide's core effectively form etalons which extend along the axial direction, i.e., along the direction of propagation. It is thus the core walls themselves, rather than the overall structure, that contribute to the guiding mechanism.

While the reflectivity of glass at normal incidence is very low, at grazing incidence it increases dramatically[3]. The grazing-incidence finesse of the thin Fabry-Perot cavities surrounding the waveguide core, and consequently their reflectivity, can therefore be very high. This kind of waveguide can be seen as a modified capillary in which

[3]Recall that this is how hollow capillaries confine light.

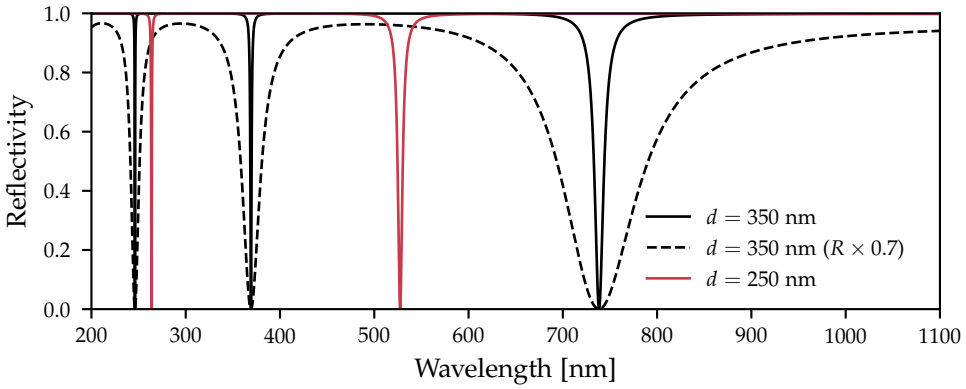

Figure 3.3.: A simple model for the reflectivity of the Fabry-Perot etalon formed by the cladding of an anti-resonant waveguide. The longitudinal and transverse wave vectors are calculated from the Marcatili model, and the phase difference calculated using the refractive index of silica and the transverse wave vector [83, 84]. The reflectivity of the cavity ends is taken as the glancing-incidence Fresnel reflectivity at the vacuum-silica interface. The black solid line shows the reflectivity of a 250 nm long cavity, and the black dashed line the same but with the Fresnel reflectivity reduced by a factor of 0.7, showing the effect of lowering the finesse of the cavity. The red line shows the reflectivity of a 350 nm long cavity, demonstrating the shift of the guidance bands in anti-resonant PCF with changing strut thickness.

the cladding is formed not of solid glass but of Fabry-Perot etalons, increasing the reflectivity of the core-cladding interface at off-resonant (or anti-resonant) frequencies. The resulting wavelength-dependent reflectivity as calculated using a very simple model is shown in fig. 3.3.

As a result of the increased reflectivity, propagation losses in anti-resonant PCF are generally much lower than in capillaries, especially when the core diameter is small; using eq. (3.1.8), the loss at 790 nm in a perfectly straight capillary of 50 µm core diameter can be calculated as around 37.5 dB m^{-1}, whereas an anti-resonant fibre may have a loss below 1 dB m^{-1} [85]. In addition, bend-losses are present but, for many

common anti-resonant structures, significantly less severe than for capillaries, making them easier to use in practice [86].

However, this significant enhancement of the guiding properties when using anti-resonant reflection comes at a price. As with any Fabry-Perot cavity, the cladding structure has resonance wavelengths at which it transmits light very efficiently. Close to these *cladding resonances*, anti-resonant PCF does not guide light. The resonances are visible in fig. 3.3 as narrow dips in the reflectivity. By changing the thickness of the cladding struts, the resonances can be shifted in wavelength to lie outside the regions where confinement is required. However, since this involves chemically etching the fibre or fabricating a different fibre altogether [87], it is substantially harder to switch between different wavelength regions than when using a capillary. Furthermore, the cladding struts have to be very thin, around hundreds of nanometres, for the fibre to exhibit very broadband guiding, because the distance in frequency between cladding resonances increases with decreasing strut thickness [83,88]. This makes manufacturing anti-resonant fibre challenging in general.

Anti-resonant PCF, like PBG fibre, exists in several different forms. The first such fibre manufactured is known as kagomé PCF, since the pattern formed by its glass struts resembles a Japanese basket-weaving pattern of the same name [89]. It is characterised by a Star-of-David pattern in the cladding structure (see fig. 3.2b). Other structures with very similar guiding properties are possible, some of which are easier to manufacture than the kagomé variant. The common feature among them is that the core is surrounded by very thin walls to create the anti-resonant effect [90].

The similarity between anti-resonant guiding and confinement in a capillary has an important, somewhat unexpected consequence: away from the cladding resonances, the dispersion characteristics of anti-resonant fibre are very well approximated by the capillary model eq. (3.1.5) [83]. While empirical corrections can be made to correct for errors in the approximation at long wavelengths [91], all the general features of the model apply except for the propagation loss, facilitating comparisons to capillary waveguides.

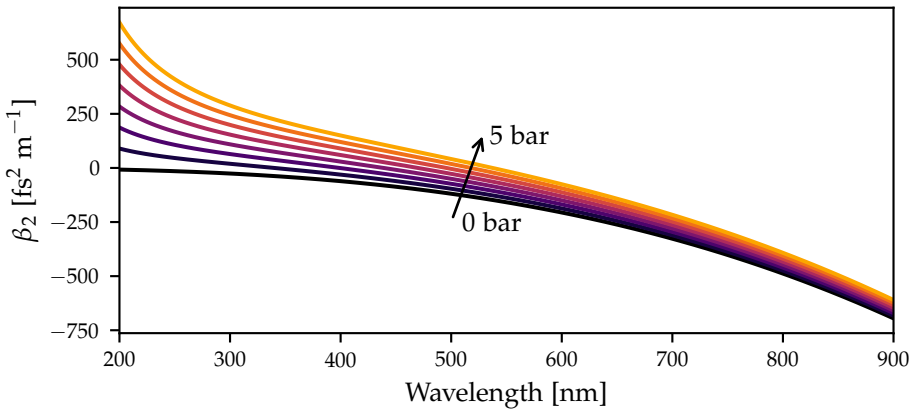

Figure 3.4.: The GVD β_2 as a function of wavelength in a capillary or anti-resonant PCF with a core diameter of 33 µm filled with argon gas to 8 different pressures from 0 bar to 5 bar.

The scaling of the anomalous GVD with the core diameter shown in eq. (3.1.7) is of particular importance. Since the superior loss characteristics of anti-resonant waveguides makes the use of relatively small core diameters feasible, it enables the propagation of laser pulses in strongly anomalously dispersive hollow-core waveguides. Moreover, the exact amount of dispersion can be tuned by filling the fibre with gas (see fig. 3.4). In combination with the nonlinearity offered by such a gas fill, which is also enhanced by the smaller core diameter and higher intensity, this paves the way for precisely controlled nonlinear dynamics, the basis of the deep-UV pulse generation mechanism discussed in this thesis (see section 3.3).

3.2. Nonlinear optics in hollow-core fibres

On a microscopic scale, the nonlinear optical effects introduced in section 2.2 are present in the same way for light confined in a waveguide. However, the long interaction length as well as the particular spatio-temporal and dispersion characteristics open up a new range of effects which cannot be easily observed using free-space nonlinear optics. In order to understand and model these effects, it is useful to derive a propagation equation, which can be used to numerically model the nonlinear processes inside a hollow-core fibre with high accuracy [7, 92–94]. With some approximations, the propagation equation can be further simplified to create a model which captures many of the phenomena discussed in this thesis on a more intuitive level.

The forward Maxwell equation

Only one additional approximation is required for the scalar nonlinear electromagnetic wave equation to become useful to numerically model the nonlinear dynamics in gas-filled hollow-core fibres. By choosing a direction of propagation (commonly z) and neglecting all waves that move opposite, it becomes possible to use a moving frame of reference for the propagation equation. This greatly reduces the computational cost of solving the equation, since only a small window of time that travels with the pulse has to be considered [94].

Starting from the scalar wave equation including the linear and nonlinear polarisation,

$$\nabla_\perp^2 E + \partial_z^2 E - \frac{1}{c^2}\partial_t^2 E = \mu_0 \partial_t^2 (P_\mathrm{L} + P_\mathrm{NL}) \,, \tag{3.2.1}$$

where $\nabla_\perp^2 = \partial_x^2 + \partial_y^2$ is the transverse Laplacian, we first Fourier transform in time as well as both transverse coordinates x and y, and introduce the linear polarisation as in eq. (2.1.20). By further defining the refractive index as in eq. (2.1.3), we arrive at

$$\partial_z^2 E(k_x, k_y, z, \omega) - k_\perp^2 E(k_x, k_y, z, \omega) + k^2(\omega) E(k_x, k_y, z, \omega) = -\frac{\omega^2}{\varepsilon_0 c^2} P_\mathrm{NL}(k_x, k_y, z, \omega) \,, \tag{3.2.2}$$

where $k^2 = \omega^2/c^2 n^2(\omega)$ as before and $k_\perp^2 = k_x^2 + k_y^2$. Similar to eq. (3.1.5), we can define

the forward wave vector (i.e., the propagation constant) as

$$\beta = k_z = \sqrt{k^2 - k_\perp^2}\,, \tag{3.2.3}$$

and thus the wave equation becomes

$$\left(\partial_z^2 + \beta^2\right) E(k_x, k_y, z, \omega) = -\frac{\omega^2}{\varepsilon_0 c^2} P_{\text{NL}}(k_x, k_y, z, \omega)\,. \tag{3.2.4}$$

To isolate the forward-propagating part of the solution, we re-write the factor $\partial_z^2 + \beta^2$ as $(\partial_z + i\beta)(\partial_z - i\beta)$. If we then assume that the solution $E(\omega)$ can be written as the superposition of forward- and backward-propagating parts [94]:

$$E(\omega) = A_+ \mathrm{e}^{i\beta z} + A_- \mathrm{e}^{-i\beta z}\,, \tag{3.2.5}$$

and further that the backward wave is much weaker than the forward one, that is, $|A_+| \gg |A_-|$, then we can approximate the z-derivative using only the forward part. Therefore, the second-order derivative with respect to z disappears:

$$\partial_z^2 + \beta^2 \approx 2i\beta(\partial_z - i\beta) = 2i\beta\partial_z + 2\beta^2\,, \tag{3.2.6}$$

which leads to the forward Maxwell equation (FME):

$$\partial_z E(k_x, k_y, z, \omega) = i\beta E(k_x, k_y, z, \omega) + \frac{i\omega^2}{2\varepsilon_0 c^2 \beta} P_{\text{NL}}(k_x, k_y, z, \omega)\,. \tag{3.2.7}$$

While this way of rejecting the backward wave is intuitive, it is not mathematically rigorous. A more complete derivation can be found in the literature [93]. It should be noted that since the second-order derivative is neglected, this procedure is equivalent to the *slowly evolving wave* approximation. The connection lies in the fact that a backwards-travelling component would lead to fast-varying oscillations in the field due to forward-backward interference. A slowly evolving wave therefore necessarily travels predominantly in one direction. We can now transform to a frame of reference that moves with the group velocity of the pulse, $v_\text{g} = \beta_1^{-1}$:

$$t \to t - \beta_1 z \tag{3.2.8}$$

$$E(\omega) \to E(\omega)\mathrm{e}^{i\beta_1 z\omega} \tag{3.2.9}$$

$$\partial_z E(\omega) \to \partial_z E(\omega)\mathrm{e}^{i\beta_1 z\omega} + i\beta_1\omega E(\omega)\mathrm{e}^{i\beta_1 z\omega}\,, \tag{3.2.10}$$

where eq. (3.2.9) is simply a consequence of the shifting property of the Fourier transform. With this transformation, the FME becomes

$$\partial_z E(k_x, k_y, z, \omega) = i\left(\beta - \frac{\omega}{v_g}\right) E(k_x, k_y, z, \omega) + \frac{i\omega^2}{2\varepsilon_0 c^2 \beta} P_{\text{NL}}(k_x, k_y, z, \omega)\,. \tag{3.2.11}$$

For propagation in free space or a waveguide, each transverse spatial frequency (k_x, k_y) has its corresponding propagation constant β as given by eq. (3.2.3). In a waveguide, however, the values that $k_\perp$ can take are restricted by their corresponding to discrete modes. Therefore, eq. (3.2.11) effectively splits into a discrete set of equations coupled by the nonlinear polarisation, with one equation for each mode and the corresponding value of β given by eq. (3.1.5).

In any case, using the FME requires calculating the nonlinear polarisation in reciprocal transverse and frequency space, $P_{\text{NL}}(k_x, k_y, z, \omega)$. Both in free space and waveguides, this is most easily achieved by calculating P_{NL} in real space and time and then Fourier transforming in time as well as projecting onto the spatial basis functions. In the case of free space propagation, the latter is equivalent to the transverse Fourier transform. For a waveguide, it is instead the projection onto the modes of the waveguide, and the dependence on $k_\perp$ turns into a modal index:

$$P_{\text{NL}}^{(m)}(z, \omega) = \int_0^a \int_0^{2\pi} P_{\text{NL}}(r, \theta, z, \omega) \hat{e}_m^*(r, \theta)\, r \mathrm{d}r\, \mathrm{d}\theta\,. \tag{3.2.12}$$

Here, we have switched to cylindrical coordinates again as the natural basis for a hollow-core waveguide, and $\hat{e}_m(r, \theta)$ is the transverse field profile of the m^{th} mode as in eq. (3.1.1), normalised to form an orthonormal set with the other modes:

$$\int_0^a \int_0^{2\pi} \hat{e}_m^*(r, \theta) \hat{e}_n(r, \theta)\, r \mathrm{d}r\, \mathrm{d}\theta = \delta_{mn}\,, \tag{3.2.13}$$

where δ_{mn} is the Kronecker delta. The nonlinear interaction as encapsulated in P_{NL} is the only part of the system of equations that couples the modes of the waveguide; that is, without the nonlinearity, there is no transfer of energy from one mode of the waveguide to another. At moderate intensities, the nonlinear contribution to the overall polarisation of the medium is small. In that case, very little energy is transferred to higher-order modes (assuming only the fundamental mode is present at the entrance

of the waveguide), and their contribution to the nonlinear polarisation can be safely ignored, leading to substantial computational speed-up.

For the Kerr effect, the nonlinear polarisation P_{NL} takes a particular functional form in that it only involves the third power of the electric field (see section 2.2.1). Therefore, the spatial projection in eq. (3.2.12) results in a constant factor multiplying the nonlinear polarisation. This factor can be eliminated by re-scaling the electric field E, leaving us with the *single-mode* FME:

$$\partial_z E(z,\omega) = i\left(\beta(\omega) - \frac{\omega}{v_g}\right) E(z,\omega) + \frac{i\omega^2}{2\varepsilon_0 c^2 \beta(\omega)} P_{\text{NL}}(z,\omega)\,, \tag{3.2.14}$$

where $\beta(\omega)$ is the propagation constant of only the fundamental mode (and thus depends only on frequency), and $P_{\text{NL}}(z,t)$ can be calculated as a function of only time (and not space) and Fourier transformed to yield $P_{\text{NL}}(z,\omega)$. Importantly, because the field E has been re-scaled to incorporate the constant resulting from the projection to the fundamental mode, the total power in the field is given by

$$W(t) = \frac{\varepsilon_0 c}{2} A_{\text{eff}} |E(t)|^2\,, \tag{3.2.15}$$

where A_{eff} is the *effective area*, which is purely a function of the modal field distribution [71]. For a detailed derivation of the single-mode equation and the necessary re-scaling of the electric field leading to the definition of the effective area, see appendix A.

The nonlinear Schrödinger equation

While it accurately reflects the nonlinear dynamics in a hollow-core waveguide with very few approximations, the FME offers very little intuition or analytical insight. To gain these, it is necessary to add further assumptions and approximations. The resulting equation, known as the nonlinear Schrödinger equation (NLSE), can be used to explicitly derive many of the phenomena which were explored and utilised during the work for this thesis. It should be noted that starting from the FME is not the most common way of deriving the NLSE, nor is the derivation presented here completely rigorous. However, it directly shows the connection between the more fundamental

propagation equation and the analytical model. More complete treatments can be found in the literature [71, 93, 94].

The first approximation on the way to the NLSE is to only consider the Kerr effect and ignore any harmonic generation processes. As a consequence, we can express $P_{\mathrm{NL}}(\omega)$ relatively simply, and the FME (in the laboratory frame) reads:

$$\partial_z E(\omega) = i\beta(\omega)E(\omega) + \frac{i\omega^2}{2\varepsilon_0 c^2 \beta(\omega)} \frac{3}{4} \varepsilon_0 \chi^{(3)} \mathcal{F}_t \Big[|E(t)|^2 E(t) \Big](\omega) \,. \tag{3.2.16}$$

Instead of the field $E(t)$, we consider the complex power envelope $A(t)$ by transforming to a frame rotating at the central frequency ω_0 and integrating over the transverse plane. Due to the rescaling of the field when deriving the single-mode FME, this integral involves the effective area (see appendix A). The field after the frame transformation and transverse integration is thus given by

$$A(t) = \sqrt{\frac{\varepsilon_0 c}{2} A_{\mathrm{eff}}}\, \mathrm{e}^{i\omega_0 t}\, E(t) \,, \tag{3.2.17}$$

where ω_0 is the central frequency of the field. $|A(t)|^2$ is thus the time-dependent power of the pulse. In the frequency domain, this leads to

$$E(\omega) = \left(\frac{c\varepsilon_0}{2} A_{\mathrm{eff}} \right)^{-\frac{1}{2}} A(\omega - \omega_0) \,. \tag{3.2.18}$$

The FME thus becomes

$$\begin{aligned} \partial_z A(\omega - \omega_0) = {} & i\beta(\omega) A(\omega - \omega_0) \\ & + \left(\frac{c\varepsilon_0}{2} A_{\mathrm{eff}} \right)^{-1} \frac{i\varepsilon_0 \omega^2 n^2 n_2}{2c\beta(\omega)} \mathcal{F}_t \Big[|A(t)|^2 A(t) \Big](\omega - \omega_0) \,, \end{aligned} \tag{3.2.19}$$

where we have introduced the nonlinear refractive index $n_2 = 3/4 \left(\varepsilon_0 c n^2\right)^{-1} \chi^{(3)}$, and the Fourier transform in the last term is now evaluated at $\omega - \omega_0$ due to the shifting property.

The next approximation is to Taylor expand the propagation constant β to second order around ω_0 as detailed in section 2.1.3:

$$\beta(\omega) \approx \beta_0 + \beta_1 (\omega - \omega_0) + \frac{1}{2} \beta_2 (\omega - \omega_0)^2 \,. \tag{3.2.20}$$

In the nonlinear term, we can instead approximate $\beta(\omega)$ to be the wave vector in the core, $\beta(\omega) \approx n\omega/c$. By using the baseband frequency $\tilde{\omega} = \omega - \omega_0$, we can then re-write the propagation equation as

$$\begin{aligned}\partial_z A(\tilde{\omega}) = i\beta_0 A(\tilde{\omega}) + i\beta_1 \tilde{\omega} A(\tilde{\omega}) + i\frac{\beta_2}{2}\tilde{\omega}^2 A(\tilde{\omega}) \\ + i\frac{n_2}{A_{\text{eff}}}\frac{n}{c}(\tilde{\omega}+\omega_0)\mathcal{F}_t\Big[|A(t)|^2 A(t)\Big](\tilde{\omega})\,,\end{aligned} \tag{3.2.21}$$

where we have used the fact that many of the constants cancel. After the rotating frame in time, we also transform to a rotating frame along the propagation direction:

$$A(\tilde{\omega}) \to A(\tilde{\omega})\mathrm{e}^{i\beta_0 z} \tag{3.2.22}$$

$$\partial_z A(\tilde{\omega}) \to \left[i\beta_0 A(\tilde{\omega}) + \partial_z A(\tilde{\omega})\right]\mathrm{e}^{i\beta_0 z}\,, \tag{3.2.23}$$

which removes the term proportional to β_0. Furthermore, we use the retarded time as in eqs. (3.2.8) to (3.2.10) to remove the term proportional to β_1:

$$t \to t - \beta_1 z \tag{3.2.24}$$

$$A(\tilde{\omega}) \to A(\tilde{\omega})\mathrm{e}^{i\beta_1 z\tilde{\omega}} \tag{3.2.25}$$

$$\partial_z A(\tilde{\omega}) \to \left[i\beta_1\tilde{\omega}A(\tilde{\omega}) + \partial_z A(\tilde{\omega})\right]\mathrm{e}^{i\beta_1 z\tilde{\omega}}\,. \tag{3.2.26}$$

This leaves us with an equation involving only GVD and the nonlinear term:

$$\partial_z A(\tilde{\omega}) = i\frac{\beta_2}{2}\tilde{\omega}^2 A(\tilde{\omega}) + i\frac{n_2}{A_{\text{eff}}}\frac{n}{c}(\tilde{\omega}+\omega_0)\mathcal{F}_t\Big[|A(t)|^2 A(t)\Big](\tilde{\omega})\,. \tag{3.2.27}$$

Any terms proportional to $\tilde{\omega}$ will have the same effect as the group velocity represented by β_1, since β_1 is the linear term in the Taylor expansion of $\beta(\omega)$. The term $\tilde{\omega}|A|^2$ is therefore the nonlinear shift to the group velocity, leading to self-steepening, whereas the term $\omega_0|A|^2$ is the shift to the phase, leading to self-phase modulation.

Fourier transforming to the time domain, the terms involving $\tilde{\omega}$ and $\tilde{\omega}^2$ turn into first- and second-order derivatives with respect to time[4]. Furthermore, for sufficiently narrow-band pulses, so that $\tilde{\omega} \ll \omega_0$, we can ignore the frequency dependence of the

[4]With the Fourier transform as defined in eq. (2.1.14), $\mathcal{F}_t[\partial_t f(t)](\omega) = -i\omega\mathcal{F}_t[f(t)](\omega)$.

refractive index for the nonlinear term, leading to

$$\partial_z A(t) = -i\frac{\beta_2}{2}\partial_t^2 A(t) + i\frac{n_2}{A_{\text{eff}}}\frac{n_0}{c}(\omega_0 + i\partial_t)|A(t)|^2 A(t) \tag{3.2.28}$$

$$= -i\frac{\beta_2}{2}\partial_t^2 A(t) + i\gamma\left(1 + \frac{i}{\omega_0}\partial_t\right)|A(t)|^2 A(t)\,, \tag{3.2.29}$$

where n_0 is the refractive index at the central frequency, and in the second line we have introduced the *nonlinear coefficient* $\gamma = A_{\text{eff}}^{-1} k_0 n_2$ with $k_0 = \omega_0 n_0/c$. For long pulses, we can ignore self-steepening, which in the time domain is represented by the term involving the first derivative. This approximation breaks down when the pulse becomes short, even if the peak power is the same. With this final approximation, we arrive at the NLSE:

$$i\frac{\partial A}{\partial z} - \frac{\beta_2}{2}\frac{\partial^2 A}{\partial t^2} + \gamma|A|^2 A = 0\,. \tag{3.2.30}$$

Note that without the nonlinear term, this equation is indeed formally identical to the free-particle Schrödinger equation as used in quantum mechanics, albeit with position and time swapped.

By approximating the dispersion relation $\beta(\omega)$ and considering only a subset of possible nonlinear processes, we have found a propagation equation which can be expressed succinctly in the time domain without the need for convolution integrals. This makes the NLSE much more useful in qualitative or analytical analysis. In particular, the NLSE has analytical solutions representing solitons, which describe the process underlying the DUV pulse generation described in this thesis (see section 3.3).

3.3. Soliton self-compression

An intense laser pulse propagating in a gas-filled capillary or fibre will be reshaped by several effects. In most circumstances, the dominant ones are dispersion and the Kerr effect. While THG will occur as well, it is usually not phase-matched and its overall influence is small. As mentioned in section 2.2.2, the spectral broadening caused by SPM due to the Kerr effect can be used to temporally compress pulses, as long as the positive chirp introduced during the broadening can be compensated [95]. This technique is widely used, including in the work for this thesis (see section 4.3.1). In many cases, the dispersion of the waveguide is ignored. This is because SPM pulse compression is most often carried out using capillary waveguides, which necessitates a large ($> 200\,\mu m$) core diameter to reduce propagation losses, leading to very weak dispersion.

In a smaller waveguide obeying the capillary model, such as anti-resonant fibre, the dispersion is strongly anomalous. This means that as the pulse spectrum broadens, the positive chirp due to SPM is counteracted by the dispersion of the waveguide. Since a broader spectrum can support a shorter pulse, and the waveguide dispersion moves the pulse closer to the FTL, the pulse self-compresses as it propagates. Furthermore, a shorter pulse has a higher peak power, so that the strength of the nonlinear interaction increases, leading to even faster spectral expansion. The relative strengths of the dispersive and nonlinear influences on the pulse determine the end result; if the anomalous dispersion is much stronger than the positive chirp induced by the nonlinearity, it will dominate and the pulse will stretch, whereas the opposite case will lead to spectral expansion without self-compression[5]. It is in the intermediate region, with the nonlinearity and linear dispersion of similar magnitude, where self-compression and the phenomena related to it can occur.

The above description of self-compression may appear to contain a contradiction: the influence of the nonlinear interaction increases as the pulse compresses due to the rising peak power, but the anomalous dispersion is a constant of the waveguide (as

[5]The latter is the parameter region in which traditional hollow fibre pulse compression operates.

long as the gas pressure is constant), so it stands to reason that eventually, the SPM chirp will outstrip the linear dispersion and the SPM-dominant of the two limiting scenarios is reached. However, the quantities that need to be compared are the relative phase shifts per unit length in the same domain, i.e., either in time or frequency. For a Gaussian pulse with a half width of τ, the nonlinear temporal phase shift after a distance L due to SPM is given by (see eqs. (2.2.14) to (2.2.18))

$$\phi_{\mathrm{NL}}(t,L) = L\beta_0 n_2 I(t) \approx L\beta_0 I_0 n_2 \left(1 - \frac{t^2}{\tau^2}\right), \tag{3.3.1}$$

so its magnitude becomes larger as the pulse duration τ decreases and the peak intensity I_0 increases correspondingly. The temporal phase shift caused by dispersion, on the other hand, is

$$\phi_{\mathrm{GVD}}(t,L) = \frac{L\beta_2 t^2}{\tau^4 + (L\beta_2)^2}, \tag{3.3.2}$$

which *also* increases for shorter pulses (for a derivation of eq. (3.3.2) see appendix A). This is because for a shorter pulse, the extreme edges of the frequency spectrum of the pulse are further apart, so that for a constant β_2 the GVD-induced delay between them is larger. Thus, the dispersion can continue to compete with the nonlinear chirp even after some self-compression has taken place.

Given the positive feedback mechanism of self-compression, the naïve assumption would be that it leads to an infinitely short pulse. While this is physically infeasible for many reasons – higher-order dispersion and the fibre guiding bandwidth, among others – it does not appear even in a simplified model such as the NLSE, which ignores those constraints. To better understand the evolution of a pulse as it self-compresses, it is useful to analyse the situation using the NLSE, which reduces the situation to only the influences of SPM and GVD.

What is needed initially is a more quantitative way of comparing the influence of dispersion and SPM. To this end, we can re-scale the time coordinate in the NLSE by the initial pulse duration T_0 and use the normalised time $\tilde{t} = t/T_0$, as well as normalising the pulse by its initial peak power P_0 through $A(z,t) = \sqrt{P_0}\,U(z,t)$, yielding

$$i\frac{\partial U}{\partial z} - \frac{\beta_2}{2T_0^2}\frac{\partial^2 U}{\partial \tilde{t}^2} + P_0\gamma|U|^2 U\,. \tag{3.3.3}$$

From this, we can identify two different characteristic length scales which depend on the pulse parameters P_0 and T_0, the *dispersion length* L_D and the *nonlinear length* L_NL, given by

$$L_\mathrm{D} = \frac{T_0^2}{|\beta_2|}, \quad L_\mathrm{NL} = \frac{1}{\gamma P_0}. \tag{3.3.4}$$

The NLSE for anomalous dispersion ($\beta_2 < 0$) then reads

$$i\frac{\partial U}{\partial z} + \frac{1}{2L_\mathrm{D}}\frac{\partial^2 U}{\partial \tilde{t}^2} + \frac{1}{L_\mathrm{NL}}|U|^2 U. \tag{3.3.5}$$

The physical meaning of L_D and L_NL is that the shorter the respective length, the faster the corresponding effect influences the pulse as it propagates along the fibre. Note that for a constant pulse energy, P_0 and T_0 are not independent; for simple pulse shapes they are in fact inversely proportional to each other. The two limiting cases of dominating SPM or GVD can be identified by one of the characteristic lengths being much larger than the other.

As mentioned before, the NLSE can be solved analytically. This is done using the inverse scattering transform [96,97]. In the solution, a critical parameter is the ratio of the nonlinear and dispersion lengths,

$$\frac{L_\mathrm{D}}{L_\mathrm{NL}} = \frac{\gamma P_0 T_0^2}{|\beta_2|} = N^2, \tag{3.3.6}$$

since its square root N determines the number of bound solutions, that is, those that remain the same for all points along z. More precisely, it is the integer content of N, so that always $\lfloor N \rceil$ bound solutions exist, where $\lfloor ... \rceil$ denotes the nearest-integer function. It is these solutions that represent solitons. If $L_\mathrm{NL} > 4L_\mathrm{D}$, no solitons exist, because dispersion stretches the pulse before significant SPM can occur [98]. A *fundamental* soliton is formed when the dispersion and nonlinear lengths are near-equal or equal so that $\lfloor N \rceil = 1$, and there is only one bound solution to the normalised equation. The pulse then takes a particularly simple form:

$$A_\mathrm{s}^{(1)}(z,t) = \sqrt{P_0}\,\mathrm{sech}\left(\frac{t}{T_0}\right)\mathrm{e}^{i\frac{1}{2}\gamma P_0 z} = \sqrt{P_0}\,\mathrm{sech}\left(\frac{t}{T_0}\right)\mathrm{e}^{i\beta_\mathrm{s} z}. \tag{3.3.7}$$

Since P_0 and T_0, while not independent, can vary jointly, there is a continuum of possible fundamental solitons with different peak powers and pulse widths. However,

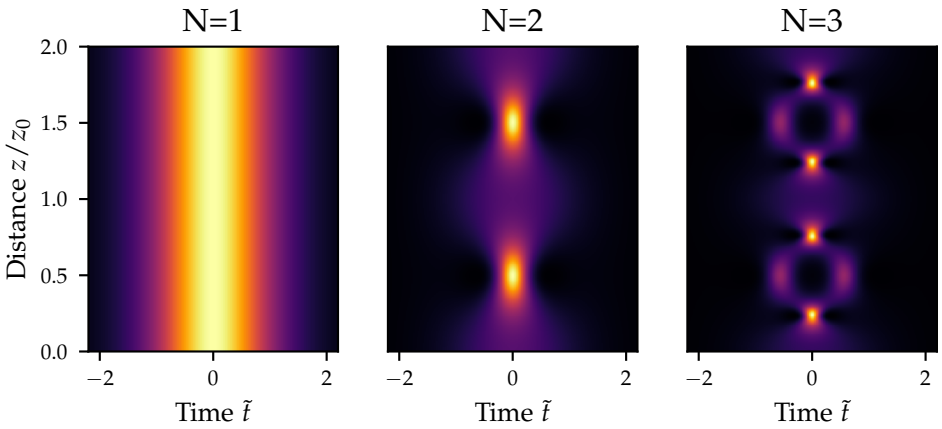

Figure 3.5.: Evolution of the fundamental ($N = 1$) and first two higher-order ($N = 2, 3$) solitons in normalised time $\tilde{t}$ and propagation direction z/z_0 for two soliton periods z_0, with the false-colour scale showing the intensity of the pulse. Note that since different soliton orders N correspond do different pulse widths T_0, the corresponding real time axis changes between plots, as does the soliton period z_0.

none of them can describe the process of self-compression, since they are stationary; the intensity envelope of the pulse does not change as it propagates. Instead, self-compression is related to *higher-order* solitons. The initial field of this class of pulses is given by

$$A_s^{(N)}(0,t) = N\sqrt{P_0}\,\mathrm{sech}\left(\frac{t}{T_0}\right) \tag{3.3.8}$$

when N is an integer. The subsequent evolution is significantly more complex than for the fundamental soliton. Instead of retaining a constant shape and accumulating a simple phase per unit length, the pulse profile undergoes periodic reshaping (see fig. 3.5). The period of the evolution is determined by the dispersion length as

$$z_0 = \frac{\pi}{2} L_D\,. \tag{3.3.9}$$

Since the behaviour of a higher-order soliton when compared with the fundamental one is determined by N, it is often referred to as the *soliton order* or *soliton number*. Importantly, the initial evolution of a higher-order soliton consists of self-compression,

with a higher soliton number corresponding to more rapid contraction and a shorter minimum pulse duration. The behaviour of these pulses shows clearly that even in the presence of only SPM and GVD, self-compression does not lead to an infinitely short pulse. However, two important issues have to be addressed: Firstly, so far we have only considered the situation of integer N. In practice, most pulses launched into a fibre will have a non-integer value for N as calculated by eq. (3.3.6). Secondly, pure higher-order solitons are created only when the pulse at the beginning of the fibre has a $\mathrm{sech}(t)$ shape as shown in eq. (3.3.8). Laser pulses are often approximated as Gaussian pulses instead, and real pulses are rarely very close to even that idealisation, often containing pre- or post-pulses as well as long pedestals.

The answers to the two questions above are related in that they both involve the evolution of pulses which do not fulfil the conditions for a pure soliton to be formed. This situation can be analysed using the direct scattering transform, whose inverse is the inverse scattering transform. For the simpler case of a pulse with an initial $\mathrm{sech}(t)$ profile in time but non-integer N, the result is that the field splits into two components: a soliton of order $\lfloor N \rceil$ and a weak non-solitonic pulse often referred to as the linear, or dispersive, component, since its intensity is insufficient to cause significant SPM, so it is stretched by GVD and walks off due to its differing group velocity [98]. As a consequence, the pulse appears to radiate away some energy and assume the pulse width and power corresponding to the closest soliton, streching or compressing in time until its soliton number is an integer.

When the pulse is of a different shape, something very similar occurs, except that in addition to width and power changes and the appearance of weak dispersive radiation, the shape of the pulse also changes and asymptotically approaches that of the nearest soliton [98]. To determine which soliton that is, the soliton number has to be calculated. Since eq. (3.3.6) involves the parameters T_0 and P_0 as related to the $\mathrm{sech}(t)$ pulse shape, an analogous definition has to be found for other pulse shapes. In practice, and for reasonably simple input pulse shapes, it is often sufficient to simply find the intensity FWHM duration of the initial pulse and use this as the FWHM duration $T_0^{(\mathrm{FWHM})}$ for

the "equivalent" $\operatorname{sech}^2(t)$ intensity profile, which is related to T_0 by

$$T_0^{(\mathrm{FWHM})} = 2T_0 \cosh^{-1}\left(\sqrt{2}\right) \approx 1.763\, T_0\,. \tag{3.3.10}$$

Using this together with the peak power P_0 yields an estimate for N [71].

The underlying reason for this asymptotic behaviour as well as the periodic evolution of higher-order solitons can be found in the direct scattering transform analysis. Through this lens, any input pulse propagated with the NLSE can be decomposed into $\lfloor N \rceil$ fundamental solitons (since $\lfloor N \rceil$ determines the number of bound solutions to the NLSE) as well as a radiative part. Any pulse with $N > 1$, integer or not, consists of several fundamental solitons, whose phase velocities are different. For non-integer N or shapes different from a sech pulse, the amplitude of the radiative part is non-zero.

The simplest example of this superposition principle is the pure first higher-order soliton, for which $N = 2$. Its initial pulse profile, $A_\mathrm{s}^{(2)}(0,t)$ is simply the sum of two fundamental solitons:

$$A_\mathrm{s}^{(2)}(0,t) = 2\sqrt{P_0}\operatorname{sech}\left(\frac{t}{T_0}\right) = 2A_\mathrm{s}^{(1)}(0,t)\,. \tag{3.3.11}$$

Due to the nonlinear nature of the propagation, this superposition then causes the pulse evolution shown in fig. 3.5. The periodic structure is thus the result of a nonlinear interference effect. For this reason, the direct scattering transform is sometimes referred to as the nonlinear Fourier transform [98].

3.4. Dispersive wave emission

The NLSE forms an analytical model in which solitons emerge as the fundamental building blocks for the nonlinear evolution of the laser pulse in a fibre. However, this model neglects a number of effects that will occur in a real fibre, chief among which is higher-order dispersion. Others are the fibre loss, self-steepening, and THG, as well as plasma effects and the potential presence of any higher-order modes and the coupling between them. Adding any combination of these to the model results in a group of propagation equations, all of which are usually referred to as the generalised nonlinear Schrödinger equation (GNLSE). The GNLSE can not be solved analytically, and the exact solutions representing solitons are no longer valid, but many qualitative features that result from soliton propagation remain.

For the process of DUV emission by self-compressing pulses, the critical addition is third-order dispersion, because it enables phase-matching between the soliton and other parts of the spectrum [99]. The GNLSE including TOD can be found by simply adding the third-order term to the expansion of the wave vector in eq. (3.2.20) and proceeding in the same way, which yields

$$i\frac{\partial A}{\partial z} - \frac{\beta_2}{2}\frac{\partial^2 A}{\partial t^2} + i\frac{\beta_3}{6}\frac{\partial^3 A}{\partial t^3} + \gamma|A|^2 A = 0\,. \tag{3.4.1}$$

For very weak pulses, the nonlinear term can be neglected. This describes linear waves, such as the radiative part of the field found for non-solitonic initial pulses in the NLSE. Their propagation is simply governed by the linear wave vector; this can be shown by inserting a plane wave as a trial solution into the GNLSE without the nonlinear term:

$$A_{\mathrm{L}} = \mathrm{e}^{i(\beta_{\mathrm{L}} z - \tilde{\omega} t)}\,, \tag{3.4.2}$$

which yields the linear wave vector β_{L},

$$\beta_{\mathrm{L}}(\tilde{\omega}) = \frac{\beta_2}{2}\tilde{\omega}^2 + \frac{\beta_3}{6}\tilde{\omega}^3\,, \tag{3.4.3}$$

as is expected from the derivation of the NLSE and GNLSE. In fig. 3.6, the wave vector β in a capillary or anti-resonant fibre for a particular combination of parameters is

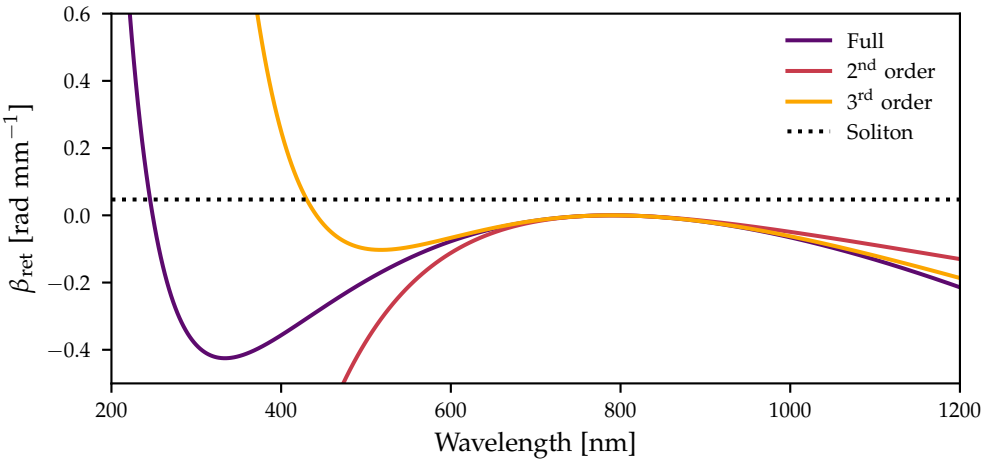

Figure 3.6.: The wave vector in the retarded frame β_{ret} for the fundamental mode in a 33 µm core diameter fibre filled with argon gas to a pressure of 4 bar, calculated using the capillary model eq. (3.1.5). The red and yellow lines show approximations to β_{ret} using 2$^{\text{nd}}$ and 3$^{\text{rd}}$-order Taylor expansions around 790 nm, respectively. The black dashed line shows the soliton wave vector β_{s} for a 15 fs pulse of 2 µJ energy centred at 790 nm, which corresponds to $N = 10.5$.

shown, transformed to the retarded frame used in the NLSE by

$$\beta_{\text{ret}}(\omega) = \beta(\omega) - (\omega - \omega_0)\beta_1 - \beta_0\,, \tag{3.4.4}$$

where ω_0 is the central frequency of the initial pulse. Also shown are the second and third-order Taylor approximations around ω_0, the latter being equivalent to the linear wave vector β_{L} in the GNLSE as given by eq. (3.4.3).

Because the soliton wave vector β_{s} as calculated from the nonlinear coefficient and the peak power according to eq. (3.3.7) is always positive ($n_2 > 0$ and thus $\gamma > 0$) and β_2 negative, there is no frequency at which $\beta_{\text{s}} = \beta(\omega)$ if only GVD is considered. Phase-matching could occur for normal dispersion, however in that case there can be no soliton dynamics. In contrast, when TOD is included, its influence dominates for frequencies far away from the expansion frequency ω_0, and the linear and nonlinear

wave vectors intersect. The phase-matching condition

$$\begin{aligned} &\beta_s = \beta_L(\tilde{\omega}) \\ \iff &\frac{1}{2}\gamma P_0 = \frac{\beta_2}{2}\tilde{\omega}^2 + \frac{\beta_3}{6}\tilde{\omega}^3 \end{aligned} \tag{3.4.5}$$

therefore has a solution [99,100]. The soliton can transfer energy efficiently to the phase-matched frequency, since the two fields remain in phase and coherent build-up can take place. This is the fundamental process underlying the DUV generation implemented in the work for this thesis. Because the phase-matching occurs between the soliton and a dispersive pulse, that is, the nonlinear phase shift of the phase-matched pulse is not considered, the process is called dispersive wave emission (DWE).

Tunability of the dispersive wave

For low soliton numbers, the nonlinear phase shift β_s is small, so that it is sometimes neglected, leading to a very simple phase-matching condition:

$$\tilde{\omega}_D = -\frac{3\beta_2}{\beta_3}. \tag{3.4.6}$$

For a more accurate estimate, the root of the polynomial equation eq. (3.4.5) can be approximated using Newton's method, which results in [99]

$$\tilde{\omega}_D = -\frac{3\beta_2}{\beta_3} + \frac{\gamma P_0 \beta_3}{3\beta_2^2}. \tag{3.4.7}$$

It should be pointed out that by the definition of $\tilde{\omega}$, this is the frequency *shift* relative to the soliton at which the dispersive wave phase-matches. Furthermore, since β_2 must be negative for this process to occur and β_3 is positive for the majority of materials including gases, $\tilde{\omega}_D$ is positive. The dispersive wave is therefore emitted at a higher frequency than the soliton, which is why it can be used to generate pulses in the DUV.

This condition is based on the assumption that the soliton wave vector can be calculated using eq. (3.3.7) even for higher-order solitons. That is an approximation, since the phase evolution for solitons with $N > 1$ is significantly more complex than this, and a simple propagation constant cannot be defined. Most importantly, the peak

power and thus the nonlinear contribution to the wave vector increases as the pulse self-compresses. Nevertheless, this form of the phase-matching equation reflects many qualitative features of dispersive wave emission, such as the blue-shift observed with higher pump powers. By numerically solving eq. (3.4.5), the phase-matched frequency can be determined even more accurately. However, a consistent blue-shift remains when simulations or experiments are compared to the analytical phase-matching calculation due to the highly nonlinear phase evolution of the higher-order soliton [101].

Estimates of the dispersive wave frequency in a real fibre exhibiting all orders of dispersion can be improved further by matching not the Taylor-expanded wave vector $\beta_L(\tilde{\omega})$, but instead the full wave vector $\beta(\tilde{\omega})$, to β_s; that is, by solving the phase-matching equation $\beta(\tilde{\omega}) = \beta_s$. The resulting difference can be seen from fig. 3.6, where the intersection points with β_s of the approximate and full linear wave vectors are different by almost 150 nm. Similar to the more simplified case, however, a blue-shift remains which can only be fully addressed by numerical modelling or experiment [101].

The phase-matching for the dispersive wave involves both the nonlinear phase of the soliton and the linear dispersion, and the dispersive wave frequency can consequently be tuned by changing either of these. However, almost all of the available experimental parameters influence both the linear and nonlinear contributions. For instance, increasing the core diameter of the fibre reduces the nonlinear coefficient γ as well as the anomalous contribution on the waveguide dispersion, while a higher gas pressure leads to more normal dispersion and a larger γ. The total input pulse energy is one of very few parameters that influence only one quantity in the phase-matching equation, namely the peak power P_0. However, it also has a strong effect on the length scale over which the self-compression process occurs, and this can adversely affect the pulse duration of the dispersive wave by adding dispersion (see section 4.2.2).

In practice, the easiest way of tuning the dispersive wave frequency is therefore the gas pressure, despite its influence on several parameters. The effect of changing gas pressure is shown in figs. 3.7a and 3.7b for a particular combination of parameters. A

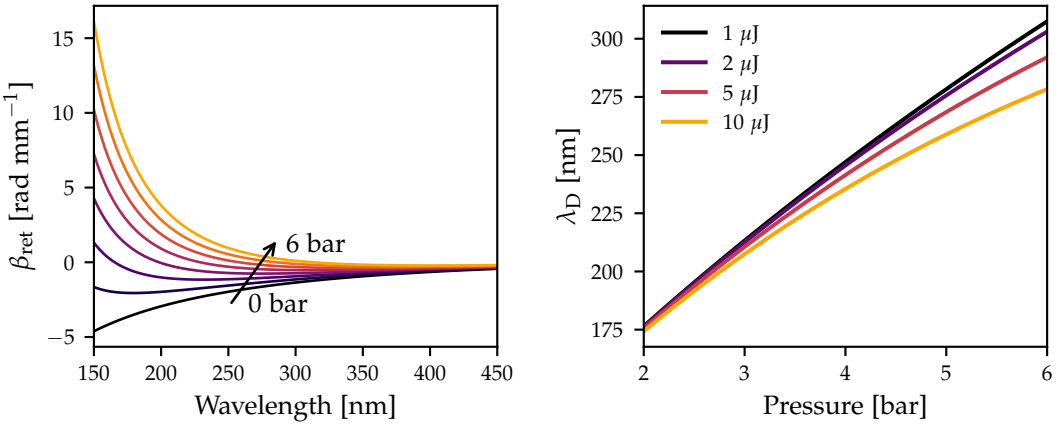

Figure 3.7.: Tunability of the dispersive wave. (a) The wave vector in the retarded frame β_{ret} for the fundamental mode in a capillary with core diameter 33 µm filled with argon to 8 different pressures from 0 bar to 6 bar. (b) The wavelength at which DWE occurs as predicted by eq. (3.4.5) in the same fibre as a function of argon pressure for an input pulse at 790 nm with 4 different energies from 1 µJ to 10 µJ, assuming a FWHM pulse duration of 15 fs and including the influence of gas pressure on both the dispersion and nonlinear coefficient γ.

higher pressure leads to phase-matching at longer wavelengths for the dispersive wave, whereas higher pulse energies lead to shorter wavelengths. Importantly, with different combinations of gas pressure and input pulse energy, the entire range from 175 nm to 300 nm can be reached.

Soliton fission

Higher-order dispersion also changes the evolution of the soliton itself. As long as the pulse is relatively long, its bandwidth is small and the influence of TOD is small enough for soliton dynamics to appear, despite the fact that the GNLSE does not admit

exact soliton solutions. When the pulse approaches the point of maximum compression, however, its bandwidth expands sufficiently for TOD to disturb it significantly. As a result, the periodic behaviour that would appear in the case of a pure soliton is interrupted. Instead of compressing, then simply stretching in time and contracting in frequency, the pulse splits in a process known as *soliton fission* [102]. The pulses that appear behave much like fundamental solitons in that they neither stretch nor expand as they propagate, and the number of pulses created is equal to the order of the soliton they originate from. They can be interpreted as the fundamental solitons which formed the original pulse [98].

Soliton fission plays an important role in the dynamics of DWE as well, since it terminates the process. Soliton fission occurs at the point of maximum temporal compression, and correspondingly maximum bandwidth. Therefore, the pulse breaks up as the dispersive wave is generated, and the distance over which it is emitted is shorter than in the absence of soliton fission. Furthermore, the relative importance of TOD and GVD depends on the bandwidth of the pulse, so that a soliton whose shortest duration is smaller will break up sooner; it reaches the point at which TOD becomes important further before its maximum compression. The shortest duration scales inversely with the soliton order, as does the propagation distance after which the first compression point is reached. This is because for $N > 1$ the nonlinear length is N times shorter than the dispersion length, so the effect of SPM is stronger than that of GVD, leading to more and faster spectral expansion and a larger bandwidth at the point of maximum compression. The distance after which soliton fission occurs can be estimated by

$$L_{\mathrm{F}} = \frac{L_{\mathrm{D}}}{N} = \sqrt{L_{\mathrm{D}} L_{\mathrm{NL}}} = \sqrt{\frac{T_0^2}{\gamma |\beta_2| P_0}}\,, \tag{3.4.8}$$

and is known as the *soliton fission length* [103].

That the fission length should scale inversely with the soliton order can be understood by considering the evolution of higher-order solitons as shown in fig. 3.5. Although the soliton period z_0 only depends on the dispersion length, and thus does not necessarily change with the soliton order, the structure within a soliton period evolves faster for

higher soliton orders. Since soliton fission occurs at the point of maximum pulse compression, and the length after which this happens becomes shorter as N rises, a soliton breaks up faster the higher its order.

An important feature of the dispersion length is that of all the fibre and pulse parameters, it scales most strongly with the pulse duration. A shorter pulse therefore requires a shorter waveguide. This is especially important when a short pulse is to be generated in the dispersive wave; if the point of soliton fission is very early, the DWE pulse is stretched by propagation in the remaining fibre.

4. Generation of Bright Ultrashort DUV Pulses

Using the range of nonlinear optical effects introduced in section 2.2, many frequency conversion schemes are possible and have been developed. While very short DUV pulses have been generated with low efficiency [104], the combined goal of tunable and efficient frequency conversion to short pulses in this wavelength range is very challenging. In particular, spectral tunability of very short DUV pulses remained out of reach until recently.

The simplest frequency conversion process, harmonic generation, limits the achievable frequencies to multiples of the driving frequency. More flexibility and shorter ultimate wavelengths can be achieved by using FWM techniques [105, 106] as well as SFG [107, 108], but still the input and output frequencies have to be commensurate. Tunability in the converted frequency is thus only possible by changing the driving frequency or frequencies, which requires a tunable laser.

One exception to this rule is formed by tunable OPAs like the one used for infrared pulse generation in our laboratory. Although the direct amplification process can only result in wavelengths longer than the initial driver, with the achievable range limited by phase-matching, frequency up-conversion within or after the parametric amplifier can yield tunable sources in many spectral regions [109]. Especially for the generation of DUV pulses, however, this approach is extremely inefficient due to the number of nonlinear conversion steps involved.

Additionally, very short DUV pulses can only be created by these processes if a nearly equally short pulse is used to drive them. The bandwidth of the initial pulse and the finite phase-matching bandwidth of any process furthermore limit the pulse

duration that can be obtained without reducing the conversion efficiency to very low values.

Dispersive wave emission overcomes several of these issues. Most importantly, a DUV pulse generated in this way can be several times shorter than the driving pulse [10,11]. In combination with reasonable conversion efficiencies ranging from 1 % to 10 % and the inherent tunability [10,87], this makes DWE a uniquely capable source. Several practical issues have to be overcome, however, with the most important one being the delivery of these short pulses to an experiment without stretching them.

4.1. Dispersive waves as a DUV source

For dispersive wave emission to be useful as a source in time-resolved experiments, it needs to be able to generate pulses that are both short in duration and sufficiently energetic. Similar to harmonic generation and other more commonly used effects, there are two components of the frequency conversion that determine the achievable pulse duration and energy; these are the phase-matching bandwidth and the mechanism by which energy is transferred from the driving pulse to the new wavelength. Both of these components work somewhat differently in dispersive wave emission than in other phenomena, however.

The dispersive wave is not generated by a direct nonlinear frequency conversion process in the same way as, for instance, harmonic generation or sum-frequency fields. This means that the nonlinear response does not immediately create a component at the dispersive wave frequency, with phase-matching determining whether efficient conversion can take place. Instead, phase-matching allows coherent build-up of the part of the pulse spectrum that extends to the dispersive wave frequency.

This mechanism is most clearly understood by considering the spectrum as a series of discrete frequency components; the SPM experienced by the pulse is then equivalent to cascading FWM interactions between all of these components, with new frequency components either side of the original set of frequencies allowing the generation of ever further-removed sidebands and thus broadening the spectrum into a wide comb as the field propagates [110]. Once the comb extends to the dispersive wave frequency, phase-matching leads to the coherent build-up of any components generated there. Crucially, even though the dispersive-wave component is the result of a cascade of FWM interactions, the phase-matching condition for this process is identical to that of the (non-existent) direct process [111].

This mechanism of transferring energy from the pump to the dispersive waves has important implications for the efficient generation of DUV pulses. In particular, the amount of energy transferred to the dispersive wave is set by how much of the original pulse energy is shifted to that frequency by SPM. Strongly asymmetric broadening is

therefore advantageous. This is in contrast to, for example, harmonic generation, where it is only the strength of the nonlinearity in combination with the interaction length that sets the efficiency at the phase-matched frequency.

Asymmetry in SPM-induced spectral expansion is caused by self-steepening (see section 2.2.2), since it breaks the time symmetry. Self-steepening, in turn, is more pronounced for shorter input pulses since their bandwidth is larger (see section 3.2). The conversion efficiency is thus higher when using a shorter input pulse [7]. Decreasing the initial pulse duration has additional benefits. Since the soliton fission length decreases proportional to T_0, a shorter fibre can be used, which reduces propagation losses before the point of DUV, further increasing the overall efficiency. Furthermore, the soliton number N decreases for a shorter pulse, even if the total pulse energy is maintained. Previous numerical work has shown that this significantly improves the quality of the pulse created in DWE, resulting in less structure in the output spectra and higher pulse contrast in the time domain [7].

Similarly to the conversion efficiency, the duration of the pulse generated in the dispersive wave is influenced by a different combination of factors when compared to other frequency conversion mechanisms. As shown in section 2.2.4, a common assumption in harmonic generation or frequency mixing schemes is that the driving field remains largely the same throughout the process. This is partly motivated simply by the usefulness of the resulting analytical phase-matching equation [13]. The interaction medium is also usually treated as unchanging, with the dispersion setting the phase-matching bandwidth and the effective nonlinearity and interaction length determining the energy conversion.

For the use of DWE to generate DUV pulses for time-resolved experiments, these assumptions are not appropriate. Most importantly, the driving pulse is reshaped during the nonlinear interaction, and the emission of a DUV pulse is contingent on this reshaping being very strong. While the driving pulse characteristics at the beginning of the fibre, in combination with the fibre parameters, can give reasonable estimates of the DUV wavelength, nothing in the analysis presented so far can be used to make

statements about the phase-matching bandwidth.

The reasoning that led to the phase-matching efficiency function in eq. (2.2.29) can be applied to the dispersive wave, and the approximations involved – that only small amounts of energy are converted and that the two fields are collinear – are sound. It should therefore be possible to calculate the phase-matching bandwidth. However, both of the two key ingredients present a challenge.

The wave vector mismatch Δk can be obtained from the soliton characteristics and the linear dispersion, but the large discrepancy between the resulting prediction of the phase-matching wavelength and numerical or experimental data implies that it is not an accurate reflection of the actual conditions at the point of DUV emission. The interaction length, on the other hand, is essentially impossible to define. With soliton fission, higher-order dispersion and potentially ionisation all gaining in importance just as the driving pulse bandwidth increases sufficiently to allow for DWE, there is no straightforward way to determine how long the soliton exists before it breaks up and energy transfer ceases.

If the gas pressure in the fibre is not constant along its length, this furthermore invalidates the assumption that the medium does not change during the frequency conversion process. The dynamics surrounding the point of dispersive wave emission are therefore even more complicated, and analysing the system in terms of its initial conditions becomes less accurate. In combination, these issues mean that a more accurate numerical model is required.

4.2. Numerical modelling

To better understand the dynamics underlying soliton formation and DWE in real systems and explore the parameter space accessible in experiments, it is necessary to go beyond the NLSE and GNLSE and simulate the situation with as few approximations as possible. To do this, I have implemented a numerical solver for the single-mode FME in a hollow waveguide, including the effects of the third-order nonlinearity as well as photoionisation. The single-mode FME, explicitly including the loss, reads:

$$\partial_z E(z,\omega) = i\left(\beta(\omega) + \frac{i}{2}\alpha(\omega) - \frac{\omega}{v_g}\right) E(z,\omega) + \frac{i\omega^2}{2\varepsilon_0 c^2 \beta(\omega)} P_{\mathrm{NL}}(z,\omega)\,. \tag{4.2.1}$$

While numerical calculations as well as measurements of the propagation loss exist for some fibres [7,86], the loss in an anti-resonant fibre is low, so that it can be approximated as a constant without affecting the propagation dynamics too much. Note that this is decidedly not the case for hollow capillaries, where the loss has to be calculated using eq. (3.1.8). The propagation constant $\beta(\omega)$ can be obtained from the capillary model using eq. (3.1.5). The cladding resonances can be included using an empirical model [112] if more precise agreement with experiment for a particular fibre is desired, however for general modelling this is not necessary.

4.2.1. Implementation

A first-order ordinary differential equation (ODE) like the FME can be solved numerically using many different methods [113, 114]. A very intuitive approach for a propagation equation is to "march" the electric field of the laser pulse forward in space in discrete steps, updating it at each step using the nonlinear and linear polarisation terms. This is where the moving reference frame introduced in section 3.2 is very useful.

The basic problem when numerically solving the FME is that the nonlinear polarisation is most easily evaluated in the time domain, while the linear propagation is more easily carried out in the frequency domain. In the respective other domain, each of these involve one or more convolution operations, which are much more computationally

expensive than multiplication. The simple solution to this problem is to evaluate the nonlinear polarisation $P_{\text{NL}}(z,t)$ in the time domain and subsequently Fourier transform it to obtain $P_{\text{NL}}(z,\omega)$.

The numerical solver for the FME uses a *Runge-Kutta* algorithm to march the solution for the electric field forward along the fibre. All Runge-Kutta methods are based on sub-dividing the discretisation step into smaller distances and evaluating the derivative (in this case, the right-hand side of the FME) at each of these points, using the result of a first-order Taylor approximation as the initial value for the function at the first sub-step [114]. As in all such methods, the size of the steps is critical for the accuracy of the model. While simply choosing a very small step size will ensure that the solution is accurate, this can make the numerical propagation unnecessarily slow. To avoid this, I use the *Dormand-Prince* method, which includes an estimate of the local truncation error (the error resulting from using a discrete step) at very little additional computational cost. The step size is then adjusted such that the error remains below a certain tolerance, if necessary by repeating iterations after decreasing the step size [115].

When applied directly to the FME, the Dormand-Prince method makes no use of the fact that the FME in the absence of nonlinear interactions,

$$\partial_z E(z,\omega) = i\left(\beta(\omega) + \frac{i}{2}\alpha(\omega) - \frac{\omega}{v_{\text{g}}}\right)E(z,\omega)\,, \tag{4.2.2}$$

has an analytical solution:

$$E(z_2,\omega) = E(z_1,\omega)\,\mathrm{e}^{i\left[\beta(\omega)+\frac{i}{2}\alpha(\omega)-v_{\text{g}}^{-1}\omega\right](z_2-z_1)}\,, \tag{4.2.3}$$

so if the field at any point in the fibre, $E(z_1,\omega)$, is known, then the field at any other point can be obtained in a single step. A common way of including this fact in a numerical simulation is to use a split-step method, in which the linear and nonlinear parts of the equation are treated separately [113]. However, this cannot be done in conjunction with the Dormand-Prince method, since the splitting itself is an approximation that places constraints on the step size; there is no way of telling whether the adaptively chosen step size is sufficiently small for the split-step to remain

accurate without testing different step sizes, negating the benefit of the embedded error estimate.

The linear part of the equation can be taken into account by using it as a *pre-conditioner* to the equation, which includes the linear propagation in the field itself [116]. This means that the Dormand-Prince method is used to solve not the FME itself, but instead the pre-conditioned equation:

$$\partial_z \bar{E}(z,\omega) = \frac{i\omega^2}{2\varepsilon_0 c^2 \beta(z,\omega)} \bar{P}_{\text{NL}}[\bar{E}(z,\omega)]\,. \tag{4.2.4}$$

The pre-conditioned field $\bar{E}(z,\omega)$ is given by

$$\bar{E}(z,\omega) = \exp\left\{-\int_0^z \mathcal{L}(z')\,\mathrm{d}z'\right\} E(z,\omega)\,, \tag{4.2.5}$$

where $\mathcal{L}(z)$ is the linear operator in the FME,

$$\mathcal{L}(z) = i\left(\beta(z,\omega) + \frac{i}{2}\alpha(\omega) - \frac{\omega}{v_{\text{g}}(z)}\right)\,, \tag{4.2.6}$$

and we have retained any potential dependence on z caused by a pressure gradient changing $\beta(\omega)$ and v_{g} along the waveguide. $\bar{P}_{\text{NL}}$ is given by

$$\bar{P}_{\text{NL}}[\bar{E}(z,\omega)] = \exp\left\{-\int_0^z \mathcal{L}(z)\,\mathrm{d}z'\right\} P_{\text{NL}}\left[\exp\left\{\int_0^z \mathcal{L}(z)\,\mathrm{d}z'\right\} \bar{E}(z,\omega)\right]. \tag{4.2.7}$$

In this way, the linear propagation is included in the pre-conditioned field $\bar{E}(z,\omega)$ and the benefits of the embedded error estimation can be used together with the analytical solution to one part of the equation. This method was originally developed for numerical solutions of the time-dependent Schrödinger equation [116] and later adapted for use in nonlinear optical pulse propagation [117]. For a derivation of the pre-conditioned equation, see appendix A.

To investigate the dynamics of soliton self-compression and DWE in detail, it is sometimes necessary to obtain values for the field $E(t)$ at very closely spaced points along z. This slows down the numerical solver if the distance between the desired points is smaller than the step size determined from the local error estimation. To retain the speed benefits of the adaptive step sizing, the step size should not be forced to be smaller than necessary. This can be done by computing the dense samples using

a Runge-Kutta triple, also known as an interpolant [118,119]. In this method, a third Runge-Kutta method is embedded in the Dormand-Prince algorithm, providing values of the solution at points between the propagation steps without requiring further evaluations of the (computationally expensive) nonlinear polarisation term.

Calculation of P_{NL}

The self-compression dynamics that occur in anti-resonant fibres can lead to extremely short pulses, and thus high intensities, leading to strong-field ionisation. A realistic model therefore has to take into account the response of both bound and free electrons. The nonlinear polarisation is calculated in the time domain and given by

$$P_{\mathrm{NL}}(t) = P_3(t) + P_{\mathrm{e}}(t) \tag{4.2.8}$$

where $P_3(t)$ is the third-order nonlinear response of the medium, and $P_{\mathrm{e}}(t)$ is the polarisation due to free electrons created by photoionisation. The third-order response, as introduced in section 2.2.1, is given by

$$P_3(t) = \frac{1}{4}\varepsilon_0\chi^{(3)}\left[3|E(t)|^2 + E(t)^2\right]E(t)\,. \tag{4.2.9}$$

The crucial parameter here is the third-order susceptibility $\chi^{(3)}$. It is calculated from the gas density and the third-order hyperpolarisability of the gas, values for which are available in the literature [120,121].

The free-electron term, on the other hand, is not as easily obtained. The most common model for the effect of plasma in ultrafast pulse propagation methods is a semi-classical one, which treats the ionisation event through one of the approximate ionisation rates introduced in section 2.2.5, and the subsequent motion of electrons with classical mechanics; the polarisation is then obtained from the resulting current [122]. Given an ionisation rate $w(E)$ (the solver uses the ADK rate), the free-electron density $\rho_{\mathrm{e}}(t)$ can be calculated as

$$\rho_{\mathrm{e}}(t) = \rho_0\left[1 - \exp\left\{-\int_{-\infty}^{t} w\left[\mathcal{E}(t')\right]\,\mathrm{d}t'\right\}\right], \tag{4.2.10}$$

with ρ_0 the gas density, as long as only neutral and singly ionised atoms are considered. In the semi-classical model, the free-electron polarisation is then given by

$$\mathcal{P}_{\mathrm{e}}(t) = \int_{-\infty}^{t} \frac{I_{\mathrm{P}}}{\mathcal{E}(t')} \frac{\mathrm{d}\rho_{\mathrm{e}}}{\mathrm{d}t'} \mathrm{d}t' + \frac{e^2}{m_{\mathrm{e}}} \iint_{-\infty}^{t} \rho_{\mathrm{e}}(t') \mathcal{E}(t') \mathrm{d}t' \mathrm{d}t' \,, \tag{4.2.11}$$

where I_{P} is the ionisation potential of the gas, e is the elementary charge, and m_{e} is the mass of the electron. The first term in eq. (4.2.11) represents the energy imparted on the electrons by the field and thus leads to energy loss. The second term represents the effect of the plasma on the phase of the field. Note that this second term is the same as the simpler expression for the free-electron polarisation for a single plane wave in eq. (2.2.35); when written in the frequency domain, the double integral over time becomes division by a factor of $-\omega^2$. Additionally, in eq. (4.2.11) the electron density is now explicitly time-dependent.

One important complication when including the effect of photoionisation is that the single-mode FME treats the spatial profile of the nonlinear polarisation with the assumption that only third-order effects are present. This assumption enables the projection onto the fundamental mode to be carried out analytically. The significantly more complex form of the free-electron term, on the other hand, does not allow this. Therefore, the single-mode FME as presented here treats the free-electron polarisation $P_{\mathrm{e}}(t)$ as if its time dependence is given by eq. (4.2.11), but its spatial dependence is the same as that of the third-order polarisation P_3. This is a significant additional approximation, however it dramatically speeds up computation and is acceptable as long as ionisation effects are not very strong [123].

Sampling for the nonlinear response

The output spectrum of a typical DWE process leading to DUV emission from an anti-resonant fibre may extend from below 200 nm to above 2000 nm, so the sampling frequency in the simulation must be sufficient to represent these frequencies accurately. At the same time, the time window must be large enough to contain all components of the field that are of interest; if the pulse stretches significantly or several components

with different group velocities are created, this must be taken into account. The minimum number of sampling points needed to satisfy these constraints as well as the Nyquist-Shannon sampling theorem [59] is given by

$$\mathcal{N} = \frac{\Delta t \Delta \omega}{2\pi}, \tag{4.2.12}$$

where Δt and $\Delta\omega$ are the required window widths in time and frequency, respectively.

Note that while real-valued signals need to be sampled at a rate of twice the highest frequency of interest, for the analytic field E the frequency window width is equal to the highest frequency of interest, $\Delta\omega = \omega_{\max}$, since the analytic signal only contains positive frequencies. This does not lead to any increase in efficiency, however, since the complex-valued field E takes up twice as much memory as the real field $\mathcal{E}$.

The sampling criterion eq. (4.2.12) is insufficient for propagation involving a third-order nonlinearity and ionisation. Because the free-electron polarisation must be calculated using the real electric field $\mathcal{E}(t)$ [122], the analytic field must first be converted to the real field by adding the complex conjugate. The conjugate contains negative frequencies, so the frequency axis has to be extended to do this.

In addition, the nonlinear polarisation contains components outside of the frequency window. This is easiest to see for the third-order polarisation, which contains frequency components extending to three times the highest frequency in the window due to the THG term. While the components outside of the window of interest are unimportant by definition, they can only be removed by first extending the frequency axis to $3\omega_{\max}$ (in both the positive and negative directions), calculating the polarisation P_3, and then discarding the added parts. For high intensities, the exponential scaling of the ionisation rate with the field strength may add even higher frequencies. Discarding the high-frequency components, rather than simply permanently extending the frequency axis and attenuating them, has the advantage of reducing the computation time required for the linear propagation, which has to be carried out at each sub-step of the Dormand-Prince method for the pre-conditioned FME.

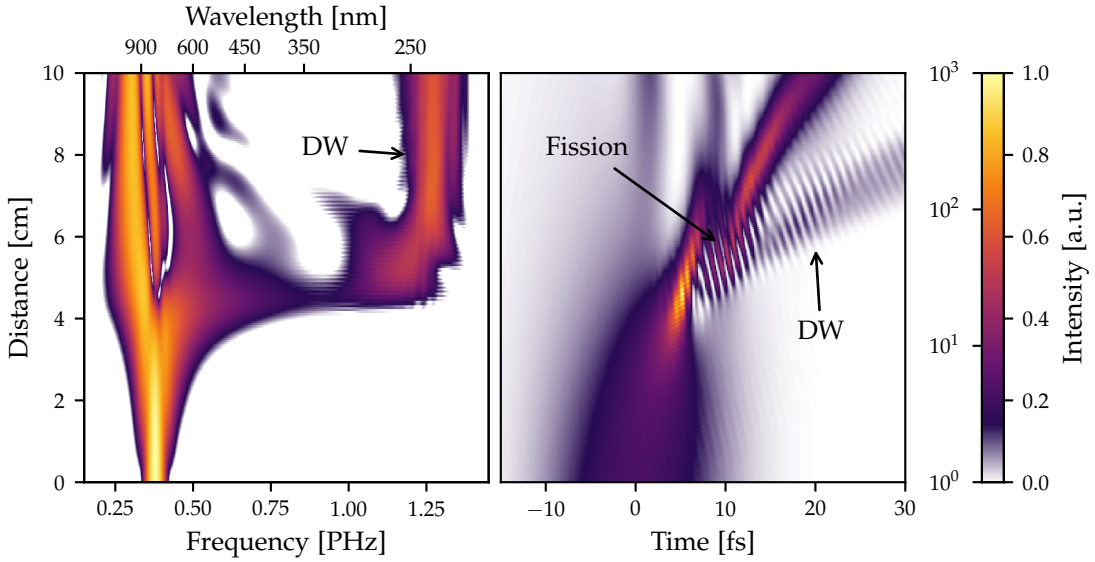

Figure 4.1.: Evolution in the frequency (left) and time (right) domains of a pulse undergoing self-compression leading to the emission of a dispersive wave (DW). In this example, a 15 fs pulse at 790 nm with 1.25 µJ energy is launched into a 10 cm long fibre with a core diameter of 33 µm filled with argon to a pressure of 5 bar. The colour scale is logarithmic in the frequency-domain plot and linear in the time-domain plot.

4.2.2. General features of dispersive wave emission

A particular advantage of numerically modelling the DUV generation process is that it offers a view of the whole process along the waveguide, rather than just the input and the output. A typical example of self-compression and dispersive wave emission is shown in fig. 4.1. Initially, the pulse evolves in a very similar manner to that of a higher-order soliton as shown in fig. 3.5, in that it is compressed in time and correspondingly its bandwidth expands. The moment the frequency spectrum extends to the DUV, however, both higher-order dispersion and the resulting phase-matching to the dispersive wave become important. The dispersive wave appears around 238 nm, and simultaneously the main pulse begins to break up in soliton fission. The point at which this occurs, around 5 cm along the fibre, agrees very well with the fission

length as calculated using the GNLSE model for these parameters, which is 5.1 cm. The wavelength at which DWE occurs, however, is not accurately predicted, with the simple model yielding a phase-matched wavelength of 277 nm.

The time-domain plot shows another important feature of the dispersive wave, which is that the group velocity of the DUV pulse is different to that of the soliton. This can be seen by the weaker pulse appearing around 5 cm moving off towards later times as it propagates, meaning it falls behind the main pulse. The same is true of the soliton that remains after fission, however its group velocity is closer to that of the initial pulse.

As a consequence of the temporal walk-off between the DUV pulse and the soliton, the spectrum of the dispersive wave does not change appreciably after about 6 cm of propagation; the DUV pulse itself is not intense enough to cause a significant nonlinear response. The pulse does accumulate phase, however, specifically it is stretched by dispersion. At the point of soliton fission and DWE, the DUV pulse has a duration of only 4.4 fs. At the end of the fibre, this has increased to 16.4 fs. To preserve a short pulse, the fibre can be cut to the correct length such that DWE occurs near the end of the fibre.

The tunability of the dispersive wave is shown in fig. 4.2. While the DWE wavelength is shorter than the prediction from the analytical model at all pressures, the trend of higher pressures leading to longer wavelengths for the dispersive wave is clearly evident.

4.2.3. Pressure gradients

For the delivery of a pulse created by DWE in a fibre to an experiment without additional dispersion, it is necessary to keep the exit of the fibre under vacuum (see section 5.2). As a result, the gas pressure along the fibre is not constant. The pressure distribution in the fibre can be derived from the condition that the flow rate be the same everywhere, yielding [124]

$$p(z) = p_0 \sqrt{1 - \frac{z}{L}}, \tag{4.2.13}$$

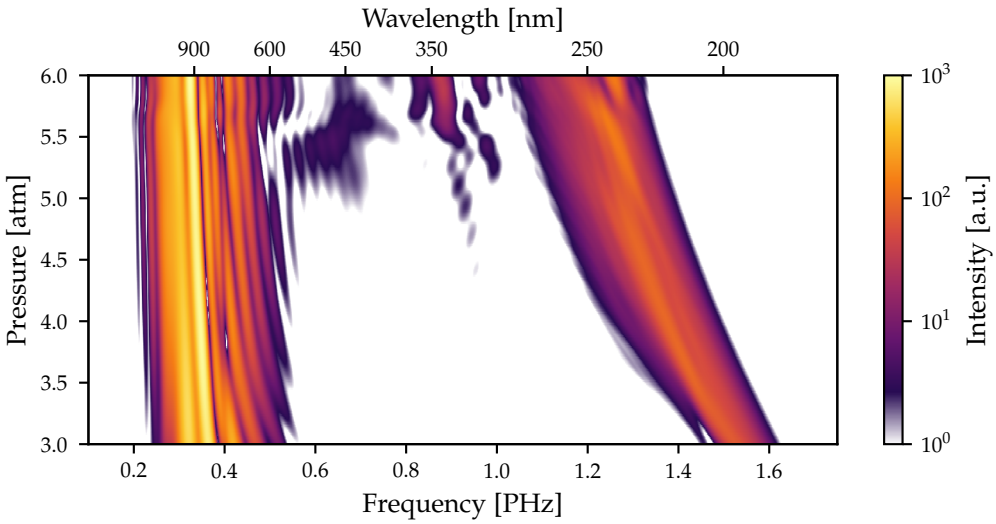

Figure 4.2.: Output spectra created by launching the same pulse as shown in fig. 4.1 into a 15 cm long, 33 µm diameter fibre filled with argon to different pressures between 3 bar and 6 bar.

where z is the distance along the fibre, L is the length of the fibre, p_0 is the pressure at the fibre entrance, and we have assumed that the pressure at the exit is negligibly small. The shape of this pressure distribution is shown fig. 4.3a. The pressure is a critical parameter influencing the dispersion of the gas-filled waveguide as well as the nonlinearity, and thus the soliton dynamics and the phase-matching conditions for DWE change as the pulse propagates. The total dispersion at the soliton wavelength is negative, so the dispersion length decreases as the pressure drops; the dispersion is becoming more negative and $|\beta_2|$ increases. The effect on the nonlinear length is the opposite, as the nonlinear coefficient γ is proportional to the pressure through the nonlinear susceptibility $\chi^{(3)}$. As a result, the soliton number is reduced as the pulse propagates (see fig. 4.3b).

A lower pressure also leads to a shorter phase-matching wavelength for the dispersive wave, in line with the results from a static pressure (fig. 4.2). The effect is slightly reduced by the fact that the nonlinear phase-shift is also reduced, which adds a red-shift

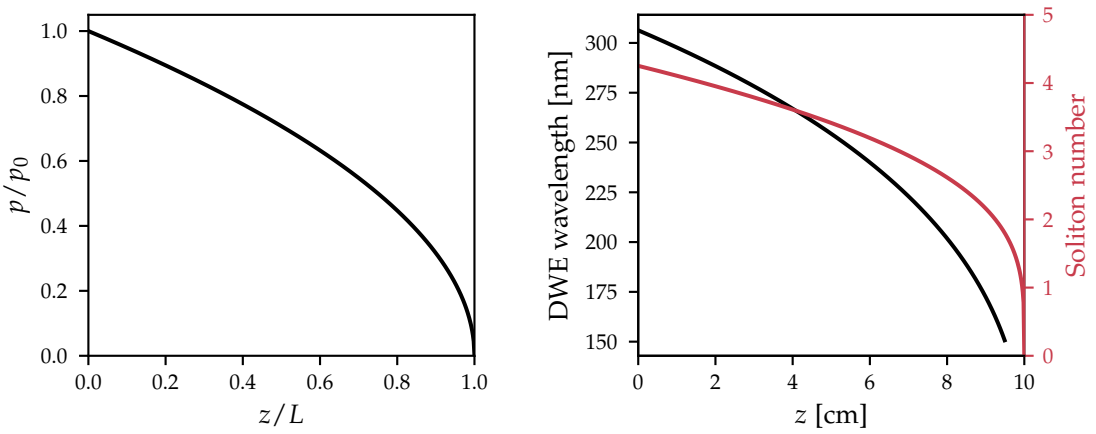

Figure 4.3.: (a) Pressure distribution along the fibre in the case of a negative pressure gradient, normalised to the pressure at the fibre entrance.
(b) Effect of the pressure gradient on the wavelength of the dispersive wave (left, black axis) and the soliton number (right, red axis) for the same pulse and fibre used in figs. 4.1 and 4.2 with the entrance of the fibre at 6 bar.

to the dispersive wave. This variation in the phase-matching condition means that the wavelength of the dispersive wave now depends on where in the fibre it is generated. If the point of maximum self-compression is close to the entrance of the fibre, the pressure is high and the dispersive wave appears at a longer wavelength. Conversely, slower self-compression leads to a shorter UV wavelength.

However, the changing phase-matching condition leads to a more stable situation when a variation of the driving pulse energy is considered. As shown in fig. 4.4, the blue-shifting effect of the higher nonlinear phase shift for larger input pulse energies is reduced somewhat by the concurrent increase in pressure at the point of DWE. The input energy therefore has less of an effect on the wavelength of the dispersive wave [10].

In terms of the overall features, such as the energy at which DWE first occurs as well

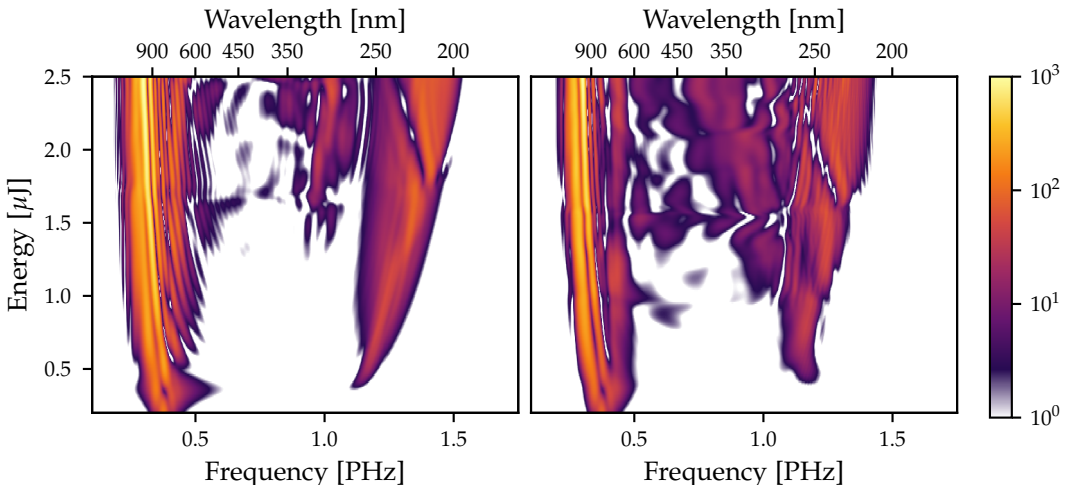

Figure 4.4.: Output spectra for different input energies in the same fibre as in fig. 4.2, filled with a static pressure of 5 bar (left) and a negative pressure gradient with an initial pressure of 7.5 bar (right).

as its general spectral location, the two situations shown in fig. 4.4 are quite similar. The relationship between the two gas pressures for which this agreement is observed is not arbitrary. Consider the total nonlinear phase accumulated along the fibre,

$$\phi_{\text{NL}} = \int_0^L P_0 \, \gamma(z) \, \mathrm{d}z \,. \tag{4.2.14}$$

With γ proportional to the nonlinear refractive index n_2 and thus the third-order susceptibility $\chi^{(3)}$, which in turn is proportional to the pressure, the integrand effectively has the same form as the pressure in eq. (4.2.13).

Demanding that the total nonlinear phase shift be equal for the case of a pressure gradient and a static pressure then reduces to finding the average value of $\sqrt{z}$ in the interval from $z = 0$ to $z = 1$, which is $2/3$. Therefore, by increasing the pressure by a factor of $3/2$ as compared to a static fill, the same general features should be obtained with a pressure gradient, as is indeed the case. The equivalence is only approximate, since the variable dispersion also has an effect on the evolution of the pulse, and is not included in this consideration. However, it can serve as a guide as to how much the fill pressure has to be increased in order to obtain, for instance, a dispersive wave at

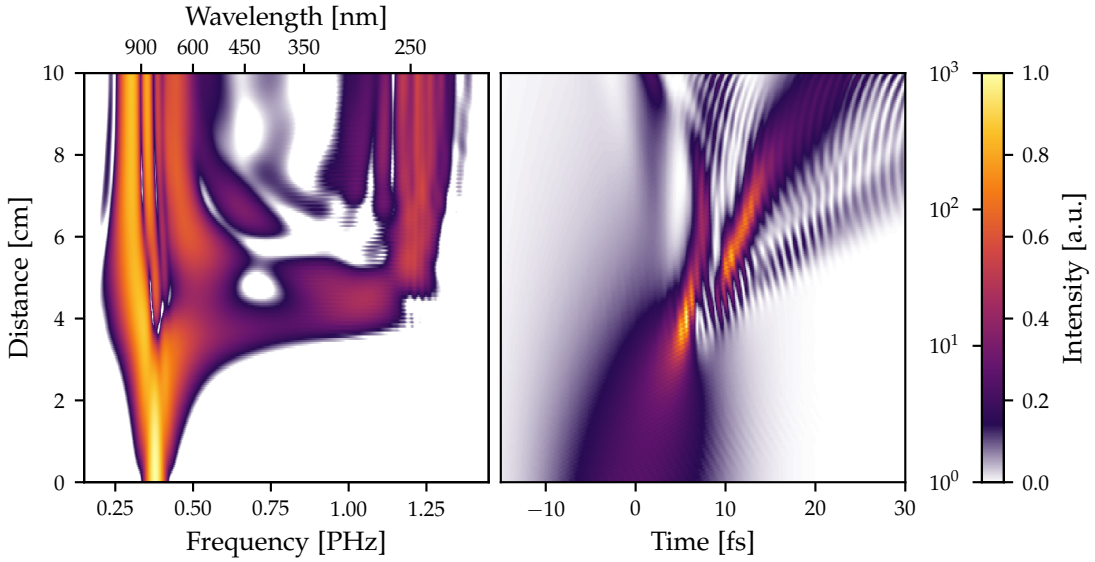

Figure 4.5.: Evolution of the same pulse as in fig. 4.1 but launched into a fibre that is filled with a negative pressure gradient from 7.5 bar to vacuum.

approximately the same wavelength.

The dynamics inside the fibre are also very similar if this condition is fulfilled. Figure 4.5 shows the evolution of the same pulse as in fig. 4.1, but this time launched into a fibre filled with a negative pressure gradient. Most of the overall features remain, including the spectral location of the dispersive wave and the distance after which soliton fission occurs. The spectrum has more structure, however. This is likely due to the appearance of other dispersive waves emitted by the soliton remaining after soliton fission. The presence of multiple pulses in the same spectral region then leads to interference fringes in the frequency domain.

One important consequence of the pressure gradient is that the DUV pulse stretches significantly less after it is created. In the example shown here, the initial pulse is 3.8 fs long and only stretches to 6.5 fs by the time it reaches the end of the fibre. It is thus much easier to retain a short pulse when using a pressure gradient.

The correspondence between a static fill and a pressure gradient is only this clear if the point of soliton fission is relatively close to the beginning of the fibre. With this

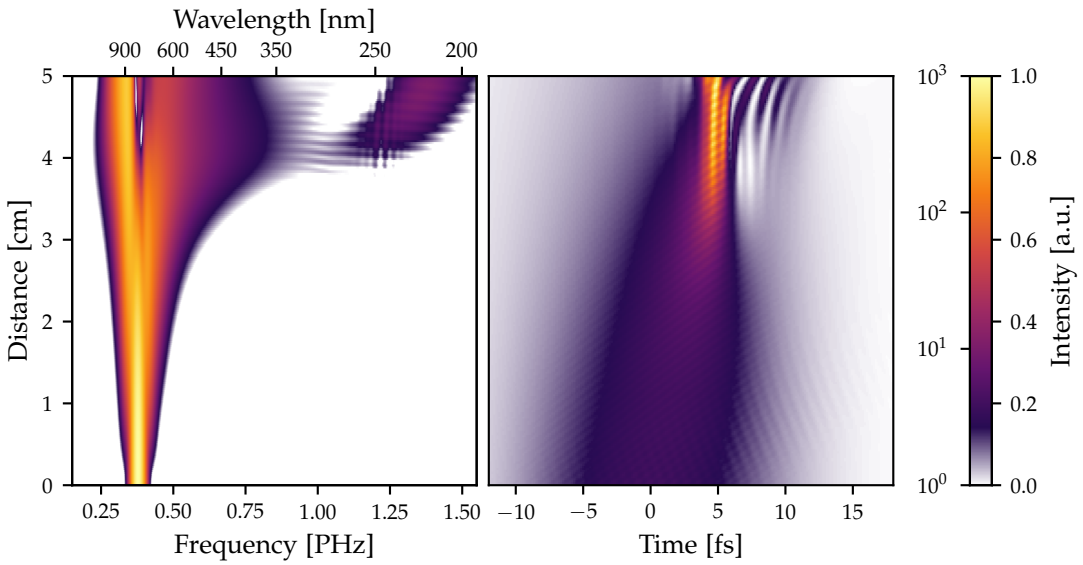

Figure 4.6.: Same as fig. 4.5 but for a fibre of half the length.

point near the fibre exit, the pressure is very different between the two cases, and the phase-matching conditions are altered significantly. As in the case of a static fill, it is beneficial to choose parameters that place the point of DUV generation close to the exit of the fibre, so the influence of the pressure gradient has to be taken into account carefully.

For a pressure gradient, the length of the fibre has more of an influence than for a static fill. The shorter the fibre, the faster the change in pressure, and the greater the change in the nonlinear and dispersion lengths as the pulse propagates. The effect of this is shown in fig. 4.6. In contrast to the longer fibre, the overall evolution is much cleaner, with less structure in both the infrared and the ultraviolet parts of the output spectrum. Correspondingly, the output pulses are also of significantly higher quality. Since the point of maximum self-compression is at a very low pressure and consequently weak higher-order dispersion, soliton fission does not occur. The infrared part of the spectrum forms a single pulse of sub-2 fs duration, albeit with a large pedestal. This situation, where the self-compressed driving pulse remains intact, forms the basis of the two-pulse source proposed in chapter 6.

The lower pressure at the point where the DUV pulse first appears results in a shorter wavelength – 4 cm into the fibre, the pressure has dropped to 3.5 bar, markedly lower than the 5 bar used in the static fill example. The rate at which the pressure drops is also much faster near the end of the fibre than in the middle (see fig. 4.3a) and as a consequence the phase-matched wavelength in the DUV changes while the dispersive wave is being generated. This is visible in the spectral shift as the DUV pulse propagates, with its central wavelength changing by over 20 nm from around 240 nm to below 220 nm. The shifting phase-matching conditions furthermore increase the bandwidth of the dispersive wave: in the example shown it exits the fibre with a duration of 2.6 fs, shorter than even the fully compressed DUV pulse in the example of a static fill.

As with all numerical modelling of complex systems or dynamics, it is very difficult to achieve the exact same results in an experiment as on a computer. Beside the approximations inherent to the propagation equation, there are also uncertainties in the material characteristics used in the simulation. Chief among these are the refractive index of the filling gas [125] and the nonlinear susceptibility [121], the values of which play critical roles in the model. The use of approximate ionisation rates is another limitation, albeit it an unavoidable one. This is unlikely to impact the results significantly, since plasma effects play only a minor role for the parameters considered here; the ionisation fraction does not exceed 10^{-4} in any of simulations shown.

A more important factor is the shape of the input pulse. In all simulations shown so far, the input pulse was taken to have an ideal Gaussian shape. This makes definition of the pulse duration and changing the input pulse characteristics straightforward. To test what influence a more realistic pulse profile has, fig. 4.7 shows the comparison between using an idealised pulse and the input pulse to the fibre used in the experiment (for details of the pulse measurement, see section 4.3.3). Because the measured pulse is somewhat longer and has some post-pulse structure, the total pulse energy has to be increased to achieve DUV generation. Beside this, however, the pulse evolves in a very similar way to the idealised situation. It stands to reason that this would not be the

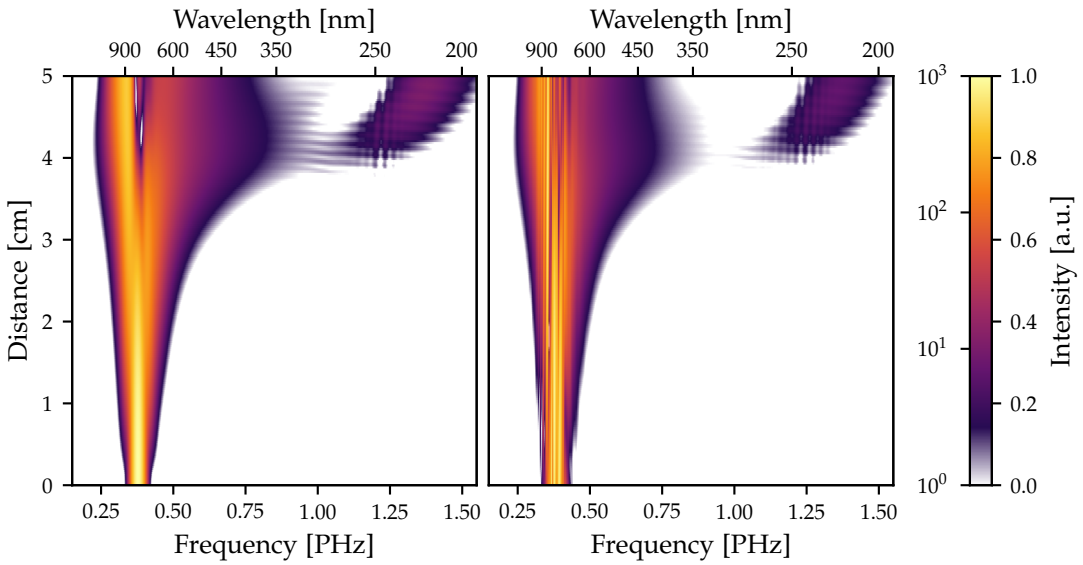

Figure 4.7.: Pulse evolution for the same fibre and gas fill as in fig. 4.6, with an ideal Gaussian pulse of 15 fs duration and 1.25 µJ energy as the input (left) and the measured driving pulse scaled to 1.9 µJ energy (right).

case for very complicated or strongly chirped input pulses, however neither of those cases are relevant for the work in this thesis.

4.3. Experimental implementation

The experimental apparatus has several tasks, with the first part mostly related to shaping the laser beam and pulse to optimise the DUV generation conditions, as well as ensuring long-term stability of the source and protecting the anti-resonant fibre from laser damage. After this, the pump pulse is coupled into the PCF and the output analysed. Here I will describe in detail how each part of the apparatus fulfils its function.

4.3.1. Pre-compression

As mentioned in section 2.3, the majority of the laser output, around 8 mJ, is used for the source of attosecond SXR pulses, with only a small amount of energy available to drive the DUV generation process. Transmitting a short high-energy pulse through the substrate of a beamsplitter would lead to self-focusing of the beam, so instead a dielectric mirror with a reflectivity of 95 % (Eksma Optics) is used to split the beam, with the transmitted beam forming the driving pulse for subsequent DUV generation.

Some energy is split off of this beam for laser diagnostics using Fresnel reflection at near-normal incidence on a wedged piece of Fused Silica (Thorlabs). As a result, the driving pulse has passed through 8 mm of glass (the substrate of the 95 % beamsplitter is 3 mm thick and the wedged splitter 5 mm) in addition to ca. 4 m of air path before it reaches the beginning of the DUV generation apparatus. This means that it is not perfectly compressed. The dispersion present in the pulse cannot be compensated using the grating compressor in the laser system since the tunable optical parametric amplifier is also very sensitive to the spectral phase of its input pulse. Given that both generation schemes are required simultaneously for a pump-probe experiment, this fixes the compressor setting.

As shown in section 3.4, a shorter pulse is beneficial for the DWE process, leading to higher conversion efficiency and a cleaner spectrum in the dispersive wave. Therefore it is useful to compress the laser pulse beyond its initial FTL before using it to drive

DWE.

The pulse compression proceeds in several stages. First, the GDD already present in the pulse is removed by reflecting the beam off of 6 dispersion-compensating mirrors, which induce $-45\,\mathrm{fs}^2$ of GDD per reflection. To expand the spectrum and enable further compression, the beam is then coupled into a 1 m long hollow capillary of 250 µm inner diameter using an $f = 1\,\mathrm{m}$ mirror. The capillary is filled with argon using a positive pressure gradient: the exit is held at 250 mbar, while the entrance is evacuated. To reduce the effect of vibrations caused by the vacuum pump, the last 30 cm of the connection to the gas cell holding the capillary entrance consist of flexible tubing with an outer diameter of 6 mm. Nevertheless, the pumping speed is sufficient (and the flow rate through the capillary low enough) to keep the entrance at a pressure below 1 mbar. This removes self-focusing and ionisation-induced defocusing at the entrance and thus improves the coupling efficiency and mode quality [126].

The GVD at 790 nm of such a large-core capillary is only $-8.2\,\mathrm{fs}^2\,\mathrm{m}^{-1}$ when evacuated and $-3.4\,\mathrm{fs}^2\,\mathrm{m}^{-1}$ when filled with 250 mbar of argon. The dispersion length for a 30 fs pulse is therefore more than 35 m even in the evacuated capillary. The nonlinear length must be much shorter than this for significant spectral broadening to take place over the capillary length of 1 m, so self-compression does not occur; the capillary forms the first half of a traditional hollow-fibre pulse compressor, the second half consisting of dispersive mirrors.

To increase the stability and repeatability of the compression, the gas cells holding the capillary itself are fixed to the optical table, so that it forms a permanent reference. Alignment is carried out using the last two mirrors before the entrance window, which are housed in motorised mounts (Newport). After the second one of these, Fresnel reflection from a 1 mm thick piece of calcium fluoride (CaF_2) at an angle of 45° picks off a small part of the pulse. This weak beam is split further by reflection from another piece of CaF_2. The more energetic transmitted part is used to re-image the plane of the spherical mirror onto the chip of a camera (Logitech) with a magnification of 0.1 using an $f = 100\,\mathrm{mm}$ lens, while another camera is placed in the focal plane of the weaker

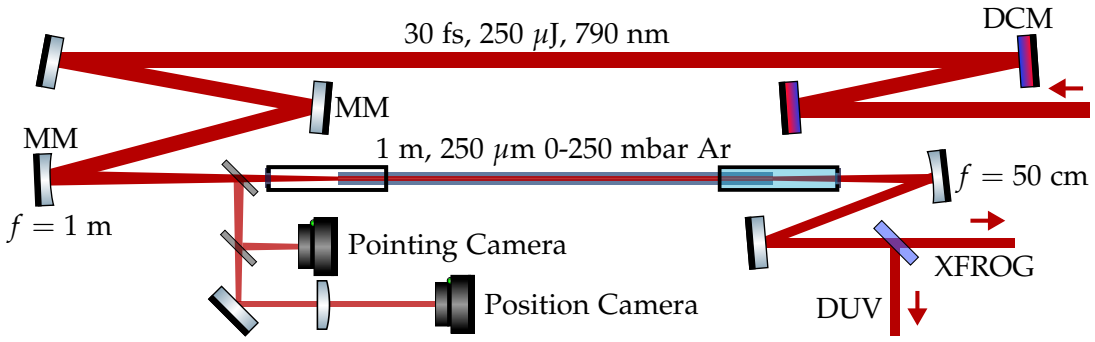

Figure 4.8.: Schematic of the hollow capillary apparatus for pre-compression of the driving pulse. The "position" lens and camera image the beam in the plane of the $f = 1\,\mathrm{m}$ in-coupling mirror. The "pointing" camera is situated in the focal plane of the in-coupling mirror and thus measures the pointing of the beam. This is information is then used to stabilise the alignment into the capillary with the two motorised mirrors (MM). After recollimation by an $f = 50\,\mathrm{cm}$ mirror, the beam is split, with part of the energy used for the DUV generation and part for the characterisation (see chapter 5). DCM: dispersion-compensating mirror

reflected part.

Together, the camera images completely identify the position and pointing of the laser beam as it impinges on the entrance face of the capillary. This allows for active stabilisation. Once the input to the waveguide is aligned the first time, the corresponding camera images are saved. By comparing subsequent images to these references, an error signal is calculated, which is used as the input to a proportional-integral-differential (PID) algorithm controlling the movement of the steering mirrors. The comparison is done by calculating the cross-correlation between the references and later images and finding its peak location, a method which is robust to noise and small changes in the beam profile [127]. I have written software which carries out the data acquisition as well as the PID control in real time (the typical frame rate is 8 Hz) and includes a graphical user interface for ease of use, based on an earlier version created by

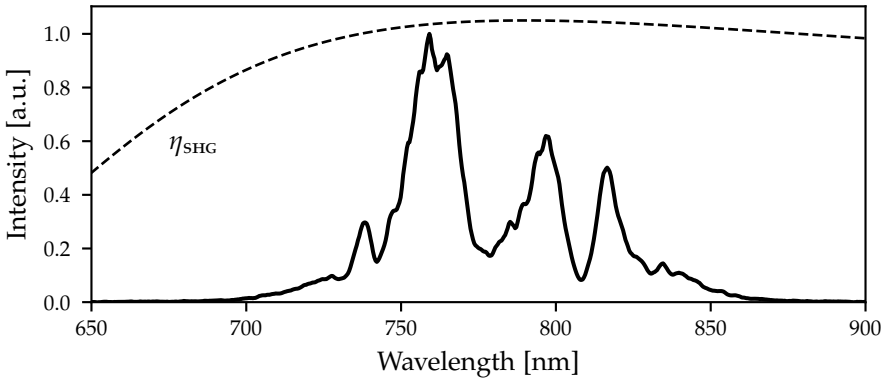

Figure 4.9.: Spectrum of the laser pulse after spectral expansion in the hollow capillary. A 3 pixel wide boxcar smoothing filter has been applied to the data to reduce random noise from the spectrometer and make the overall shape of the spectrum more obvious. The dashed line shows the type I SHG phase-matching efficiency η_{SHG} on an arbitrary scale for a 10 µm BBO crystal at $\theta = 29.55°$ (see fig. 2.6), as used in the pump pulse measurement (see section 4.3.3).

Dane Austin. The active stabilisation can correct for slow drifts caused by temperature fluctuations or air currents automatically, allowing for long-term use of the apparatus without manual intervention.

The beam at the output of the capillary is re-collimated with an $f = 500$ mm spherical mirror. Its spectrum is shown in fig. 4.9 and yields an FTL pulse duration of 11.2 fs. Some of the pulse exiting the capillary is split off by a 65 % beamsplitter, with the majority of the energy being used for the DUV generation, and the smaller part as the second pulse in the DUV pulse characterisation apparatus (see chapter 5). The total pulse energy after the capillary is around 160 µJ. This corresponds to a transmission through the waveguide of ca. 60 %.

For the DUV generation, the positive GDD induced by the SPM and the exit window of the gas cell is compensated by reflection off of further chirped mirrors. In a conventional hollow-fibre compression system, the mirror set usually over-corrects

(adds more negative GDD than necessary for compression) and a pair of glass wedges is used to fine-tune the spectral phase of the pulse. In the present apparatus, the pulse has to pass through several transmissive optics before it is coupled into the PCF for DUV generation (see section 4.3.2). This additional dispersion has to be compensated, and the beam therefore reflects off of the chirped mirrors 18 times, with each reflection adding ca. $-45\,\mathrm{fs}^2$ of GDD. The reflectivity of the dispersive mirrors is ca. 97 %; this reduces the pulse energy from around 100 µJ to 60 µJ.

4.3.2. Spatial filtering and coupling

While the mode-dependent loss of the capillary (see section 3.1) lets it act like a spatial filter by itself, the beam profile at the exit is not perfect. This is partly due to residual higher-order modes and the imperfect straightness of the capillary itself as well as astigmatism induced by the spherical re-collimating mirror. Most importantly, however, the surface quality of dispersion-compensating mirrors is poor when compared to, for instance, steering mirrors, and consequently the beam profile after compression is afflicted by small-scale distortions (see fig. 4.10). To improve the spatial profile of the beam, it passes through a pinhole-based spatial filter.

The principle of spatial filtering is based on the Fourier transform property of focusing optics: in the paraxial approximation, the intensity profile in the focal plane of a beam is equal to that of its two-dimensional transverse Fourier transform with the coordinates scaled by the focal length and wavelength [72]. Therefore, small-scale spatial distortions, which correspond to high transverse spatial frequencies, will appear in the focal plane as components far away from the centre of the focal spot.

By placing a hard aperture in the focal plane, these components can be eliminated, and the beam in the far-field will emerge without the distortions. As the Fourier transform relationship holds, this is equivalent to blurring the far-field beam profile by a two-dimensional convolution with the (appropriately scaled) Fourier transform of the aperture function[1]. The amount of filtering is set by the ratio of the focal spot size to

[1] In the case of a circular pinhole, this blurring function is the Airy pattern. [72]

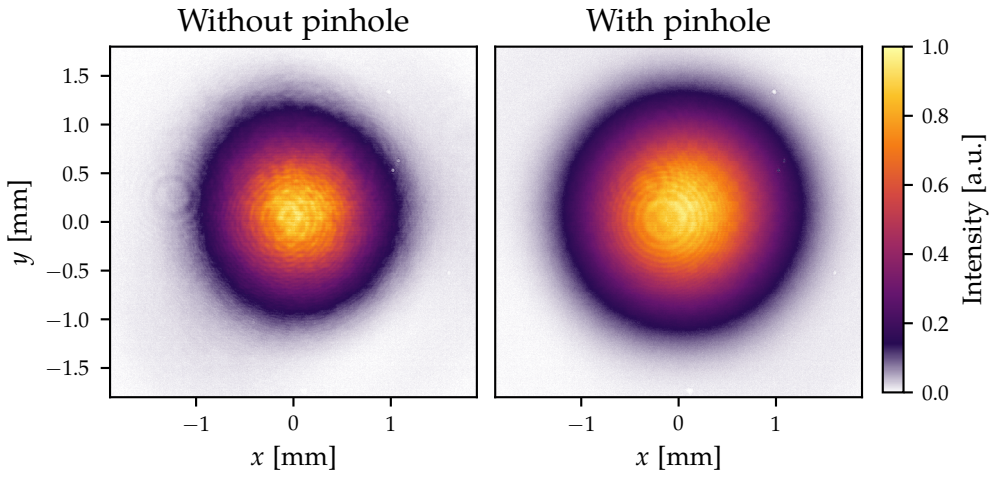

Figure 4.10.: Transverse profile of the beam in a plane ca. 8 cm downstream of the spatial filter pinhole, with the pinhole removed (left) and inserted (right). As well as showing a cleaner overall profile, the beam is larger since the pinhole crops the focal spot to a smaller size. The circular diffraction patterns near the centre of the beam are caused by dust on the attenuators before the detector.

the pinhole size. Since the pinhole is removing energy from the pulse, there is also a trade-off between transmission through the filter and the quality of the resulting beam.

If the spatial filter is implemented using only transmissive optics, it confers another advantage. Contrary to reflective optics (i.e., mirrors), any movement of a transmissive optic caused by air currents or accidental impact is not magnified by the subsequent propagation distance – a small tilt of a mirror may move a laser beam by many times the amount the mirror moved if it propagates a large distance. By building the spatial filter completely in-line using lenses, and afterwards re-imaging the pinhole directly onto the entrance of the PCF, unintended movement occurs only before the pinhole, rather than between the spatial filter and the PCF. Misalignment due to any movement will not cause damage to the delicate cladding structure of the anti-resonant waveguide as the misaligned beam will be blocked by the pinhole.

To harness this advantage, the spatial filter is located just before the beam is coupled

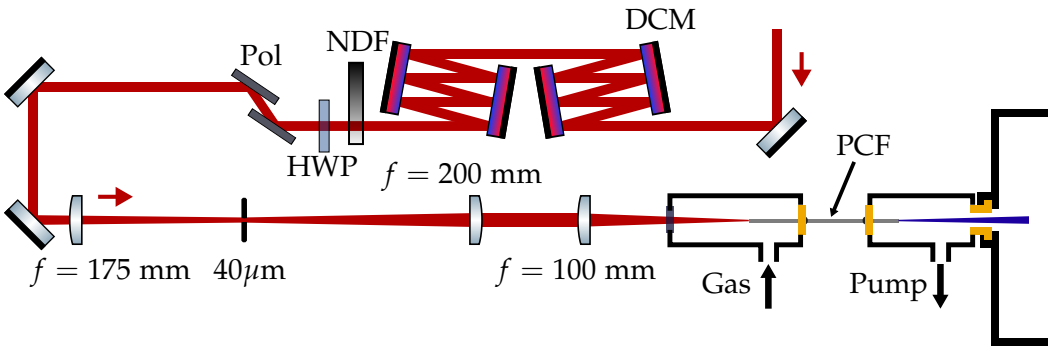

Figure 4.11.: Schematic of the DUV generation apparatus. DCM: dispersion-compensating mirror; NDF: neutral density filter; HWP: half-wave plate; Pol: polariser.

into the PCF. After the dispersive mirrors, the beam traverses a tunable attenuator formed of a variable neutral density filter wheel (Thorlabs) as well as a half-wave plate (B.Halle) and thin-film polariser (Eksma) combination. The half-wave plate is mounted in a motorised rotation stage (Thorlabs) which allows precise and repeatable scanning of the pulse energy. To reduce the total amount of GDD induced, the polariser is used in reflection, and as a consequence the beam is vertically polarised after the attenuator. The filter wheel functions as added attenuation during alignment; the extinction ratio of the polariser is insufficient to reduce the initial pulse energy of over 50 µJ to a level that is safe for alignment into the PCF.

The three-dimensional translation stages holding the spatial filter pinhole, fibre-coupling lens and input gas cell (Elliot Scientific) are mounted such that their optical axes are parallel by using a straight edge. The beam is then aligned onto this common optical axis using the steering mirrors before the first lens by mounting apertures on the stages. To enable the use of pressure gradients, the PCF is glued into both gas cells.

The spatial filter is formed of an $f = 175$ mm plano-convex focusing lens (Thorlabs), a diamond pinhole of 40 µm diameter (Fort Wayne Wire Dies) in its focal plane, and an $f = 200$ mm plano-convex re-collimating lens (Thorlabs). The beam is then re-focused to couple into the PCF by an $f = 100$ mm plano-convex lens (Thorlabs). An image of

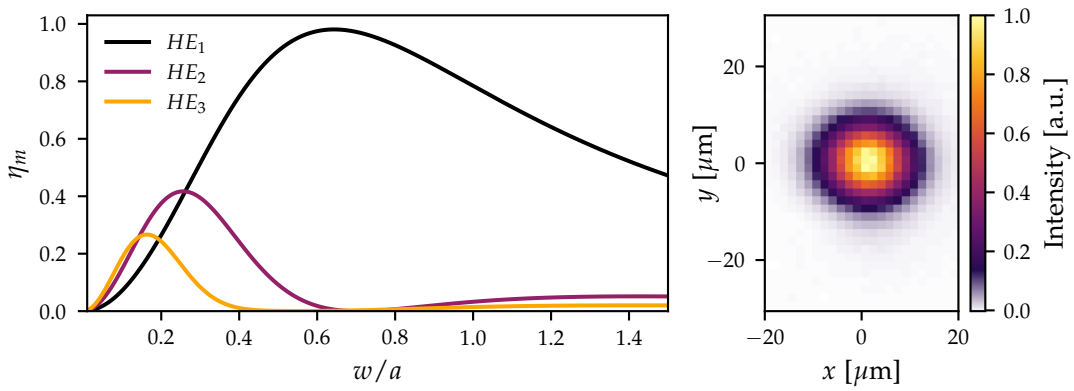

Figure 4.12.: Left: Coupling efficiency η_m for the three lowest-order modes in a capillary as a function of the ratio between input spot size w and the core radius a. The peak efficiency for the fundamental mode is 98 % and occurs at $w/a = 0.64$.

Right: Image of the focus in the plane of the entrance face of the PCF.

the focus is shown in fig. 4.12.

In principle, it is possible to combine the latter two lenses and directly re-image the pinhole plane onto the PCF entrance face, reducing the amount of dispersion to be compensated. However, this has the disadvantage that the focal spot size and the distance between the pinhole and the PCF are coupled. The particular arrangement necessary for direct coupling to vacuum (see section 5.2) means that the fibre itself cannot move once it is installed. Thus, whenever the entrance face of the fibre is moved (for instance by cleaving), the entire spatial filter system would have to be moved as well. With a collimated beam between the two lenses, this is not the case.

Using two lenses has the further advantage of providing two degrees of freedom for changing the focal spot size; this is critical when coupling the beam into the fibre. For a Gaussian beam impinging on the fibre in perfect alignment (that is, without any tilt and with the beam waist at the fibre entrance), the field profile at the entrance face is given by

$$E(r,\theta) = E_0\, \mathrm{e}^{-(r/w)^2}, \tag{4.3.1}$$

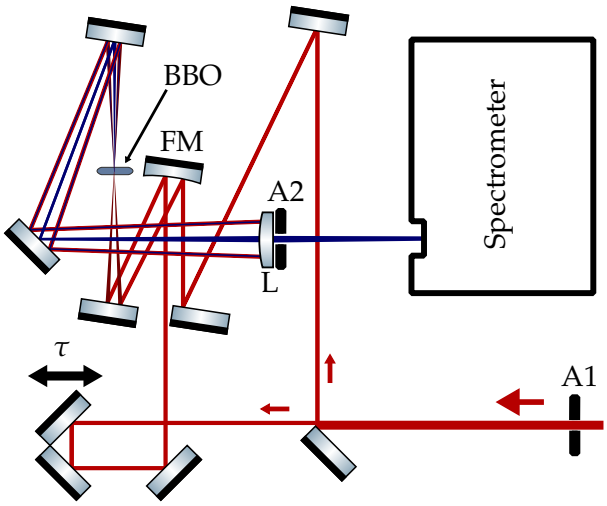

Figure 4.13.: Optical layout of the SHG FROG used to measure pulses at 790 nm. A1: entrance aperture; τ: delay stage; FM: focusing mirror; L: re-imaging lens; A2: aperture to remove the fundamentals and their second harmonics.

where w is the $1/e^2$ radius of the beam waist. The coupling efficiency for a waveguide following the capillary model according to eq. (3.1.9) then reduces to

$$\eta_m = \frac{4\left[\int_0^a J_0(u_m r/a)\, e^{-(r/w)^2}\, r dr\right]^2}{w^2 \int_0^a J_0^2(u_m r/a)\, r dr}. \tag{4.3.2}$$

This depends critically on the ratio between the core radius a and the beam waist radius w, as shown in fig. 4.12. It is therefore very important to match the spot size of the incoming beam as closely as possible to the optimal value. Note that while this is also very important for the larger hollow capillary, any higher-order modes present at the entrance are attenuated in that case, leading to nearly single-mode output. This does not happen in anti-resonant fibre, so that more care must be taken not to excite any higher-order modes.

4.3.3. Pump pulse measurement

In order to optimise the pulse compression and to verify that a compressed pulse is used to drive the DUV generation process, I have constructed a scanning SHG FROG

device for the pump pulse. It is designed to perform the measurement with no phase added prior to the nonlinear interaction, resulting in a faithful representation of the pulse in the FROG trace.

The optical layout of the device is shown in fig. 4.13. The pulse is split with wavefront division by reflection off of a half mirror, with one half of the beam passing the edge and the other being reflected. One replica then passes through a delay-line formed of a retro-reflecting pair of mirrors on a motorised delay stage (Thorlabs). The two arms are recombined on parallel beam paths by respectively passing and reflecting off of another half mirror before being focused into the nonlinear medium[2], a 10 µm thick BBO crystal on a 1 mm thick substrate (Lightkey Technology), by an $f = 150$ mm spherical mirror. The crystal is cut at $\theta = 29.2°$, which is the type I phase-matching angle for 800 nm, however with a small rotation of the crystal the correct angle of 29.55° can be obtained.

Of the resulting three beams, the two fundamentals and their co-propagating second harmonics are removed by a variable aperture (Thorlabs), which is mounted as close as possible to the re-imaging element, an $f = 75$ mm lens (Thorlabs). This is because it is in this plane that the beams are furthest apart and thus easiest to separate. The remaining beam, which contains the FROG signal, is then analysed by a spectrometer (Ocean Optics) with a response window from 190 nm to 1100 nm, more than wide enough to capture the entirety of the trace. Acquisition of a FROG trace consists of scanning the delay using the motorised delay stage and acquiring spectra at each delay.

Note that after the nonlinear interaction has taken place, only the power spectrum of the signal pulse is important for the measurement, so the use of a transmissive (and thus dispersive) imaging element is not an issue, as long as its transmissivity is sufficiently uniform across the bandwidth of the FROG signal.

The crossing angle between the two pulse replicas is approximately 1°. The amount of temporal smearing induced by this depends on the size of the focal spot in the crystal

[2]Spatial overlap was ensured by first placing a camera in the focal plane. The crystal was then placed in the focal plane of the beams by scanning it along the propagation direction and maximising the amount of SHG signal.

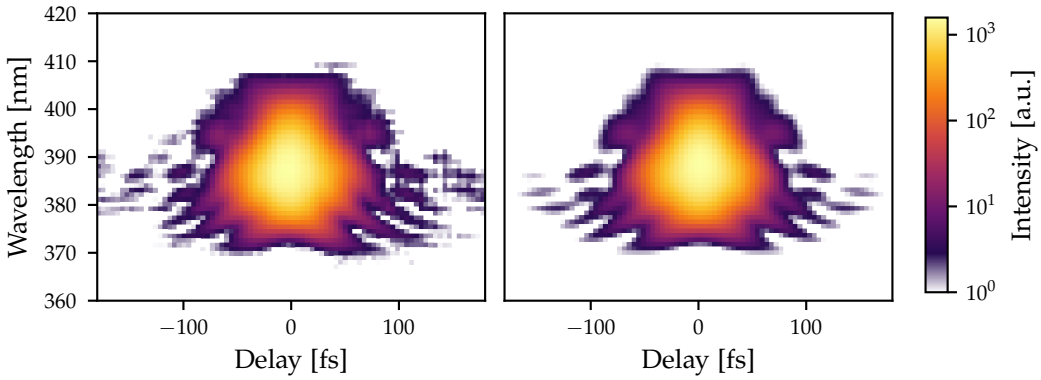

Figure 4.14.: Left: measured FROG trace of the pulse at the entrance of the hollow capillary after thresholding, on a logarithmic colour scale to make low-intensity parts of the signal visible. The trace consists of 115×115 points. At each delay sample, the signal acquired is the average of 100 laser shots (the average of 10 acquisitions of 10 ms duration). Right: reconstructed trace after 200 iterations of the PCGPA.

as well as the magnification onto the detector, with the latter setting the upper bound. Since the input beam size and aperture setting can vary between measurements of the pulse at different points in the apparatus, it is safest to calculate the maximum smearing. The entrance slit of the spectrometer is 25 µm wide and the lens is 20 cm away from the crystal plane. Since the spectrometer is positioned in the image plane of the lens[3], the magnification can be calculated as $M = 0.6$ using the thin lens formula [128]. The spectrometer thus captures light from a spot in the crystal 42 µm across, resulting in a maximum temporal blurring according to eq. (2.4.6) of 2.4 fs. For a 15 fs pulse, this means the measurement will overestimate the pulse duration by around 190 as, which is small enough to be ignored.

Since there are two nonlinear processes involved in the DUV generation – spectral

[3]To establish this, the beam size at the spectrometer slit was measured at several distances from the lens by scanning the whole spectrometer across the beam using a translation stage and recording the signal strength. The imaging plane is then where the spot size is minimal.

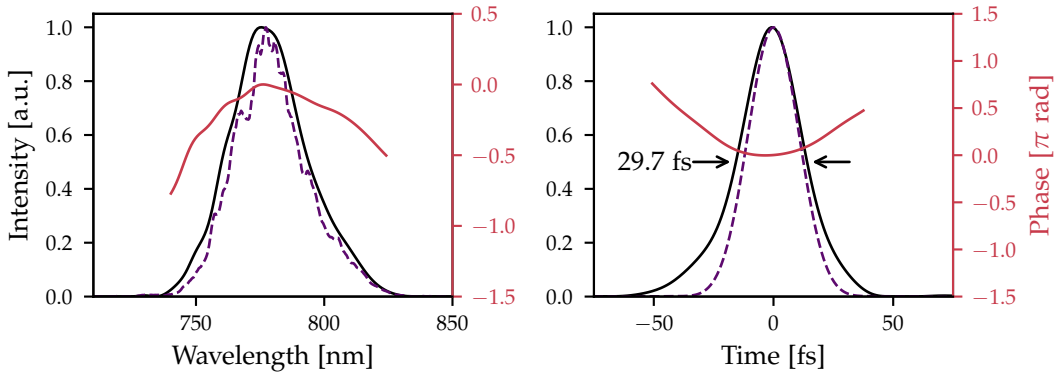

Figure 4.15.: Input pulse at the entrance of the capillary as measured using the SHG FROG apparatus, represented in both the time and frequency domains. Left: the power spectrum (black) and spectral phase (red) of the pulse as reconstructed from the FROG measurement, and the independently measured power spectrum (purple dashed). Right: time-dependent intensity (black) and temporal phase (red) as well as the intensity profile of the transform-limited pulse (purple dashed) using the power spectrum as reconstructed from the FROG measurement.

expansion in the large capillary and DWE in the anti-resonant PCF – there are two points where the pulse should be compressed for optimal conditions; these are at the entrance of the capillary and the PCF, respectively. The pulse is therefore measured at both of these points. To get the most accurate measurement, as much of the dispersion affecting the pulse in practice should be included. Therefore, the pulse is picked off as close as possible to the entrance of the capillary and PCF, respectively.

For the capillary, a mirror is inserted after the piece of CaF_2 serving as a diagnostic pick-off. Since the beam is converging at this point, it is re-collimated after coming to a focus using a concave spherical mirror and subsequently steered into the FROG apparatus. Similarly, a mirror is inserted after the last focusing lens before the PCF entrance and the beam re-collimated. In both cases, the only dispersive element not taken into account is the respective entrance window to the gas cell – a 1 mm thick,

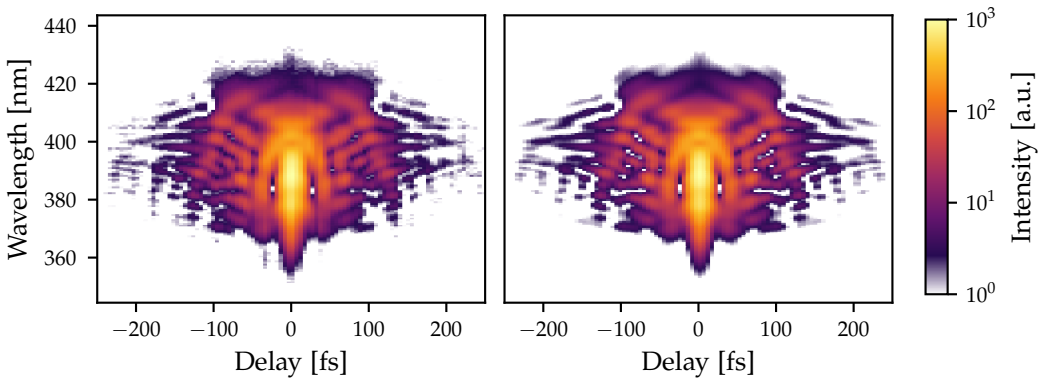

Figure 4.16.: Same as fig. 4.14 but for the pulse at the entrance of the PCF. Here the trace contains 149 × 149 points.

anti-reflection coated piece of fused silica in the case of the capillary, and a 1.5 mm thick uncoated piece of CaF_2 for the PCF. For the measurement of the PCF input pulse, an equivalent piece of glass is inserted into the beam path before the FROG device to alleviate this problem. Note that the beam also propagates through some amount of air both before and inside the FROG device before the nonlinear interaction takes place. This can be taken into account by phase-correcting the measured pulse if required.

Reconstruction of the acquired SHG FROG traces is carried out using the PCGPA. The software for this was implemented by Dane Austin. To increase the reliability of the reconstruction, a threshold is applied to the trace first to reduce noise. Furthermore, the PCGPA requires that the data is interpolated onto a square grid of delays and frequencies which are Fourier-conjugate axes to one another[4]. The measured and reconstructed FROG traces are shown in figs. 4.14 and 4.16. The FROG error Δ is 0.0092 for the measurement at the capillary entrance and 0.0062 for that at the PCF entrance, indicating excellent agreement between the two traces.

Because the two pulses interacting in an SHG FROG measurement are identical, the

[4] A pair of time and frequency axes are Fourier conjugates if their sampling rates exactly match the minimum required for the window size in the other domain according to the Shannon-Nyquist sampling theorem [59].

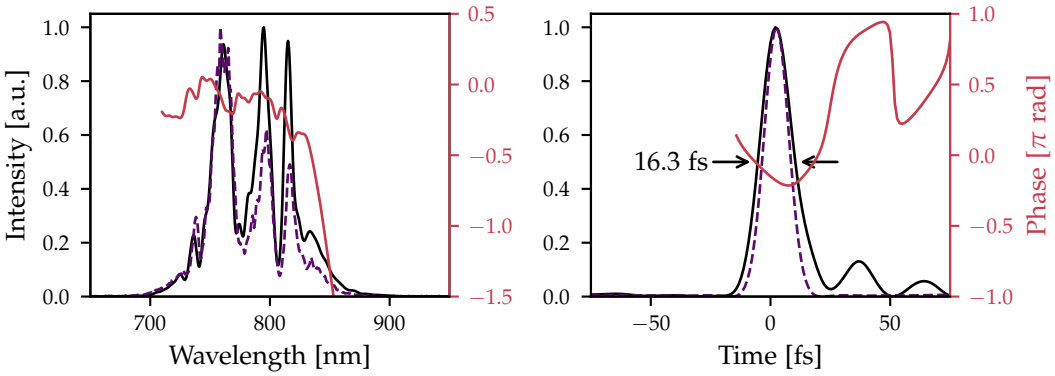

Figure 4.17.: Same as fig. 4.15 but for the pulse at the entrance of the PCF.

trace contains several ambiguities, i.e., questions that cannot be answered without further information. The unknown quantities are the arrival time of the pulse, its absolute phase, and most importantly, the sign of the spectral phase, or equivalently the direction of time [60]. The symmetry of the trace means that it is the same if the delay axis is reversed, which also reverses the direction of time in the reconstructed pulse. A single SHG FROG measurement is therefore incapable of distinguishing between a pulse that is positively chirped and one that is negatively chirped.

To alleviate this problem, the symmetry can be broken by adding a known amount of dispersion to the pulse and taking a second measurement. There is only one combination of signs for the phases from the two measurements for which their difference will match the added phase, the sign of which is known. The sign ambiguity is lifted here by propagating the pulse through a 6 mm thick piece of fused silica, which adds around 220 fs^2 of GDD to the pulse.

The measurement of the pulse at the entrance of the capillary is shown in figs. 4.14 and 4.15, and that of the PCF input pulse in figs. 4.16 and 4.17. The pulse at the entrance of the capillary is slightly negatively chirped since no equivalent piece of glass was inserted for this measurement. By numerically adding propagation through a 1 mm thick piece of silica glass, the measured pulse is compressed by about 1 fs.

As shown in figs. 4.16 and 4.17, the pulse compression in the capillary has reduced

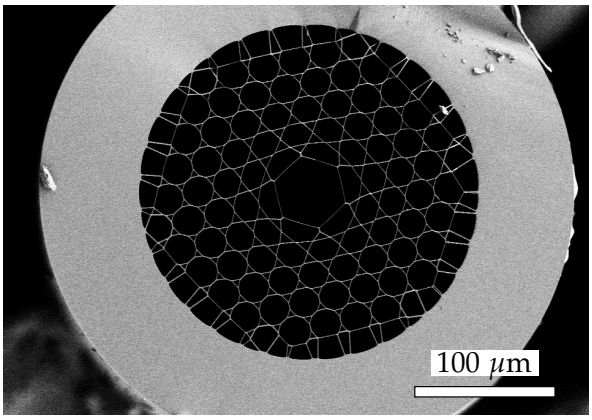

Figure 4.18.: Cross-section scanning electron micrograph of the 46 µm core diameter kagomé PCF used for the results in this section. Image courtesy of Francesco Tani, Max Planck Institute for the Science of Light, Erlangen, Germany.

the pulse duration by a factor of around 1/2. The pulse is also very close to its transform limit in duration, evidencing that a compressed pulse can be delivered to the DUV generation process. The weak post-pulses at around 40 fs and 60 fs are the consequence of the steep phase gradient around 850 nm visible in the frequency domain plot. This is likely caused by the bandwidth extending slightly past the bandwidth of the dispersion-compensating mirrors.

4.3.4. DUV output energy and spectrum

The main quantities of interest for the characterisation of the output of the apparatus are the spectral location and energy of the DUV pulse, as well as its duration. Additionally, the stability of the output is also important for the source to be useful in a transient absorption experiment. The measurement of the duration is the main experimental result of the work for this thesis, and is addressed in detail in chapter 5.

To measure the spectrum of the DUV pulse, the beam exiting the PCF was analysed with a commercial spectrometer (Ocean Optics). Since the beam in the fibre is contained largely in the fundamental mode, the beam size at all wavelengths is the same as the

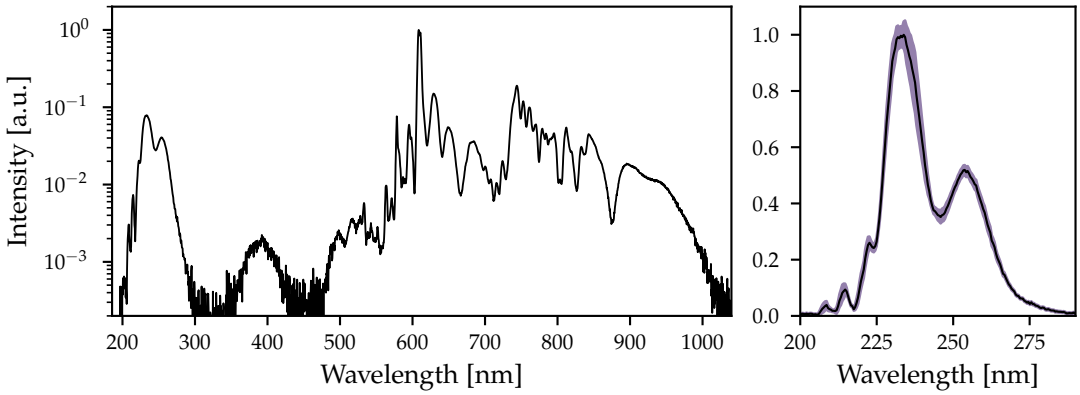

Figure 4.19.: Typical output spectrum of a DWE process driven by a 16 fs pulse of 8 µJ energy in a 46 µm core diameter fibre filled with 23.5 bar of neon. Left: The total output spectrum on a logarithmic scale. Right: the DUV part of the spectrum on a linear scale. The shaded area shows the standard deviation at each wavelength over 1000 shots taken over ca. 5 minutes.

pulse leaves the waveguide. The divergence angle of each wavelength is therefore different, leading to significant spatio-spectral coupling in the far field. This could be taken into account either by focusing the beam onto the entrance slit of the spectrometer, or by using a diffusing sphere. However, since a full spectral sensitivity calibration of the spectrometer was unavailable, a significant source of uncertainty remains about the shape of the overall spectrum even then. To simplify the apparatus, the spectrometer was therefore simply placed in the collimated beam.

A typical output spectrum of a DWE process is shown in fig. 4.19. The dominant features are the extreme spectral broadening around the pump pulse wavelength of 790 nm and the peak at 250 nm, which shows the dispersive wave. The sharp edge of almost two order orders of magnitude in spectral intensity near the peak at 600 nm is due to a cladding resonance in the fibre. The fibre in this case is larger than that used in previous examples (see fig. 4.18), since a larger core diameter allows for more energy to be coupled into the fibre while preserving the balance of nonlinear and dispersive effects required for DWE. As a result of the large core diameter, phase-matching in the

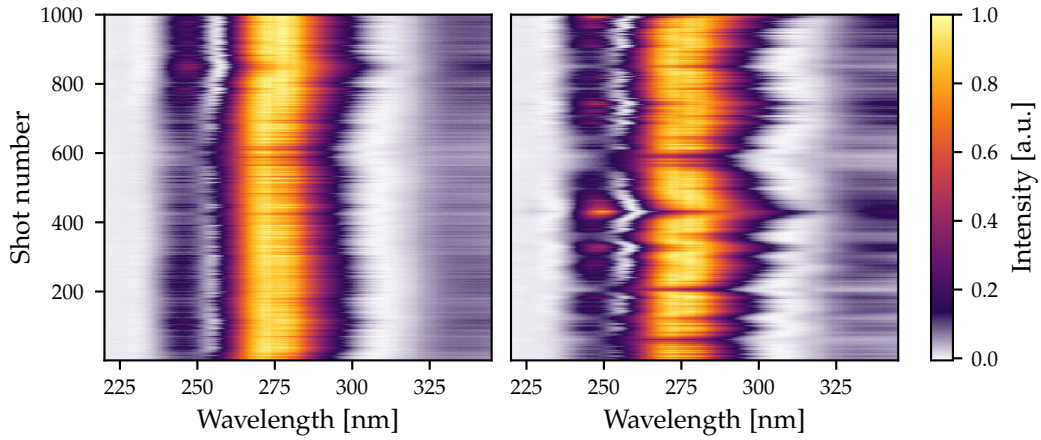

Figure 4.20.: Output spectra from the DUV generation apparatus acquired over ca. 5 minutes with (left) and without (right) measures to increase the stability of the system.

DUV spectral region is best achieved by using high pressures of neon gas.

The spectrum of the dispersive wave is shown in the right-hand part of fig. 4.19. The transform-limited duration of a pulse with this spectrum is below 3 fs. By inserting a prism to disperse the output beam spectrally, the DUV component could be isolated in space and the pulse energy measured with a photodiode power meter (Thorlabs). The total energy in the spectral region between 200 nm and 300 nm in this case was 250 nJ, taking into account losses caused by the prism as well as the window on the exit gas cell. This corresponds to a conversion efficiency of 3.2 %. While higher conversion efficiencies have been achieved in DWE [87], this is likely already enough for some experiments, and the efficiency is sufficient for the generation of µJ-level pulses by appropriate scaling of the input pulse energy and fibre length.

Stability

For the DWE-based source to be useful in a transient absorption experiment, it must be capable of delivering a stable pulse, that is, one that does not change significantly in either duration or spectral location. While also important, the DUV pulse energy is

less critical, since fluctuations in both the pump and probe intensities can and must be addressed by the design of the experiment, given the intrinsically more unstable probe pulse generation mechanism of HHG. As shown in fig. 4.19, under good conditions the DUV part of the spectrum is very stable, with the standard deviation over 1000 shots not exceeding 5 % at any point in the spectrum. The standard deviation of the total pulse energy in the DUV region is 3.5 %.

The most important parameter to achieve this is the stability of the laser system, however the design of the generation apparatus itself is critical as well. There are two aspects to this. The first, as already discussed in section 4.3, is to ensure that any thermal drifts, air currents or accidental movements cannot lead to damage to the PCF. This is crucial especially for experiments using a pressure gradient in the fibre, since the bonding of a new fibre into the gas cells takes more than 12 hours. The second is that pointing instabilities must be kept to a minimum; through the spatial filter and the coupling into the PCF, beam movements cause power fluctuations in the fibre in addition of those already present in the laser beam.

The biggest detrimental factor to overall stability of the apparatus is the presence of air currents in the laboratory. To mitigate this effect, the simplest solution is to enclose the setup. The effect of this is shown in fig. 4.20, where the stability of the DUV spectrum is compared between an enclosed and open apparatus. Two conclusions can be drawn from this observations. The first is that minimisation of air currents is indeed critical to achieving reliable operation of the DUV source. The second is that even in the presence of strong air currents and the attendant instability in the whole apparatus, the PCF is protected by the spatial filter.

5. Characterisation of DUV Pulses

As outlined at the beginning of chapter 4, a source of DUV pulses for transient-absorption experiments must fulfil three central criteria if it is to be useful – it needs to be bright, tunable in wavelength across the DUV range and short in duration, significantly below 5 fs if possible. DUV generation by dispersive wave emission can satisfy these in principle. The first two, pulse energy and spectral tunability, are relatively straightforward to measure. As with any ultrashort pulse, however, establishing its duration is substantially more difficult.

When the second pulse in the intended experiment is generated by HHG, it is effectively guaranteed to be very short. The time resolution is then chiefly determined by the DUV pulse, and measuring it accurately is critical for the experiment. In addition, the pulse that is characterised must faithfully reflect the one that is used in the experiment. While an *in situ* measurement would be ideal for this purpose, in many cases this is not possible, and instead large differences between the pulse characterisation and experimental set-ups should be avoided.

An accurate measurement of the pulse duration is especially important for pulses generated by DWE, since the strong coupling between various experimental parameters precludes reliable predictions of the pulse shape from its spectrum alone. This is in contrast to, for instance, bulk-crystal harmonic generation, where the fixed medium provides a degree of confidence that, after initial characterisation, a pulse with the same spectrum will have the same duration.

In this chapter, I describe the design and implementation of a pulse characterisation apparatus, based on FROG, to measure DUV pulses generated in an anti-resonant fibre. The measurement is carried out in near-identical conditions to those of a transient-

absorption experiment, and therefore demonstrates what pulse duration can be delivered to an experimental target. In particular, I address the challenge of preserving the pulse duration of an initially short DUV pulse both for the pulse characterisation and an experiment.

5.1. Cross-correlation FROG

At first glance, the problem of measuring the DUV pulses generated by DWE is very similar to that of measuring any other ultrafast laser pulse, for instance the one driving the process. Therefore it should be possible to fully characterise the pulses with a FROG measurement. However, the particular properties of the unknown pulses in this case complicate the situation.

Firstly, the DUV pulses can be very short, with durations below 5 fs predicted [10, 11]. Although pulses of this duration have been measured using FROG [65], it is significantly more challenging than for longer pulses, and the frequency response of the measurement apparatus as well as the nonlinear process have to be taken into account with great care.

The more immediate issue, however, is that the most common variant, SHG FROG, cannot be used in the DUV. This is because the crystals normally used as the nonlinear medium are not transparent at sufficiently short wavelengths. The lowest transmitted wavelength is 190 nm for BBO [17], which is currently the only nonlinear crystal for this wavelength region that can be manufactured in thin enough samples to achieve the phase-matching bandwidth necessary for very short pulses [129]. The second harmonic of any pulse with a centre wavelength below 380 nm would be absorbed in the crystal and could not be observed, precluding SHG FROG measurements in the DUV.

To generate a FROG signal from a DUV pulse, another nonlinear interaction has to be employed. One option is to use self-diffraction [63] or a transient grating [11], since these do not up-convert the pulse in frequency in order to create the signal. The downside of these methods is that they are based on third-order interactions, so that significantly more pulse energy is required to make the measurement. In the case of a transient-grating FROG, three beams need to be crossed in the nonlinear medium, increasing experimental complexity. This is not required for self-diffraction FROG, however it suffers from increased background in the signal as a consequence [60].

Another issue that influences the choice of nonlinear interaction is that beam transport is challenging in the DUV: the protective coating on most broadband UV laser mirrors,

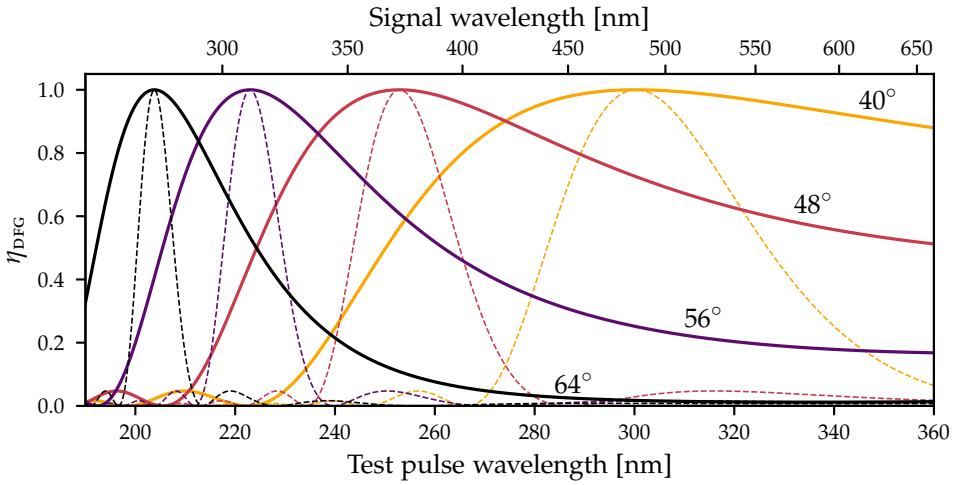

Figure 5.1.: Type I phase-matching efficiency for DFG between a DUV test pulse and a gate pulse at 790 nm in a 5 µm thick BBO crystal used at four crystal angles between 40° and 64°. The dashed lines show the same, but for a 20 µm thick crystal. The top axis shows the wavelength of the DFG signal resulting from down-conversion of the corresponding wavelength on the bottom axis. The change in effective thickness at different angles is taken into account, assuming a crystal angle at normal incidence of 48°.

which are made of aluminium, absorbs light at wavelengths below 250 nm, so that aluminium mirrors coated with magnesium fluoride (MgF_2) have to be used instead. These are both more expensive and significantly more sensitive to environmental damage [130]. When using a Kerr nonlinearity, this disadvantage applies to both the test pulse and the FROG signal.

It would thus be preferable to *down-convert* the test pulse in frequency in a nonlinear process to generate a signal further towards the visible part of the spectrum, and using this as the FROG signal. Given that the DUV generation process is driven by a pulse at 790 nm, this is most easily achieved by using DFG between this and the test pulse. As shown in fig. 5.1, the DFG signal resulting from a pulse centred anywhere in the DUV spectral range lies above 250 nm. This is longer than the cut-off wavelength both

for transmission through BBO crystals as well as for the cheaper and more robust type of aluminium mirror. Since there are two pulses involved in the interaction, this type of measurement is best described as a frequency-resolved cross-correlation of the two pulses, and is appropriately known as cross-correlation frequency-resolved optical gating (XFROG) [131, 132].

As the DUV source is widely tunable, the measurement must be applicable to pulses of different wavelengths. This capability is provided by angle-tuning of the nonlinear crystal. Figure 5.1 shows the type I phase-matching efficiency for DFG in BBO between a DUV pulse and one at the wavelength of the driving field for several different angles between the propagation direction and the crystal axis, taking into account the changing effective thickness as the crystal is rotated away from normal incidence. The wavelength of peak phase-matching efficiency can be tuned across the entire range from 200 nm to 300 nm by rotating the crystal by only 24°. Since BBO, like most materials, becomes more dispersive at short wavelengths, the phase-matching bandwidth decreases at the low end of this range, however it is still sufficient for the measurement of sub-5 fs pulses when using a thin (5 µm) crystal.

In principle, it is possible to achieve larger phase-matching bandwidths for a given crystal thickness by using type II phase-matching. However, for type II phase-matching the effective nonlinearity of BBO scales with the crystal angle as $\cos^2\theta$ rather than $\cos\theta$ as in the case of type I phase-matching, and additionally θ is larger than for type I phase-matching by around 10° [17]. In combination, these lead to a decrease in signal intensity by a factor of ca. 10 for the DFG process used here. Importantly, for type I phase-matching it is the two fields at longer wavelengths that are in the ordinary wave [21]. In the case of DFG, these are the input pulse at the longer wavelength and the signal pulse. Therefore the two pulses crossing in the crystal have to be polarised in orthogonal directions. This is satisfied by the polarisation rotation in the attenuator for the driving pulse in the DUV generation apparatus.

The first demonstrated technique to measure an unknown pulse using XFROG requires the second pulse, often termed the *gate*, to be well characterised [131]. In this

case, the reconstruction algorithm is altered to generate the FROG trace amplitude $\sigma(\omega, \tau)$ (see section 2.4.3) using the gate pulse and the current guess for the test pulse. The disadvantage of this method is that a precise characterisation of the gate pulse at point of the nonlinear interaction is necessary; any errors in this measurement affect the reconstructed test pulse. Furthermore, for an XFROG measurement of this type to yield reliable results, the two pulses interacting need to be of approximately equal duration [131, 132]. For the measurement of potentially sub-5 fs pulses, this poses a major technical hurdle. Removing knowledge of the gate pulse profile, however, makes the phase retrieval problem significantly harder. In fact, it can be shown that without additional information, many ambiguities exist in such a so-called *blind* FROG trace, precluding useful pulse measurements [133].

5.1.1. Ptychographic reconstruction

The most successful solution to the issue of blind XFROG is to borrow a technique from the field of lens-less imaging. The process of retrieving the object function, which describes the sample, from the data in these experiments is a phase retrieval problem as in a FROG trace, and many algorithms to solve it have been developed [134].

One relatively recent method is the ptychographic iterative engine (PIE) [135, 136]. The algorithm was developed for ptychography experiments, in which lens-less imaging is performed with the illumination beam scanned across several overlapping positions on the sample [137]. This is analogous to an XFROG measurement, because the field produced in the sample plane can be written as the product of the illumination field, or probe, $P(\mathbf{r})$ and the object function $O(\mathbf{r})$ [135]. Scanning the probe across the target by displacements $\mathbf{R}$ yields an exit field ψ, given by

$$\psi(\mathbf{r}, \mathbf{R}) = O(\mathbf{r})P(\mathbf{r} - \mathbf{R})\,. \tag{5.1.1}$$

Since the image is acquired in the far field, Fraunhofer diffraction can be applied [72], which means that the image is the intensity of the Fourier transform of ψ:

$$\Psi(\mathbf{k}, \mathbf{R}) = \left| \int_{-\infty}^{\infty} O(\mathbf{r})P(\mathbf{r} - \mathbf{R})\mathrm{e}^{-i\mathbf{k}\cdot\mathbf{r}}\mathrm{d}\mathbf{r} \right|^2 , \tag{5.1.2}$$

which has a very similar form to that of a FROG trace as shown in eq. (2.4.5). The PIE was later improved upon to allow the reconstruction of both the object and the probe functions, with the updated version named the extended ptychographic iterative engine (ePIE) [138, 139]. This opened the door for blind XFROG measurements to be carried out using ptychographic methods.

The ePIE was recently used to retrieve the temporal profile of laser pulses in an XFROG measurement [140, 141]. The additional information needed to break the ambiguities in a blind XFROG trace can be supplied in the form of the power spectrum of one of the two unknown pulses [142, 143]. Crucially, the gate pulse can be much longer than the test pulse for this reconstruction method; in fact, since the "illumination" samples of the pulse need to overlap for the reconstruction to succeed, a longer gate pulse allows for larger delay steps. The time resolution of the measurement is not compromised by this, since the ePIE is capable of retrieving features on a timescale much faster than the sampling step in delay [140, 143].

The ePIE was further improved by adding a regularisation term to the algorithm. This essentially changes the cost function which the algorithm seeks to minimise such that large changes in the reconstructed fields are penalised. This leads to faster convergence and a more reliable reconstruction [144]. The resulting algorithm is correspondingly called the regularised ptychographic iterative engine (rPIE).

All measurements of ultrashort laser pulses using ptychographic reconstruction in the literature have been carried out using SFG as the nonlinear interaction [140–143, 145–147]. This is likely partly because SFG maps most closely onto the form of the signal which the ePIE was developed for, that is, the Fraunhofer diffraction pattern shown in eq. (5.1.2). The signal in an SFG XFROG measurement is given by

$$S(\omega, \tau) = \left| \int_{-\infty}^{\infty} E(t) G(t - \tau) \mathrm{e}^{i\omega t} \mathrm{d}t \right|^2 , \tag{5.1.3}$$

where $G(t)$ is the field of the gate pulse. Note that we have dropped the subscript from the test pulse to maintain legibility in the rest of this section. This is exactly eq. (5.1.2), with the frequency ω taking the place of spatial frequency $\mathbf{k}$ and the delay τ taking

that of the displacement **R**. For a DFG XFROG, on the other hand, the signal is

$$S(\omega,\tau) = \left| \int_{-\infty}^{\infty} E(t)G^*(t-\tau)\mathrm{e}^{i\omega t}\mathrm{d}t \right|^2 , \tag{5.1.4}$$

that is, the gate pulse enters the signal field as the complex conjugate. Recall that the analytic signal of the gate field, $G(t)$, contains only positive frequencies; the complex conjugate therefore contains only negative frequencies, leading to subtraction rather than addition of the frequencies of $E(t)$ and $G(t)$. The PIE algorithm as used in the literature has to be adapted to reflect this difference. For a derivation of the update functions used below, see appendix A.

The ptychographic algorithm for reconstructing XFROG traces differs from other methods, such as the PCGPA, in that it treats each delay sample in the data separately. Each delay sample is a slice of the XFROG trace, $S(\omega,\tau_j)$. One complete iteration of the algorithm consists of using each delay sample in turn, with the order in which they are used randomised at each iteration [142]. In the following we will denote the outer loop index by n and the inner loop index by j, since the reconstructed fields are updated with each delay slice. Furthermore, we will treat the time and frequency domains as continuous to keep expressions simple, however in practice time and frequency are sampled at discrete points and Fourier transforms must be carried out using the discrete Fourier transform [114]. For each step in the inner loop, first the DFG field is calculated using the current guess for the test pulse, $E_{n,j}(t)$, and the gate pulse, $G_{n,j}(t)$, with the latter shifted by the value of the delay τ_j:

$$\psi_j(t) = E_{n,j}(t)G^*_{n,j}(t-\tau_j)\,. \tag{5.1.5}$$

This is then Fourier transformed to yield the XFROG amplitude $\psi(\omega)$, and its magnitude is replaced by the square root of the measured signal $S(\omega,\tau_j)$. This yields the updated DFG field $\psi'(t)$ by the inverse Fourier transform:

$$\psi'_j(t) = \mathcal{F}^{-1}_{\omega}\left[\frac{\mathcal{F}_t\left[\psi_j(t)\right]}{\left|\mathcal{F}_t\left[\psi_j(t)\right]\right|} \sqrt{S(\omega,\tau_j)} \right] . \tag{5.1.6}$$

The difference between $\psi(t)$ and $\psi'(t)$ is then used together with a weight function to

update the test pulse:

$$E_{n,j+1}(t) = E_{n,j}(t) + \frac{G_{n,j}(t-\tau_j)}{(1-\alpha)\left|G_{n,j}(t-\tau_j)\right|^2 + \alpha\left|G_{n,j}(t-\tau_j)\right|^2_{\max}}\left[\psi'_j(t) - \psi_j(t)\right], \quad (5.1.7)$$

where α is a real-valued parameter, and $Q_{\max}$ denotes the maximum value of the quantity Q. Note that in the version of this algorithm for SFG XFROG, the complex conjugate of the gate pulse appears in the numerator of the weight function [142]. In the ePIE, only the term proportional to $\left|G_{n,j}(t-\tau_j)\right|^2_{\max}$ is present in the denominator of the weight function. The term $(1-\alpha)\left|G_{n,j}(t-\tau_j)\right|^2$ is due to the regularisation modification introduced in the rPIE.

The gate pulse is updated in a symmetrical fashion to the test pulse:

$$G_{n,j+1}(t-\tau_j) = G_{n,j}(t-\tau_j) + \frac{E_{n,j}(t)}{(1-\alpha)\left|E_{n,j}(t)\right|^2 + \alpha\left|E_{n,j}(t)\right|^2_{\max}}\left[\psi'_j(t) - \psi_j(t)\right]^*. \quad (5.1.8)$$

Here, the different form of the interaction for DFG leads to conjugation of the weight function as well as the difference term $\psi'_j(t) - \psi_j(t)$ as compared to the SFG case. Note also that the gate pulse has to be shifted by $-\tau_j$ after this procedure to retrieve the updated guess $G_{n,j+1}(t)$. The gate pulse is furthermore projected onto its power spectrum $I_{\mathrm{G}}(\omega)$, which has to be measured independently:

$$G_{n,j+1}(t) = \mathcal{F}^{-1}_{\omega}\left[\frac{\mathcal{F}_t\left[G_{n,j}(t)\right]}{\left|\mathcal{F}_t\left[G_{n,j}(t)\right]\right|}\sqrt{I_{\mathrm{G}}(\omega)}\right]. \quad (5.1.9)$$

This way, the added information is used at each step and convergence can be achieved [142]. Note that apart from the asymmetry in the interaction for the case of DFG XFROG, the only difference in the algorithm between the gate and test pulses is the additional projection on the power spectrum for the gate. After each iteration through all delay slices, the complete reconstructed XFROG trace is generated according to eq. (5.1.4), and the FROG error Δ is calculated as in eq. (2.4.7). The algorithm is run until Δ converges to a minimum.

Implementation details

Given that the interaction in DFG XFROG involves pulses at three different wavelengths, the necessary sampling window in the frequency domain is very large, extending from

the DUV to the near infrared (NIR). The window in the time domain, on the other hand, must be able to accommodate the gate pulse when shifted to both extremes of the delay scanning range. This range, in turn, is determined chiefly by the gate pulse duration, as this is the dominant contribution to the cross-correlation width. The time window must therefore be at least twice as wide as the gate pulse duration.

The FWHM is an inappropriate measure for the cross-correlation width and pulse duration in this case, since it excludes the contribution of the weaker wings of the signal. Instead, the entire time over which there is detectable signal has to be used; this is several times longer than the FWHM. Depending on the experimental signal-to-noise ratio, the resulting required window width can easily run to several hundred femtoseconds, leading to a large number of time and frequency samples and correspondingly slow execution of the algorithm.

To avoid this issue, it is useful to shift all three pulses to the baseband. For the signal pulse, this is done by determining the central frequency of the XFROG trace as the first moment of the frequency marginal,

$$\omega_0^{\mathrm{S}} = \frac{\sum_{ij} \omega_i S(\omega_i, \tau_j)}{\sum_{ij} S(\omega_i, \tau_j)}, \tag{5.1.10}$$

then shifting the frequency axis of the trace by ω_0^{S} and interpolating the trace onto a common baseband frequency axis $\tilde{\omega}$, which is also used for both the gate and test pulses. Even without frequency shifting, interpolation is necessary in order to obtain a uniformly spaced frequency axis from the non-uniform axis in the raw data[1], which allows the use of significantly faster algorithms for the discrete Fourier transform [114].

For the gate pulse, the first moment of the separately measured power spectrum is used:

$$\omega_0^{\mathrm{G}} = \frac{\sum_i \omega_i I_{\mathrm{G}}(\omega_i)}{\sum_i I_{\mathrm{G}}(\omega_i)}, \tag{5.1.11}$$

and the power spectrum is interpolated onto the baseband frequency axis as well. The central frequency of the test pulse is then determined by the gate pulse and trace

[1] The non-uniformity in frequency space is a consequence of the operating principle of optical spectrometers: because the frequency resolution is created by diffraction off of a grating, the displacement on the detector array is (approximately) proportional to wavelength rather than frequency.

moments:

$$\omega_0^{\mathrm{t}} = \omega_0^{\mathrm{s}} + \omega_0^{\mathrm{G}} \,. \tag{5.1.12}$$

The actual gate and test pulses can then be retrieved from the baseband signals by simply shifting their frequency axis back by the appropriate central frequency. Because the reconstructed test pulse satisfies the form constraint, that is, its DFG signal with the reconstructed gate pulse yields the experimentally acquired spectrum to within the FROG error Δ, this shifting to the baseband and inferring the test pulse frequency neither adds nor removes any information to the process.

At several points in the ptychographic reconstruction algorithm, it is necessary to shift a pulse in time by some delay τ and back. This is most easily achieved by exploiting the shifting property of the Fourier transform:

$$E(t-\tau) = \mathcal{F}_{\omega}^{-1}\left[\mathrm{e}^{i\omega\tau}\mathcal{F}_t[E(t)]\right] . \tag{5.1.13}$$

The weighting parameter α is critical in the reconstruction. Viewing the PIE as an optimisation algorithm [145], it determines the size of the step taken in the direction of steepest descent; that is, in the direction towards the minimum of the cost function. To keep the process from stagnating in or around a local minimum, it is useful to choose α randomly [142]. In the implementation described here, a new random value for α between 0.2 and 1 is chosen for each iteration of the outer loop.

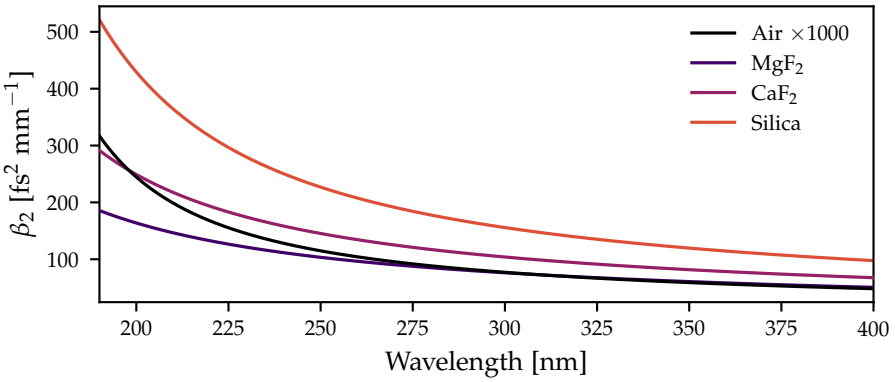

Figure 5.2.: The GVD of three common materials for vacuum windows and optics as well as air. Note that the dispersion of air has been multiplied by a factor of 1000 to display it on the same scale.

5.2. DUV XFROG apparatus

The basic principle of an XFROG apparatus is very similar to that of an SHG FROG, except that two beams enter the optical set-up and are overlapped in the nonlinear medium. Effectively, this simply moves the beam-splitting step of the FROG layout to a much earlier point, since both pulses in an XFROG measurement are ultimately derived from the same laser. In the case of XFROG in the DUV, however, there is one very important complication to this principle, and that is the dispersion of both the measurement apparatus and the optical path preceding it.

As can be seen in fig. 5.2, the GVD of materials commonly used for transmissive optics in the visible and DUV increases rapidly with decreasing wavelength, rising to several hundred square femtoseconds per millimetre – for comparison, the GVD of silica glass at 790 nm is only $36\,\mathrm{fs}^2\,\mathrm{mm}^{-1}$. Even air, while much less dispersive than any bulk material, becomes a significant source of dispersion over distances of metres. $100\,\mathrm{fs}^2$ of GDD is sufficient to stretch a 5 fs pulse to a duration of more than 50 fs, so this amount of dispersion must be either avoided or compensated.

Unfortunately, the common schemes for dispersion compensation are either impracti-

cal or unavailable in the DUV spectral region. No suitable materials with anomalous dispersion in this wavelength region exist, so bulk dispersion cannot be used. Prism compressors [148], commonly used in the infrared wavelength region, may compensate positive GDD, however the third-order dispersion of the prism material would distort the pulse significantly, erasing the benefits of removing the quadratic phase. Grating compressors do not have this issue, however the dispersion induced in these devices is usually very large, so that a stretcher-compressor combination has to be used in order to compensate for smaller amounts of GDD [15]. This leads to a large number of reflections (or more accurately, diffractions) off of gratings, greatly reducing the energy throughput. Lastly, while there is no fundamental reason why dispersive mirrors cannot be manufactured for the DUV, this is not one of the main areas of interest for these optics, and there are no commercial suppliers.

This lack of suitable means of dispersion compensation makes delivering a short pulse to an experiment a matter of avoiding dispersion wherever possible. As a consequence, no transmissive optics such as lenses or windows can be used, and the entire beam path has to be kept in vacuum. The vacuum requirement is only a minor addition to the experimental effort when using the DUV source in an experiment together with an SXR pulse, since in this case the other half of the pump-probe pair must be propagated under vacuum in any case. The inability to use windows, however, makes connecting the DUV source to the apparatus for characterisation or an experimental chamber significantly more challenging.

The phase-matching pressure for DWE at DUV wavelengths lies in the range of several atmospheres for argon gas and tens of atmospheres for neon. This high a gas pressure at both ends of the anti-resonant fibre cannot practically coexist with a vacuum environment for the rest of the beam path without a window separating the two regions. In the absence of a window, keeping the exit end of the fibre under vacuum is therefore the only solution. As shown in section 4.2.3, the resulting pressure gradient along the fibre complicates the dynamics of self-compression and DWE, but similar pulses can be generated.

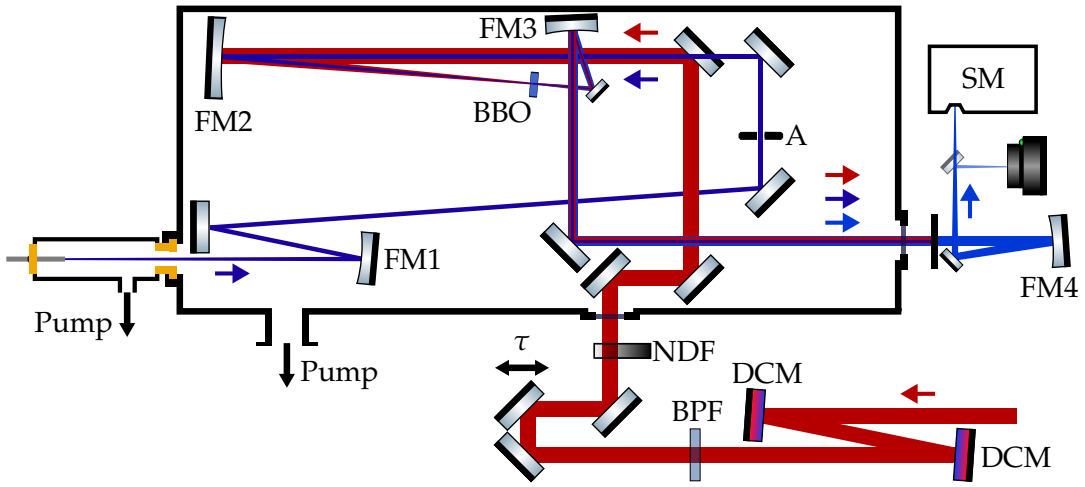

Figure 5.3.: Optical layout of the DUV XFROG system. A detailed description can be found in the text. FM: focusing mirror; A: aperture; NDF: neutral density filter; BPF: bandpass fiter; DCM: dispersion-compensating mirror; SM: spectrometer; τ: delay stage. The path of the test pulse is shown in purple, that of the gate pulse is shown in red, and that of the signal pulse is shown in blue.

Vacuum propagation and the pressure gradient along the fibre is maintained by enclosing part of the XFROG apparatus in a vacuum chamber. In particular, the entire beam path of the DUV test pulse between the exit of the fibre and the nonlinear interaction in the crystal is kept in vacuum. The delay line, which is incorporated into the path of the gate pulse, is outside, as are the parts of the system which select and analyse the signal pulse, since dispersion is no longer an issue here. High vacuum is not required; at 1 mbar, the UV dispersion of air is already reduced by a factor of 1000 and thus negligible. Since the conductance of a small-core PCF is very small [149], the pressure in the vacuum chamber does not exceed 0.5 mbar, even when the entrance of the PCF is filled with over 20 bar of gas.

The optical layout of the system is shown in fig. 5.3. The DUV pulse exiting the fibre is collimated by an $f = 30$ cm spherical mirror (FM1). Afterwards it is steered

towards another $f = 30\,\text{cm}$ spherical mirror (FM2), which focuses it into the nonlinear medium, a nominally 5 µm thick BBO crystal on a 1 mm silica glass substrate (Lightkey Technology). All optics between the exit of the fibre and the crystal are MgF_2-coated aluminium mirrors (Acton Optics). To reduce the total energy focused into the crystal, and thus the chances of damage occurring, the aperture (A) is closed to a diameter of approximately 5 mm to only allow the central part of the beam through. Due to the spatio-spectral coupling in the beam, this predominantly attenuates the visible and infrared parts of the spectrum with little effect on the DUV pulse energy.

The gate pulse reflects off of two dispersion compensation mirrors (FemtoOptics) before passing through a bandpass filter, the passband of which is centred at 790 nm and 10 nm wide (FWHM). After the delay line, consisting of a retro-reflecting pair of mirrors on a motorised delay stage (Thorlabs), the now-narrowband pulse passes through a variable neutral density filter (Thorlabs) and enters the vacuum chamber through a 6 mm thick anti-reflection coated window made of silica glass (Eksma optics). It is steered towards focusing mirror FM2, with the last steering mirror being a half-mirror with its top edge 1 mm underneath the height of the DUV beam. The two beams thus cross in the crystal with the crossing angle in the vertical plane. The crossing angle is below 1°.

The vertical crossing angle is useful for the beam steering after the crystal; since the signal beam appears on the outside of the ingoing beam pair (see fig. 5.4), subsequent optics must not be too small lest they clip the beam. Since reflection at 45° in the horizontal plane reduces the effective aperture of the steering mirrors by a factor of $\sqrt{2}$ in the horizontal dimension, propagating the three beams vertically displaced results in easier alignment.

As shown in section 5.1.1, it is useful for the gate pulse to be narrowband for the ptychographic reconstruction algorithm to converge quickly, and a transform-limited gate pulse is not necessary, but nevertheless advantageous. The bandpass filter required for this and the entrance window to the vacuum chamber both add non-negligible amounts of GDD to the pulse; the 6 mm of silica glass for the entrance window alone

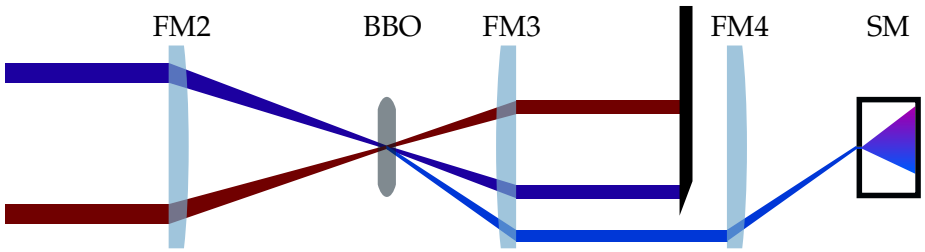

Figure 5.4.: Schematic of the XFROG interaction and subsequent filtering of the signal in a side-on view. The transverse phase-matching in the DFG interaction means that the signal beam exits the nonlinear medium at a steeper angle than either input beam, in this case below the DUV beam. Therefore it can be isolated by a simple beam block. The labels correspond to those in fig. 5.3.

add around 220 fs^2, and the filter is also around 6 mm thick and will thus account for a similar amount of GDD. Combined with the chirp already present in the pulse from the SPM in the capillary, this makes dispersion compensation necessary. The attenuator makes sure that damage to the nonlinear medium can be avoided, and also serves to check that the measurement is not distorted by XPM.

After the nonlinear interaction, the three beams – the test, gate, and signal pulses – are re-collimated by an $f = 10$ cm mirror. Because they are collimated rather than re-imaged straight away, the mask to remove the test and gate beams can remain outside of the vacuum chamber, where it is possible to adjust it. This is very important to enable diagnostics of the system (see below). The exit window is an uncoated 1.5 mm thick piece of MgF_2. Imaging onto the detector is done with an $f = 50$ cm spherical mirror. A flip mirror can be inserted after the re-imaging mirror to instead send the beam towards a UV-sensitive camera (EHD Imaging), which is also located in the focal plane of the beam.

The two most critical factors when aligning any two-pulse system, be it a transient-absorption experiment or an XFROG measurement, are the spatial and temporal overlap between the pulses. Due to space constraints in the vacuum chamber, the spatial overlap

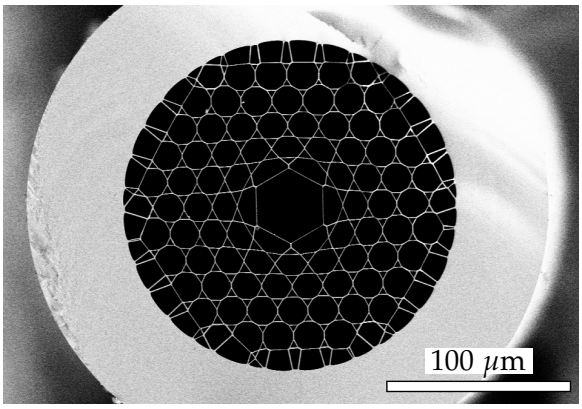

Figure 5.5.: Cross-section scanning electron micrograph of the 33 µm core diameter kagomé PCF that was in use for the characterisation of DUV pulses. Image courtesy of Francesco Tani, Max Planck Institute for the Science of Light, Erlangen, Germany.

cannot be ascertained by simply placing a camera in the focal plane of the focusing mirror. Instead, the re-imaged focus is used by inserting the flip mirror as shown in fig. 5.3. This has the additional advantage of being usable with the chamber under vacuum, so that any misalignment caused by deformation of the chamber during the pumping process can be diagnosed and corrected.

A common method of finding temporal overlap is to increase the intensity of one or both of the interacting pulses, focusing both into a thick third-order nonlinear medium such as glass, and looking for the disturbance of one beam by the nonlinear refractive index profile induced by the other. Unfortunately, this method cannot be used here. Only the gate pulse is sufficiently energetic to cause a strong nonlinear refractive index variation, however it is significantly larger than the DUV pulse both in space (because of its longer wavelength) and in time (because of the spectral bandpass filter). Any induced refractive index changes will therefore be experienced by the entirety of the DUV pulse, and the resulting average phase shift will not be easily detectable.

The simplest way of temporally overlapping the gate and test pulses is therefore to look for the DFG signal directly. This cannot be done by eye, since the signal beam

is invisible; because it is also very weak, even UV detector cards yield no detectable signal. So as not to rely on alignment into the spectrometer, the UV-sensitive camera is used for the temporal overlap as well. With the selection mask set such that only weak scatter from the DUV beam (the middle of the three) reaches the camera, so that accidental blocking of the signal beam is avoided, the delay is scanned in steps of tens of femtoseconds until the signal is found. Initially, this is most easily done without evacuating the chamber, since the resulting air path stretches the DUV pulse to several hundred femtoseconds, increasing the range over which signal can be detected. The location of temporal overlap changes when the air is removed from the chamber, however the amount is easily predicted by measuring the length of the optical path.

One disadvantage of the apparatus is that due to the reflectivity of the mirrors after the nonlinear crystal, it is not possible to measure the spectrum of the DUV pulse with the chamber evacuated, unless the whole spectrum is above 250 nm. As a consequence, reference spectra were taken only for some example pulses.

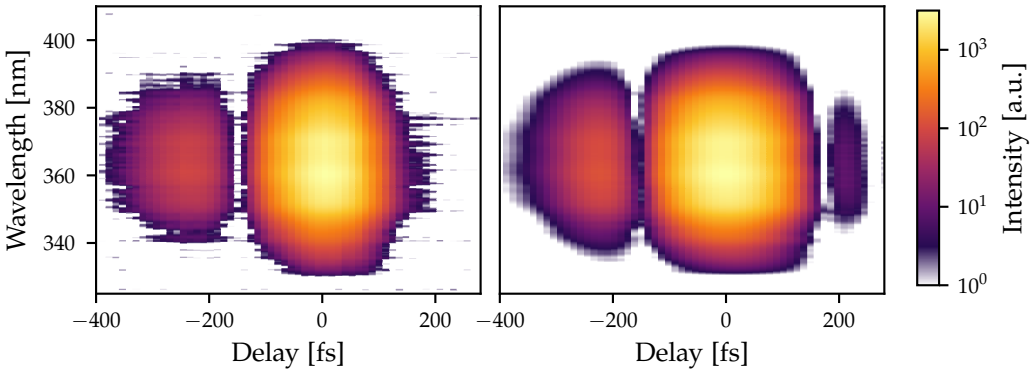

Figure 5.6.: Example DFG XFROG trace of a pulse centred around 250 nm, obtained by coupling a 16 fs, 1 µJ pulse at 790 nm into a 17 cm long, 33 µm core diameter PCF filled with a negative pressure gradient from 8 bar to vacuum. The gate pulse is approximately Gaussian in shape with a duration of 133 fs (see fig. 5.7). Left: the measured trace. Right: the reconstructed trace after ca. 300 iterations of the rPIE algorithm.

5.3. Tunable sub-10 fs DUV pulses

With the conversion efficiency sufficient to achieve the pulse energies necessary for an experiment, the XFROG measurement must primarily establish that a short pulse can be obtained at any chosen wavelength in the DUV. For that purpose, measurements were taken with several different gas pressures in the PCF to tune the phase-matching conditions.

A typical XFROG trace is shown in fig. 5.6. The most important feature of the measured trace is that there is no evidence of strong chirp in the pulse, since the trace is not visibly tilted. The test and gate pulses retrieved from this measurement are shown in fig. 5.7, along with an additional measurement of the gate pulse. Note that the gate pulse in these two cases has not experienced the exact same amount of GDD due to the different air path and the fact that a separate piece of glass was used to add approximately the same GDD as the vacuum chamber window.

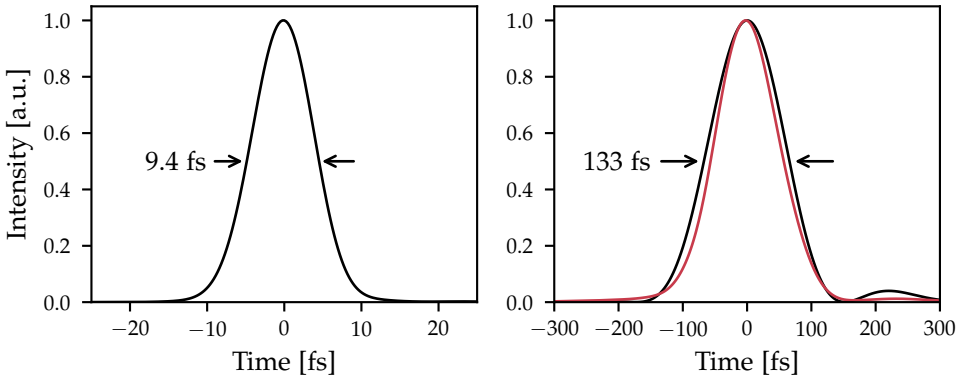

Figure 5.7.: Test (left) and gate (right) pulses as retrieved from the XFROG trace shown in fig. 5.6. The red line shows the gate pulse as measured with the SHG FROG device discussed in section 4.3.3 after propagating the pulse through a 6 mm thick piece of silica glass to approximate the influence of the vacuum chamber window.

The reconstructed test pulse is somewhat longer than would be expected from both the numerical modelling and non-vacuum measurements in the literature [11]. However, this is likely not caused by the pulse itself, but rather by the measurement; it is likely that the phase-matching bandwidth was insufficient.

Evidence that the phase-matching window was too narrow can already be seen in the reconstructed XFROG trace. The reconstruction algorithm effectively forces the convolution of the gate and test pulse spectra to reproduce the frequency marginal of the trace. Because the gate pulse spectrum is fixed in the reconstruction algorithm (see section 5.1.1), only the spectrum of the test pulse is free to vary to satisfy this constraint. If the phase-matched bandwidth, which determines the width of the marginal, is reduced, this decreases the reconstructed bandwidth of the test pulse. The width of the delay marginal, on the other hand, is set only by the duration of the two pulses and is not affected by the phase-matching as long as neither pulse has a very strong spectral phase gradient. As a consequence, insufficient phase-matching bandwidth leads to an inconsistency between the widths in delay and frequency of the measured

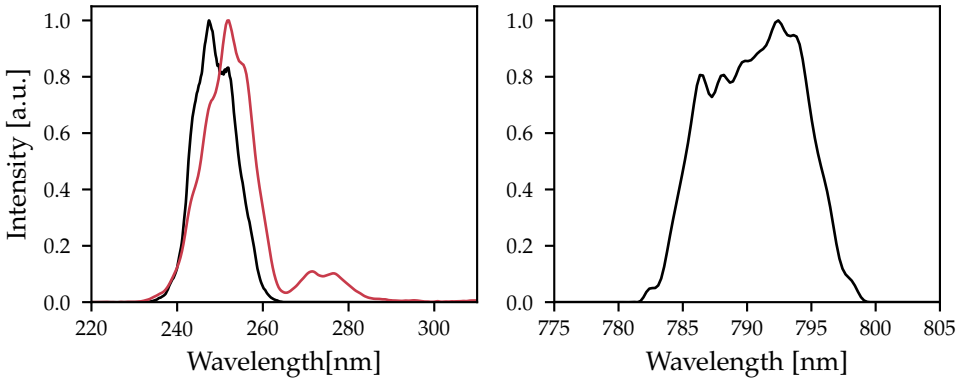

Figure 5.8.: Left: normalised spectrum of the test pulse as reconstructed from the XFROG measurement (black) and as measured independently (red). Right: spectrum of the gate pulse as used in the XFROG trace reconstruction.

trace. The reconstructed trace then appears stretched in delay when compared to the measured one, since this inconsistency can only be accommodated by assuming longer pulse durations. This effect can be seen in fig. 5.6

As shown in fig. 5.1, the phase-matching bandwidth is strongly affected by the thickness of the medium. Therefore it is likely that the observed discrepancy is caused by the nonlinear medium being too thick due to a manufacturing error.

Figure 5.8 shows the normalised reconstructed and externally measured spectra of the test pulse. The transform-limited pulse duration as obtained from the externally measured spectrum is 4.5 fs. Here, evidence of poor phase-matching is visible in a shift in the spectrum as well as a reduction in bandwidth, with the shift partly due to a small misalignment of the crystal angle leading to a shift of the phase-matching efficiency curve. The misalignment is caused by the need to set the crystal angle with the vacuum chamber at atmospheric pressure; the DUV pulse is then chirped to over 100 fs, making the crystal angle at which maximum overall signal is obtained delay-dependent and thus adding uncertainty as to the correct angle. It should be noted that with a 5 µm thick crystal, the phase-matching bandwidth would be sufficient for such a small uncertainty to have a negligible effect on the spectrum.

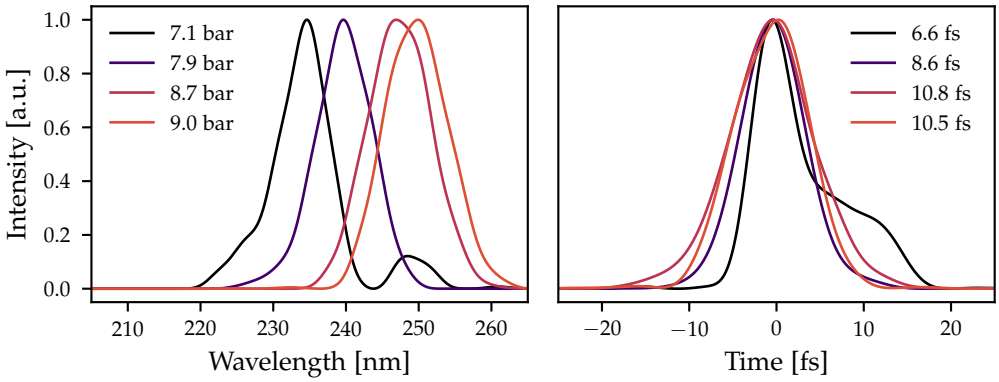

Figure 5.9.: Spectra (left) and pulse profiles (right) obtained from the XFROG measurement for a 2.4 µJ, 30 fs pulse coupled into the same PCF as in fig. 5.6, filled with a negative pressure gradient of varying initial pressures (see legend). Matching line colours in the two plots indicate data from the same measurement. The legend in the right-hand side plot shows the FWHM duration of each retrieved pulse.

In the case of SHG FROG, the mixing of all possible frequencies from both (identical) pulses makes reconstruction of traces that are strongly affected by phase-matching difficult or even impossible, and correction factors have to be applied to account for the limited bandwidth [64]. This is not the case for XFROG with a narrowband gate pulse, since here the frequency conversion can be approximately treated as a pure shift – in this case, a subtraction of the gate pulse frequency. Effectively, this view approximates the gate pulse spectrum as a delta function. Therefore, the phase-matching efficiency acts only as a bandpass filter on the trace and consequently on the reconstructed pulse. As long as the pulse is not too complex, that is, in the absence of strong spectral phase gradients, the pulse duration as measured here then forms an upper bound, with the real pulse duration likely significantly shorter.

Despite the lack of appropriate phase-matching, the measurement can yield information on the change in pulse duration as the DWE wavelength is changed through the gas pressure. The result of keeping the input pulse energy constant while changing the

pressure is shown in fig. 5.9. It was found that the DWE wavelength could be changed to a greater degree while retaining detectable signal and relatively simple spectral shapes when the input pulse was long. The pre-compression capillary was therefore evacuated to remove the spectral expansion and the resulting dispersion compensated using the grating compressor inside the laser.

Owing to the limited phase-matching bandwidth, spectra at the edge of the range of the pressure scan are likely reshaped strongly by the measurement, however their central wavelength is a guide to the tunability achieved. As shown in the time-domain plot in fig. 5.9, the pulse duration remains around or below 10 fs as its central wavelength is changed by around 20 nm. With this duration forming an upper bound, it is likely that the pulses were in fact well below 10 fs in duration at all central wavelengths, evidencing that the tunability of the source does not compromise the achievable temporal resolution.

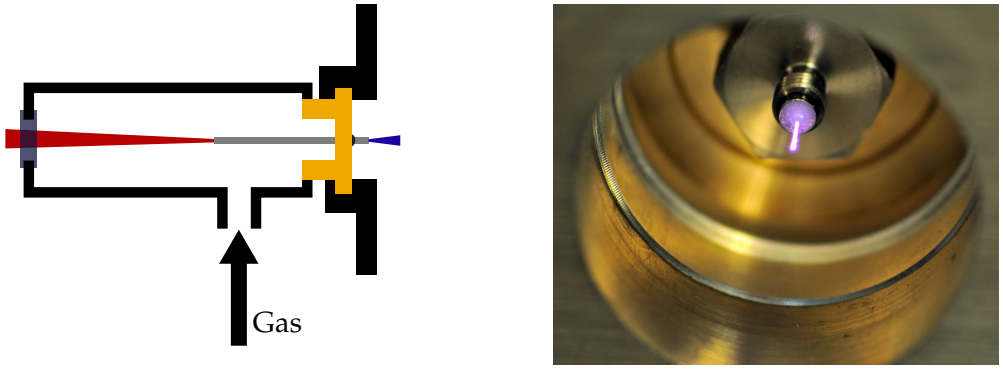

(a) (b)

Figure 5.10.: (a) Altered mounting method for very short lengths of PCF. The fibre is bonded into the same part that connects the gas cell to the vacuum chamber. This makes the second chamber unnecessary and allows for the mounting of shorter fibres.
(b) Photograph of the exit end of the fibre in operation, looking from the inside of the vacuum chamber. The purple glow is due to a cladding resonance around 400 nm leading to strong scattering of light at that wavelength.

5.4. Sub-5 fs DUV pulses

Even at their transform limit, the pulses generated in the relatively long fibre discussed in section 5.3 are not as short as required for an experiment with true few-femtosecond resolution. They are also somewhat longer than those predicted by the numerical modelling for a static fill (see section 4.2.2). The reason for this lies chiefly in the length of the fibre, with a shorter fibre required to generate very short pulses as described in section 4.2.3. The length of the fibre can be reduced by altering the DUV generation apparatus to only use one gas cell (see fig. 5.10a).

The phase-matching bandwidth using a thick crystal is clearly insufficient to measure pulses with even larger bandwidth than those discussed in section 5.3. However, this can be remedied to some degree by rotating the crystal during the measurement in a

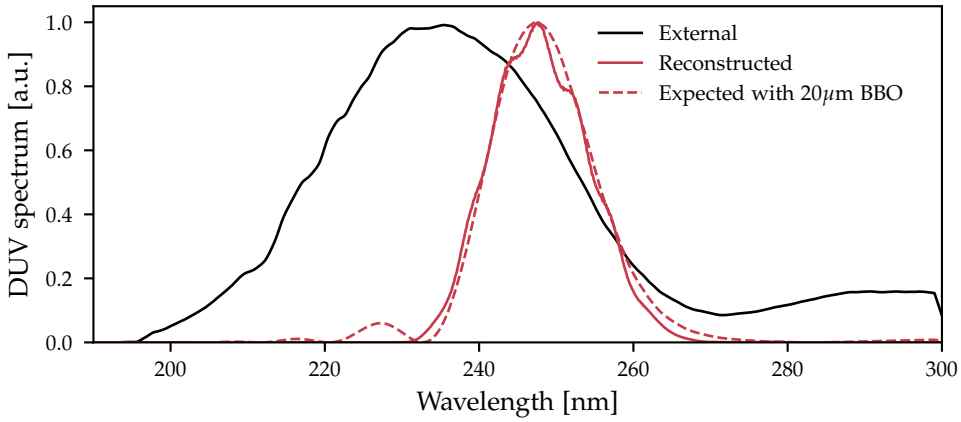

Figure 5.11.: Normalised spectra of a DUV pulse generated with 3 µJ input energy and an entrance pressure of 10.9 bar as measured externally (black) and reconstructed by an XFROG measurement without angle dithering (red). The red dashed line shows the externally measured spectrum after multiplication by the phase-matching efficiency for a BBO crystal which is 20 µm thick rather than the nominal 5 µm.

technique known as *angle dithering* [150]. The most important criterion for this to be successful is that the dithering be faster than the acquisition period at each delay point in the trace. Only then can the dithering be treated as a continuous function in the analysis, which leads to a pure increase in phase-matching bandwidth without adding artefacts.

By comparing the reconstructed pulse spectrum to an externally measured one, it is possible to determine the thickness of the nonlinear medium, since the relationship between the two can be approximated as a simple multiplication by the phase-matching efficiency. The spectrum of the more broadband pulse generated in the shorter fibre is much wider than the observed phase-matching bandwidth, so that any spectral shift due to the angle-sensitivity does not obscure the narrowing effect. Figure 5.11 shows this comparison for a particular case. The crystal angle was determined by aligning the incoming beam to normal incidence, so that the only free parameter in the calculation

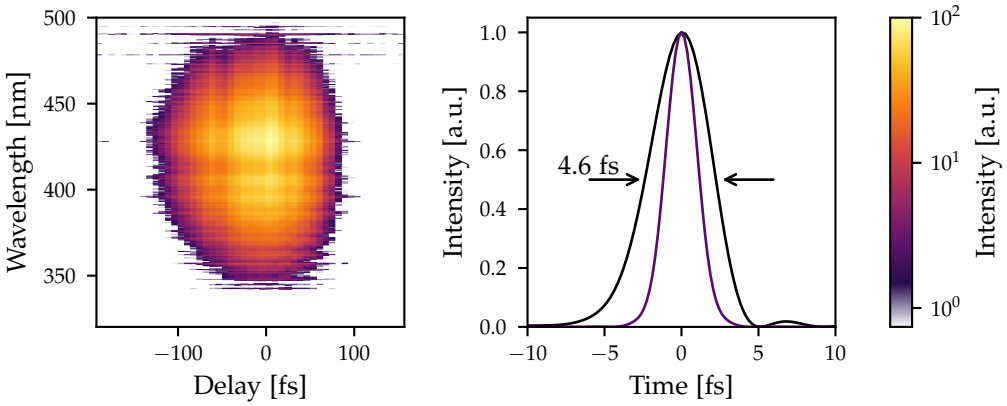

Figure 5.12.: Measured trace (left) and reconstructed test pulse (right) from an angle-dithered XFROG measurement. In this case the entrance pressure in the fibre was 11.5 bar and the 1 µJ input pulse was pre-compressed to 16 fs. The purple curve shows the transform-limited test pulse as retrieved from the externally measured spectrum shown in fig. 5.13. The duration of the transform-limited pulse is 2.4 fs

is the crystal thickness. The excellent agreement between the two red curves means it is very likely that the crystal was 20 µm thick.

To put the angle-dithering approach into practice, the rotation stage holding the crystal in the XFROG vacuum chamber was exchanged for a motorised version (Thorlabs)[2]. While this provided the necessary range of motion, the speed of the motorised mount was limited to around one degree per second. The necessary angular range to access the whole bandwidth of the pulse was about 15°, so the spectrum at each point had to be averaged over at least 15 s. To ensure that a lack of synchronisation between the crystal rotation and the data acquisition did not lead to uneven weighting in the average over angles, this was increased to one minute.

Since for any given crystal angle there is no signal at most of the wavelengths in the

[2]It should be noted that the incorrect crystal thickness was only discovered after the switch to the shorter fibre had been completed, and thus the dithering technique could not be applied to the measurements discussed in section 5.3.

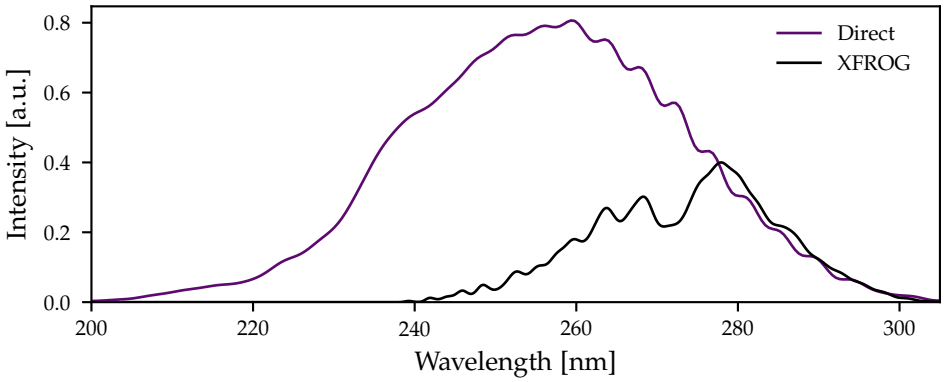

Figure 5.13.: Test pulse spectrum as retrieved from the XFROG measurement shown in fig. 5.12 (black) and as measured with a spectrometer directly (purple).

spectral window of interest, the spectra were not acquired as a single average over the dithering period. Instead, 50 shots were acquired individually, and a threshold was applied to each of these spectra before adding them up. In this way, accumulation of noise due to the lack of signal was avoided.

The rotation of the crystal leads to additional issues that are not present for a single crystal angle. Since the crystal is mounted on a 1 mm thick silica substrate, the rotation causes a displacement of the beam exiting the back surface by refraction in the substrate. The re-imaged focal spot therefore walks off of the spectrometer entrance slit as the crystal is rotated, reducing the signal at the edges of the spectral window. Even with the improved phase-matching bandwidth, the measurement using a thick crystal can thus not truly capture the extremely broadband pulses generated in the shorter fibre.

The XFROG trace acquired using this technique with the output of a 7 cm long PCF is shown in fig. 5.12. To further aid the measurement, the gas pressure in the fibre was increased in order to push the dispersive wave to longer wavelengths, where the phase-matching bandwidth is greater for any given crystal. The increased bandwidth can be seen by comparing the measured trace to that in fig. 5.6; the angle-dithered trace extends almost 150 nm despite the relatively poor signal-to-noise ratio, whereas traces in earlier measurements are around 70 nm wide. As a consequence of the improved

bandwidth, this measurement is capable of retrieving a 4.6 fs pulse.

Unfortunately, the bandwidth of this improved XFROG measurement is still severely limited by the beam walk-off issue discussed above. This can be seen by comparing the reconstructed test pulse spectrum with a reference spectrum, shown in fig. 5.13. Clearly, only part of the spectrum is retrieved, so that the test pulse duration can once again only be treated as an upper bound. The lack of any visible chirp in the trace in combination with the very large bandwidth means that the test pulse is likely below 3 fs in duration, as predicted by the numerical modelling for this combination of parameters.

With the incorrect nonlinear crystal, the measurements presented here cannot unambiguously determine the pulse duration of the dispersive wave, and significant assumptions have to be made to make any quantitative statements at all. However, the primary purpose of the characterisation apparatus is fulfilled: the absence of significant pulse chirp in any of the measurements so far is clear evidence that it is possible to deliver a very short pulse to an experiment by using a negative pressure gradient, and that dispersive wave emission can be a powerful source for ultrafast science. Exactly how short the pulses are will have to be determined in a future experiment using the correct nonlinear medium. This will consist of simply replacing the crystal, since all other parts of the measurement have been proven to work through the results presented already.

6. A Combined Source for Pump and Probe

With dispersive wave emission providing a source of ultrashort pulses that is tunable across the DUV, both the pump and probe pulses for an SXR transient absorption experiment are available. However, the question of delay stability between the two pulses has yet to be addressed.

In the apparatus described in chapters 4 and 5, the DUV pulses are generated using a driving pulse that is split from the output of the laser, while the higher-energy portion of the pulse is converted to the SWIR, compressed, and finally used to drive HHG. The SXR harmonics are then combined with the DUV pulse. The optical path from the splitting point to the recombination mirror is over 12 m long, comprising the internal path of the OPA, beam routing into, through, and out of the SWIR capillary, as well as the in-vacuum pump-probe interferometer. Over this distance, thermal fluctuations may cause femtosecond-scale drifts in the delay between the two pulses, drastically reducing the effective time resolution of the experiment. A large proportion of both paths is outside of vacuum, so that currents and temperature-dependent refractive index changes in the air are present in addition to expansion and contraction of optics and optical tables.

One possible technique to alleviate this problem is to measure the delay between the two arms, for instance via spectral interferometry, and to re-sort the data from delay scans based on this [151]. However, this requires the two pulses to overlap spectrally at least in part, which is not the case for the SWIR pulse driving HHG and the pulse driving DWE. Additionally, the timescale of delay drifts needs to be significantly slower than the acquisition time at each delay point for the re-sorting to be effective. With the

low signal levels in SXR spectroscopy using HHG sources necessitating several minutes of integration at each point[1] [53], this requirement cannot easily be met.

Because of these issues, it would be preferable for the two arms to have common optical paths for as long as possible. At the very least, this requires the DUV generation to be driven by the SWIR pulse. For example, the pulse could be split after the SWIR capillary with both parts being compressed separately and used to drive the two frequency conversion processes, HHG and DWE. However, this would still require the pump-probe interferometer to begin outside the vacuum chamber, or alternatively that a high-pressure gas cell be placed inside the chamber, neither of which are particularly desirable. Furthermore, a substantial portion of the SWIR pulse would have to be split off to drive the DUV generation process.

A better solution is to generate both the few-cycle SWIR pulse and the DUV pump pulse in the same fibre, directly connected to vacuum, and split them inside the chamber. In this way, both the splitting and recombination occur in vacuum and thus close together. Additionally, the two pulses never propagate in air, reducing both the attendant instability issues as well as dispersion, and the driving pulse energy available for HHG is reduced only by the small amount lost in the splitting process. The use of a self-compressed pulse to drive HHG has been proposed [152] and demonstrated [153] before, however so far only propagation in PCF has been considered, limiting the total pulse energy, and the potential to create two pulses has not been explored.

In this chapter I propose a scheme to simultaneously generate both a SWIR and a DUV pulse, both of sub-5 fs duration, in a single waveguide that can be connected directly to a vacuum system. Instead of a PCF, the proposed apparatus uses a simple capillary of the type already used in our laboratory, significantly reducing both the cost and the demands on the beam profile of the initial pulse. I will further demonstrate, using a simple numerical model, that the infrared pulse is capable of generating high harmonics in the SXR range.

[1] In principle, it is possible to integrate for shorter periods and take several spectra at each delay point. The degree to which this is practical, however, is limited by the read-out noise of the detector.

6.1. Practical constraints

The most important requirement for the output of the proposed source is that it must generate the correct pulses. For the DUV pulse, the desired characteristics are the same as in the case of driving at 790 nm, with one exception: because the driving pulse generated in the OPA is much more energetic, the conversion efficiency from the driver to the DUV pulse is less important. Even a decrease in the efficiency by one order of magnitude would lead to more energy in the DUV pulse than achieved so far. The pulse duration and tunability of the photon energy are needed just as before. Consequently, the requirements on vacuum propagation are also the same, meaning that a negative pressure gradient has to be used.

Further requirements are added by the intention of generating two short pulses simultaneously, which is to say that the driving pulse should remain short after the self-compression process. This means that soliton fission must not be allowed to occur. Moreover, significant stretching of the soliton after DWE has happened should also be avoided; with the dispersion of both the fibre and most optical materials being negative at SWIR wavelengths, re-compressing the pulse afterwards poses an added complication. Therefore, the point of maximum self-compression and DUV emission should be very close to the end of the fibre. This is in contrast to the process used in previous chapters, in which the DUV pulse was often created nearer the middle of the fibre.

Another parameter that has been ignored so far is the total power transmission of the apparatus including the waveguide. This is of paramount importance for the proposed source because of the extreme inefficiency of the HHG process used to generate the attosecond SXR pulse. Commonly, over 600 µJ of driving pulse energy is required with a pulse duration below 15 fs to access to the water window. Thankfully, the comparatively large pulse energy means that a fibre with a much larger core diameter can be used while still obtaining the necessary intensity for strong nonlinear effects. In the proposed scheme, a simple capillary of 400 µm core diameter is used. At this size, the propagation loss at the SWIR wavelength of 1750 nm is only 0.36 dB m^{-1}, lower

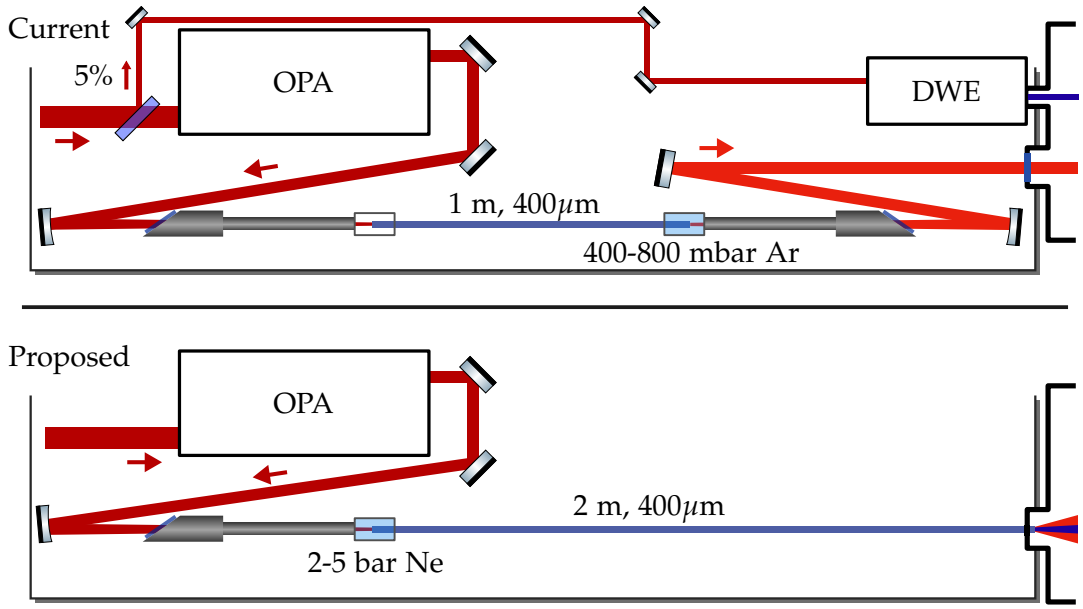

Figure 6.1.: Optical layout of the current (top) and proposed (bottom) experimental apparatus for the generation and combination of high harmonics and DUV pulses. Note that the experimental chamber (at the right edge of the image) is not the same as that containing the XFROG apparatus.

even than in most anti-resonant PCF.

The large core diameter reduces the anomalous dispersion of the waveguide when compared to the small anti-resonant PCF used in earlier chapters; the GVD of an evacuated 400 µm capillary at 1750 nm is only $-35\,\mathrm{fs^2\,m^{-1}}$, compared to around $-250\,\mathrm{fs^2\,m^{-1}}$ for a 45 µm PCF at 790 nm. Since the pulses generated in the OPA process are also comparatively long, around 35 fs, the dispersion length is very large, exceeding 10 m even in an evacuated capillary. Correspondingly, the soliton fission length for moderate soliton numbers is in the range of a few metres, meaning that a longer fibre is needed.

In practice, this cannot be scaled arbitrarily, however, for the simple reason that the optical table has a finite length. The table just outside the vacuum chamber containing the pump-probe interferometer is 3 m long. Although the current SWIR capillary is only 1 m long, it takes up nearly the entire table. This is because the beam at the

entrance and exit of the capillary is very small, and the intensity correspondingly high. The entrance and exit windows must therefore be placed sufficiently far away from the entrance and exit faces of the waveguide to avoid burning, resulting in a total length of 2 m for the apparatus including the gas cells on either side of the capillary. In addition, the focusing and collimating mirrors are around 50 cm away from the entrance and exit windows, respectively, so that the angle of incidence on the spherical optics can be minimised to avoid astigmatism (see fig. 6.1).

For a capillary directly coupled to the vacuum chamber, the added length on the exit side of the waveguide can be omitted, since the re-collimation occurs in the vacuum chamber. This puts the maximum waveguide length achievable without a major re-design of the laboratory at around 2 m. One option to increase this further would be to let the capillary extend into the vacuum chamber for some distance, however given the space constraints inside it, this may prove very difficult in practice.

Apart from lower cost and easier procurement, the use of a capillary as the waveguide additionally eases the requirements on the input beam profile. As discussed in section 4.3, a beam being coupled into a PCF needs to be almost entirely free of distortions in order to avoid damage to the delicate cladding structure. This would add significant experimental complication and introduce additional losses to the fibre-coupling of the SWIR driving pulse, since the profile of the beam generated in the OPA is very poor (see fig. 6.2a), requiring aggressive spatial filtering. Given the high pulse energy, nonlinear effects in air are likely to occur in the focal plane of the spatial filter, so that the filtering would have to be done in vacuum. When using a capillary, the extra stage of filtering can be avoided, since it acts as a spatial filter itself.

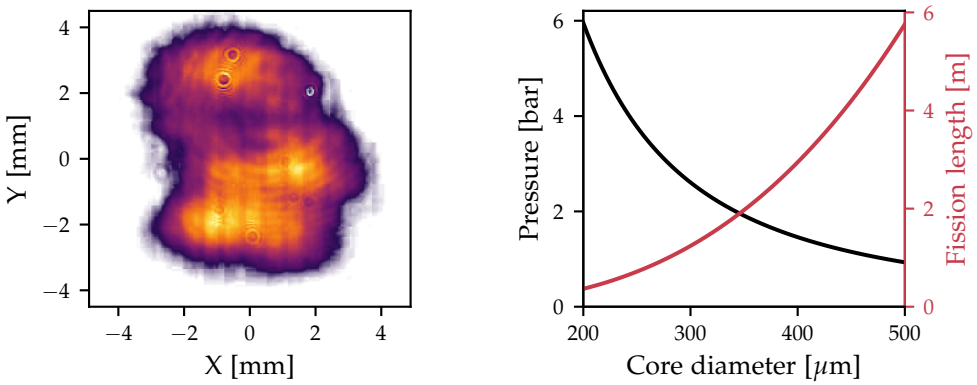

(a) (b)

Figure 6.2.: (a) The far-field beam profile of the 1750 nm pulse generated in the OPA. The small features in the top part of the beam are caused by dust on the attenuators in front of the camera, and the vertical lines are artefacts from the camera chip.
(b) The (static-fill) phase-matching pressure in neon (left axis, black) for DWE at 250 nm when driven by 700 µJ energy in a 35 fs driving pulse at 1750 nm for a variety of capillary core diameters, and the resulting fission length (right axis, red).

6.2. Numerical results

To establish that the simultaneous generation of both a driver for HHG and a DUV pump pulse is feasible both in principle and in practice, I have performed numerical simulations of an experimental set-up that could be built in our laboratory. To establish realistic parameters, the current SWIR capillary pulse compression apparatus was taken as a guide.

Since the input beam profile is very poor (see fig. 6.2a), the low coupling efficiency at the capillary entrance significantly reduces the amount of energy that actually drives nonlinear processes in the fibre. Therefore, the total pulse energy before coupling is not an appropriate value to take for the simulations. Instead, the transmitted energy

can be used. At present, a typical value for the pulse energy transmitted through the evacuated 400 µm capillary is 650 µJ. By measuring the spatio-temporal profile of the beam after the capillary, it has previously been established that most of the energy is indeed propagating in the fundamental mode [154]. Using the propagation loss from the Marcatili model as in eq. (3.1.8), this corresponds to an initial pulse energy in the fundamental mode of around 770 µJ. One consequence of this is that a conversion efficiency to the DUV of less than 0.2 % would be sufficient to obtain µJ-level pulses.

Within the analytical model for phase-matching to the dispersive wave introduced in section 3.4, the choice of the waveguide diameter is largely arbitrary because of the scaling of the various parameters. A smaller core will lead to higher loss per unit length, however the intensity is higher and the anomalous waveguide dispersion stronger, so that both the dispersion and nonlinear lengths are shorter. Correspondingly, the soliton fission length, which can be taken as a length scale for the DWE process, is also shorter, so that a shorter waveguide can be used, decreasing the overall loss. This implies that by choosing a smaller, shorter waveguide, the same overall dynamics can be achieved while making the apparatus more compact.

However, this analysis ignores both photoionisation effects and the coupling into the waveguide [23]. As mentioned, the intended application for the light source demands that a negative pressure gradient be used. Therefore, the presence of a nonlinear medium at the fibre entrance has to be taken into account. Both the self-focusing induced by the Kerr effect and the plasma defocusing will influence the coupling efficiency, especially since the spatial filtering effect of the capillary has not yet taken place, so that the intensity distribution is not perfectly uniform and the pulse energy is high. In order to avoid these issues, it is prudent to simply increase the core diameter to the largest practical size to reduce the intensity at the entrance face.

With the waveguide length limited to around 2 m, the problem of finding the biggest feasible core size reduces to making sure that the fission length under the conditions required for DUV generation is around this value. This is complicated somewhat by the presence of the pressure gradient, however the core size obtained for a static fill can

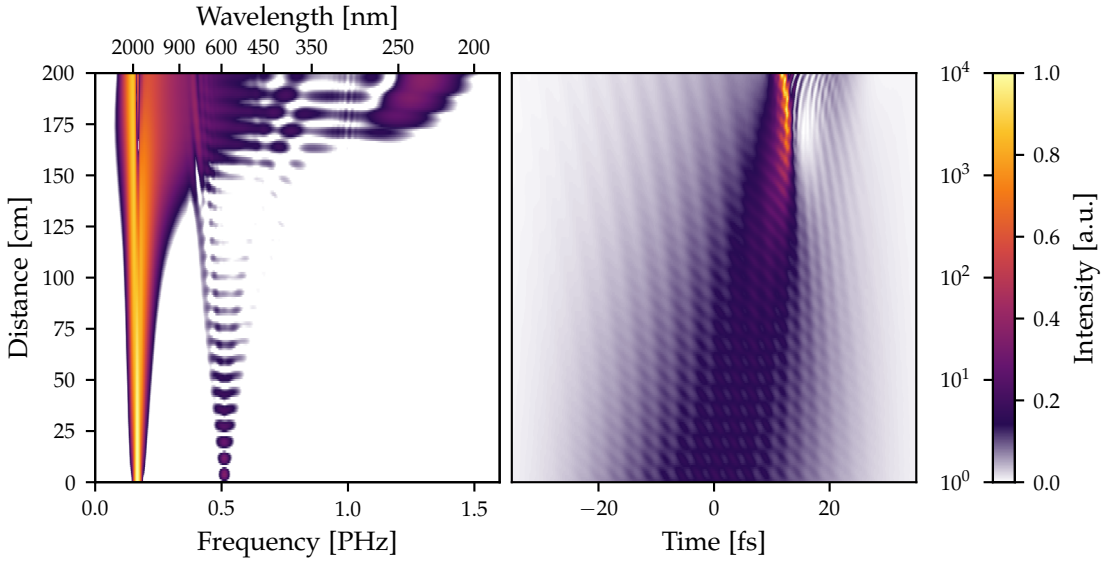

Figure 6.3.: Spectral (left) and temporal (right) evolution of the 35 fs SWIR driving pulse discussed in the text in a 2 m long capillary with a core diameter of 400 µm filled with neon using a negative pressure gradient from 3.4 bar to vacuum. The strong fringes in the time-domain plot are created by the third harmonic of the driving pulse, visible around 600 nm in the spectral-domain plot. Note that the logarithmic colour scale for the frequency domain plot spans four orders of magnitude.

serve as a guide for a numerical exploration of the parameter space.

Figure 6.2b shows that for phase-matching at 250 nm, a fission length of just under 2 m is obtained for a core diameter of about 350 µm. To obtain similar dynamics with a pressure gradient, the initial pressure needs to be higher, suggesting that the core diameter should be slightly larger than this value. Therefore, a core diameter of 400 µm is indeed reasonable.

One example of a soliton self-compression and DUV emission process in a practically feasible capillary waveguide is shown in fig. 6.3. The simulation in this case includes the loss of the capillary as shown in eq. (3.1.8). The typical features associated with the appearance of the dispersive wave are present, including the dramatic spectral

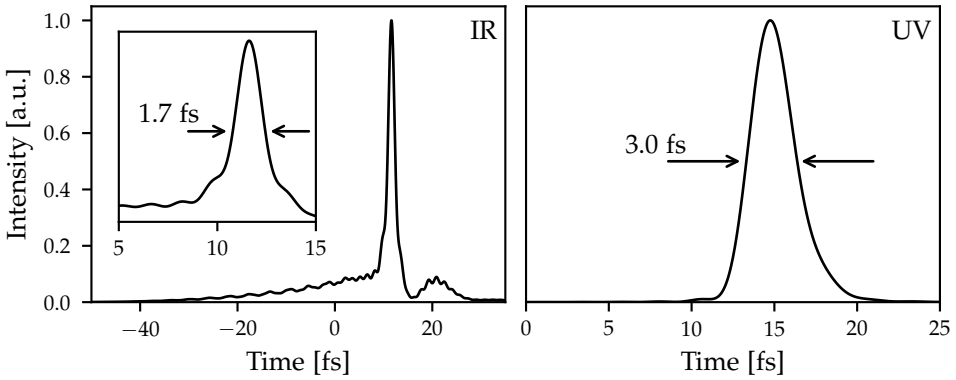

Figure 6.4.: Infrared (left) and DUV (right) pulses at the end of the fibre from the simulation shown in fig. 6.3. The bandpass windows are 400 nm to 5000 nm for the infrared pulse and 190 nm to 250 nm for the DUV pulse. The inset in the left plot shows a zoomed-in region around the peak of the infrared pulse.

expansion and temporal compression of the driving pulse. Crucially, the point of maximum pulse compression is far enough towards the end of the fibre to avoid soliton fission, and the driving pulse stays intact.

The dispersive wave appears at 225 nm in this example. The conversion efficiency to the DUV is 3.6 %, similar to what was achieved when driving with a pulse at 790 nm. However, due to the much larger input energy, the DUV pulse contains over 25 µJ of energy. The pressure at the entrance side of the fibre is 3.4 bar, but at the point of dispersive wave emission it has dropped to 1 bar. This is close to the phase-matching pressure for a dispersive wave at 225 nm driven by this input pulse as calculated using the GNLSE model, which is 1.2 bar.

Power stability

Because the pulse driving the self-compression process in the proposed scheme is itself derived from the laser output using a nonlinear frequency conversion process, the input pulse energy fluctuates more than that of the laser output itself. As already mentioned,

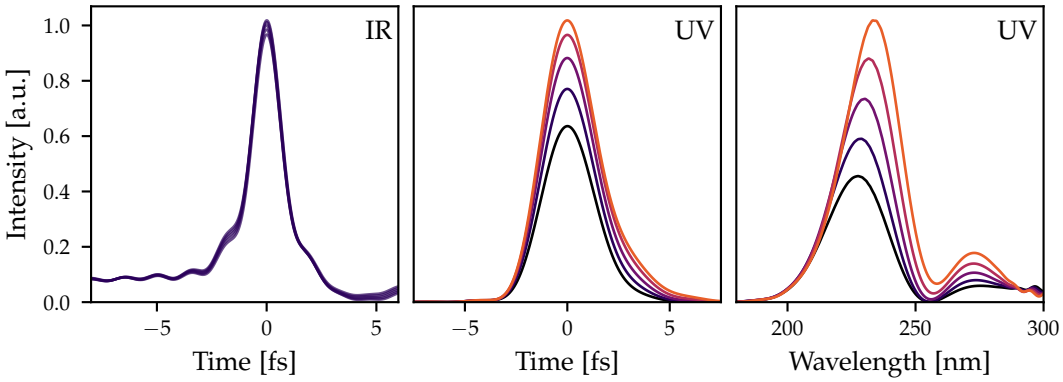

Figure 6.5.: Infrared (left) and DUV (centre) pulses as in fig. 6.4 for 5 values for the input energy between 683 µJ and 718 µJ, representing the extremes of a 5 % fluctuation in the input pulse energy around the nominal value of 700 µJ. The pulses were overlaid by shifting their peak to time $t = 0$. Right: spectrum of the DUV pulse for the same range of input energies.

the beam path inside the OPA system is also very long, so that the OPA output energy and any pointing instabilities in the laser beam are coupled through the alignment sensitivity of the apparatus. A conservative (i.e., pessimistic) estimate of the pulse energy fluctuation of the system is around 5 %.

It is thus important to consider the effect of these energy fluctuations on the self-compression process. Should a small increase of energy lead to soliton fission, for instance, or a small decrease to the disappearance of the DUV pulse, the requirements on the laser stability may be impractically stringent. These particular issues are not present when using a capillary for only spectral broadening; in this case, energy fluctuations couple to the pulse duration after compression, but do not usually lead to pulse breakup.

As shown in fig. 6.5, even the more demanding self-compression process is not affected too strongly by input energy fluctuations. A 5 % change around the mean energy leaves the shapes of both the infrared and DUV pulses essentially unchanged, with neither pulse duration changing by more than 5 %. However, the pulse energy

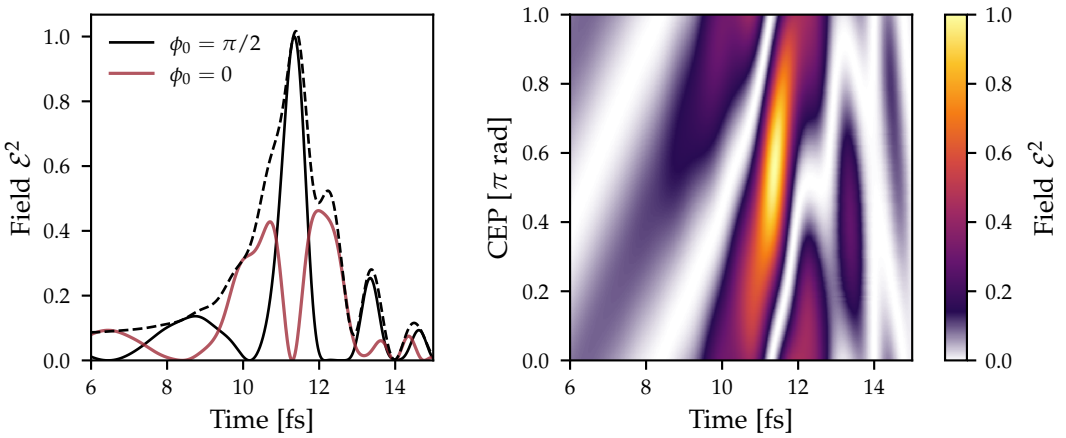

Figure 6.6.: Dependence of the electric field, rather than intensity, profile of the self-compressed infrared pulse on the CEP of the input pulse, shown at the exit of the fibre.
Left: field profiles at the CEP value resulting in the highest peak field ($\pi/2$) and the one exactly $\pi/2$ out of phase from it. The black dashed line shows the intensity profile at $\phi_0 = 0$.
Right: Evolution of the field profile with CEP.

in the DUV changes dramatically, with an overall fluctuation of about 50 % around the mean. Furthermore, the central wavelength of the DUV pulse shifts a little, with the centre of mass of the main peak changing by ca. 3 nm at the extremes. This is a consequence of the fact that the DUV emission occurs at very slightly different pressures when the input energy is changed as well as the small change in the nonlinear contribution to the phase-matching.

The frequency shift is likely inconsequential when the source is used in an experiment, since it is very small compared to the bandwidth of the pulse. The energy fluctuation is more worrisome, however, and reducing any instabilities present in the laser and OPA output as much as possible is therefore very important.

Carrier-envelope phase

As explained in section 2.3, a simple, if not easy, way of generating isolated attosecond pulses with HHG is to use a few-cycle pulse to drive the process. However, this places additional demands on the pulse as compared to the case of a multi-cycle driving field. As the pulse duration approaches the cycle time of the field, the CEP of the pulse becomes more and more important. If the peak of the carrier wave under the intensity envelope does not coincide with the peak of the envelope itself, the maximum field strength is reduced, leading to weaker harmonic emission. For a pulse envelope that is symmetric around its peak, like for instance a Gaussian pulse, the extreme case is that the maximum field strength occurs at two points in the pulse envelope, so that two attosecond pulses at lower photon energy are created.

The self-compressed pulse is extremely short, with an intensity envelope several times shorter than the cycle time of the input pulse (5.8 fs). At this pulse duration, the definition of the CEP becomes somewhat less useful, given that the carrier field is strongly reshaped by the envelope and non-sinusoidal. However, the CEP of the multi-cycle input pulse still has a dramatic effect on the output pulse. In fact, because the self-compression dynamics are sensitive to the electric field itself rather than purely the intensity envelope, the shape of the self-compressed intensity envelope also changes as the input CEP varies. This is in sharp contrast to the case of pure spectral expansion without self-compression, where the pulse is not strongly reshaped as it propagates, and the sub-cycle effects of strong-field ionisation play a much smaller role.

Figure 6.6 shows the effect of varying the input pulse CEP on the output pulse. The pulse splits from a single strong spike into two cycles of lower, near-equal strength as the phase is changed by $\pi/2$, in a similar way to a Gaussian pulse. However, the evolution of the pulse between these points, shown in the right-hand part of the figure, is more complicated, since the intensity envelope also changes shape – note that the envelope shown as a dashed line in the left-hand part of fig. 6.6 differs from that shown in fig. 6.4.

This variation of the envelope as well as the carrier has potentially important con-

sequences for the use of the proposed source to generate attosecond pulses. The numerical results show that, similar to traditional capillary-based pulse compression schemes, the CEP of the driving pulse must be stabilised to reliably generate single attosecond pulses. In the OPA system used in our laboratory to down-convert the 790 nm laser pulses to the SWIR, the CEP of the pulse at 1750 nm is passively stabilised by the frequency conversion process: since the 1750 nm pulse is the idler in the OPA process, its phase is determined by the difference between the phases of the pump (at 790 nm) and the signal (at 1480 nm). With both of the latter derived from the same laser pulse, this phase difference is constant except for small fluctuations in the delay between the pulses [155]. The fluctuations are measured and compensated for by adjusting the length of the signal-pump interferometer in the parametric amplifier in a feedback loop.

Stability of the CEP does not mean that it has the correct value at the point of attosecond pulse generation, however. Rather, the CEP needs to be scanned in order to find the optimum condition. This is partly because most measurements of the CEP can only determine phase differences rather than an absolute value of the CEP. Furthermore, the phase and group velocity of the pulse are not equal except in vacuum, so that the CEP in the harmonic generation target is determined by the precise optical path length to the target. Since the stabilisation is done in the OPA system, any tuning of the CEP is also most easily carried out there by adding an offset to the stabilisation feedback loop. For self-compression of the infrared pulse, however, this tuning affects not only the CEP but also the pulse shape, so that the combination of optimal pulse shape and CEP may be impossible to attain. This problem can be overcome by tuning the CEP after the fibre with the insertion of glass wedges. By using sufficiently thin wedges manufactured from a material that is weakly dispersive at the SWIR wavelength, such as CaF_2 (GVD of $-8.5\,\mathrm{fs}^2\,\mathrm{mm}^{-1}$ at 1750 nm), the added dispersion can be kept to a minimum.

6.3. Potential for attosecond pulse generation

To ascertain whether the self-compressed pulse discussed in the previous section is indeed capable of generating SXR harmonics, I have calculated the HHG spectrum resulting from focusing this pulse into a gas target. To do this, I used the strong-field approximation (SFA), the most common numerical model for HHG [31, 49].

The strong-field approximation

The SFA takes its name from the principal assumption of the model, namely that the laser field applied to the atom is very strong. Additional assumptions are the *single active electron* approximation, which means that only the outermost valence electron of the atom can be ionised, and the *dipole* approximation, which means that the electric field varies only with time and not space. The latter is equivalent to the assumption that even the furthest excursion of any electron is much shorter than the driving wavelength. In combination, these approximations lead to a model in which only the ground state and the field-dressed continuum (i.e., free-electron) states of a single electron have to be considered, with the electric field driving transitions between the two. Bound excited states are ignored.

The quantity of interest is the time-dependent dipole moment of a single atom-electron system,

$$\mathbf{D}(t) = \langle \Psi | \mathbf{r} | \Psi \rangle \ , \tag{6.3.1}$$

where $\mathbf{r}$ is the spatial coordinate and $|\Psi\rangle$ is the state of the electron. This is because the dipole moment is directly related to the nonlinear polarisation $\boldsymbol{\mathcal{P}}_{\mathrm{NL}}$ through the density of atoms subjected to the laser field. Therefore, the Fourier transform of the dipole moment yields the spectrum of the nonlinear response, which is to say the HHG spectrum. The utility of the SFA is that it allows for the calculation of $\mathbf{D}(t)$ without solving the time-dependent Schrödinger equation itself, decreasing the required computation time by orders of magnitude. This means that varying parameters such as the pulse duration, pulse energy or gas pressure becomes feasible.

Two further approximations are necessary to make the calculation of $\mathbf{D}(t)$ in the SFA numerically efficient. The first is that no continuum-continuum transitions are considered, which means that any scattering between a returning electron and the Coulomb potential of the ion is ignored. The second is to carry out the resulting integral over momentum using a saddle-point approximation, which simply ignores those parts of the integral where rapid oscillations in the phase lead to an average value of 0. With these approximations, the dipole can be calculated as

$$\mathbf{D}(t) = -i \int_{-\infty}^{t} \mathrm{d}t_\mathrm{b} \left(\frac{2\pi}{\tau}\right)^{3/2} g^*(t_\mathrm{b})\, \mathbf{d}^*[\mathbf{p}_\mathrm{st} + \mathbf{A}(t_\mathrm{b})]\boldsymbol{\mathcal{E}}(t_\mathrm{b})\, g(t)\, \mathbf{d}[\mathbf{p}_\mathrm{st} + \mathbf{A}(t)]\mathrm{e}^{-iS(\mathbf{p}_\mathrm{st}, t, t_\mathrm{b})} + \mathrm{c.c.}\,, \tag{6.3.2}$$

where t_b can be interpreted as the birth time of the electron and consequently $\tau = t - t_\mathrm{b}$ is the excursion time, $g(t)$ is the population of the ground state, $\mathbf{d}(\mathbf{p}) = \langle g|\mathbf{r}|\mathbf{p}\rangle$ is the dipole moment between the ground state and the continuum, $\mathbf{A}(t)$ is the vector potential of the field, $\boldsymbol{\mathcal{E}}(t)$ is the field itself, $S(\mathbf{p}, t, t_\mathrm{b})$ is the *action*,

$$S(\mathbf{p}, t, t_\mathrm{b}) = \int_{t_\mathrm{b}}^{t} \mathrm{d}t' \left(I_\mathrm{p} - \frac{1}{2}\left[\mathbf{p} + \mathbf{A}(t')\right]^2\right), \tag{6.3.3}$$

and $\mathbf{p}_\mathrm{st}$ is the stationary momentum (the momentum at which the action is stationary):

$$\mathbf{p}_\mathrm{st} = \frac{-1}{\tau} \int_{t_\mathrm{b}}^{t} \mathrm{d}t' \mathbf{A}(t')\,. \tag{6.3.4}$$

To calculate $\mathbf{D}(t)$ numerically given a field $\boldsymbol{\mathcal{E}}(t)$, the lower limit of the integral over birth times in eq. (6.3.2) has to be changed to a finite time. By choosing the limit such that the entire laser pulse is contained in the window, no information is lost. Furthermore, by assuming a linearly polarised field, we can use scalar quantities.

The SFA is essentially the quantum analogue of the semi-classical three-step model. Where the three-step model treats the trajectory of the electron outside the ion purely with classical mechanics [29], in the SFA this is encoded in the action. Furthermore, the SFA involves no "recombination" of the electron with the ion as a defined event. Instead, harmonic emission is the result of the interference between the continuum electron wavepacket and the ground state. The SFA can include the effect of arbitrary atomic potentials through the dipole moment d, and arbitrary pulse profiles through

the field $\mathcal{E}$. Furthermore, while it ignores the possibility of ion-electron scattering, it does include higher-order returns [31]; in the classical picture these correspond to electrons which pass the atom one or more times before recombination.

As with all nonlinear frequency conversion schemes, phase-matching is very important for the efficient generation of high harmonics. In the SFA as presented in eq. (6.3.2), this is not taken into account at all. It is also not possible to derive a simple phase-matching efficiency curve such as those for low-order processes [30, 156]. A complete treatment involves calculating the dipole for each plane of a propagating laser pulse and adding up the contributions, which is orders of magnitude more computationally expensive than a single SFA calculation. However, some general phase-matching effects can be approximated during the SFA calculation itself.

For a sinusoidal laser field, there are two combinations of birth and return times in each half-cycle of the field which result in the same photon energy, except for the very highest energy. These pairs are usually separated into the *long* trajectories and *short* trajectories, based on which trajectory has the longer excursion time. The long trajectories are commonly much weaker than the short ones due to the larger amount of electron wavepacket diffusion, and furthermore phase-matching for the long trajectories is more difficult to achieve [30, 156]. In addition, they are emitted with a much larger divergence and thus hard to detect [39]. It is therefore reasonable to remove the contributions of the long trajectories. This can be done by attenuating any components with excursion times longer than one half-cycle, since all short trajectories have shorter excursion times. So as not to add additional frequency content to the dipole, it is best to use a smoothly decaying apodisation function, such as the error function.

The apodisation has the further benefit of reducing the computation time required for the SFA integral. Without it, the integral for return times near the end of the time window extends over a very large range in possible birth times, making computation slow. The apodisation means that a smaller window extending to only around one cycle of the field can be considered.

The spectrum of the simulated driving field is very broadband and the spectral expansion during self-compression is asymmetric due to self-steepening, so the carrier frequency of the compressed pulse is not the same as that of the initial pulse generated by the OPA; it furthermore changes slightly over the course of the pulse, partly due to the strong reshaping of the carrier by the sub-cycle envelope. The apodisation of the long trajectories was done based on a cycle time of 5.5 fs.

There are two quantities in the SFA integral eq. (6.3.2) which are not immediately given by the SFA itself. The first is the dipole moment d between the ground state and the continuum. Values for this as well as approximate analytical expressions are available in the literature [157]. The second quantity is the ground-state population $g(t)$. This can be calculated using one of the approximate ionisation rates introduced in section 2.2.5. Through $g(t)$, the SFA includes the effects of ground-state depletion; in a very strongly ionised medium, $g(t)$ is small, so that the dipole $D(t)$ is drastically reduced, and harmonic generation ceases. Here, ground-state depletion was included by calculating $g(t)$ with the ADK ionisation rate.

It is important to note that since only a single electron is considered and the bound-state structure of the atom is largely ignored, the SFA cannot give accurate results for low-order harmonics of the laser field. Another way to view this is that the semi-classical analogue, the three-step model, clearly does not apply where the response of the atom is due to bound charges, because only a liberated electron is considered. Furthermore, while the shape and frequency extent of the harmonic spectrum calculated by the SFA is a good guide as to what to expect in an experiment, the harmonic yield or conversion efficiency is not, with the SFA commonly overestimating the yield by around two orders of magnitude [158].

Results

Figure 6.7 shows the harmonic spectrum as a function of input CEP for the driving pulse shown in figs. 6.3 and 6.4 when focused into a helium gas target. Only the on-axis field in the focal plane of an aberration-free Gaussian beam was considered. To obtain

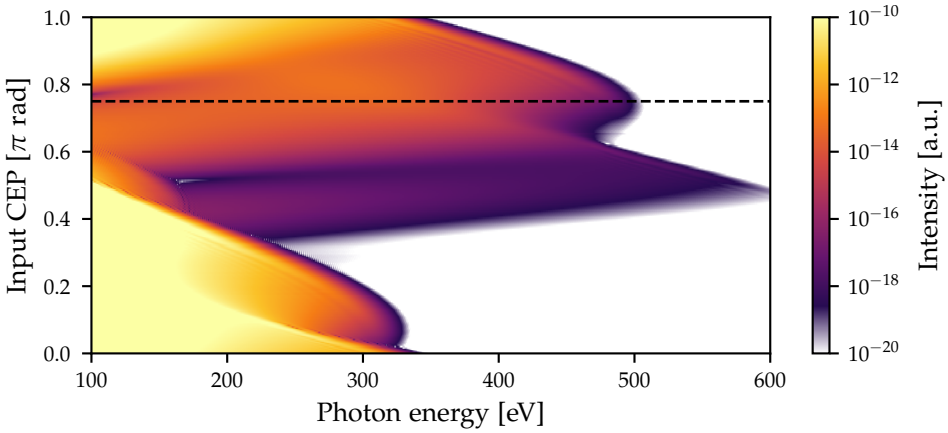

Figure 6.7.: CEP-dependent high-harmonic spectra calculated using the SFA for the self-compressed pulse described in the previous section, focused into a helium target with a magnification of $M = 0.5$, resulting in a focal spot diameter of 128 µm. Phase-matching of the short trajectories was approximated by temporal apodisation of birth times more than one half-cycle before the return time, assuming a cycle time of 5.5 fs. The dashed black line shows the CEP value for which the spectrogram is shown in fig. 6.8.

the driving field amplitude in the medium, the scaling inherent in the single-mode FME was undone and the transverse field profile of the pulse modelled as a Gaussian beam with a waist given by the value for optimum coupling efficiency, $w_0 = 0.64a$ with a the fibre core radius. Collimation and refocusing was included by re-scaling the field strength according to the magnification of the imaging system; here, $M = 0.5$ was taken to include collimation with an $f = 1$ m optic and re-focusing with an $f = 50$ cm optic.

The SFA results show that the pulse is capable of generating high harmonics across all of the water window and beyond, with the highest cutoff photon energy at around 600 eV. This is despite the relatively loose focusing with a focal spot size of 128 µm, which is more than twice the focal diameter used in recent work on SXR HHG in our laboratory. The difference arises from several factors.

Firstly, the model presented here does not include any losses or dispersion introduced

to the pulse during the transport from the capillary to the gas target. Given that the pulse is directly delivered to a vacuum chamber in the proposed apparatus, ignoring dispersion is not a very significant approximation, however the reflectivity of the beam transport optics would have to be taken into account for a more realistic model. Energy loss would make tighter focusing necessary to achieve the same intensity, however in this simple zero-dimensional model nothing else would change.

Secondly, the pulse used here is significantly shorter than that used in our previous work, where the driving pulse could not be compressed much below 12 fs due to dispersion in the beam path. The driving pulse also likely had a strong satellite after the main peak, reducing the peak intensity further [53, 154]; this is not the case for the pulse in the simulations. It should be noted that although the pulse duration is only a fraction of the cycle time (5.8 fs at 1750 nm), this does not lead to a significant reduction in the ponderomotive energy; the full width of the electric field transient itself is ca. 2.25 fs, corresponding almost exactly to a single half-cycle of the long-wavelength driving field. The pulse thus has a near-perfect shape for the generation of isolated attosecond pulses, with little ionisation or harmonic emission occurring before and after the main peak and a sufficiently long cycle time for high photon energies.

The third factor leading to the very high photon energies as compared to previous experiments is the most important, namely the lack of macroscopic effects. As mentioned before, this single-point SFA calculation does not include the effects of phase-matching beyond removal of the long trajectories. A related issue is that the reshaping of the beam by the Kerr effect and plasma defocusing in the gas target is not included. Numerical modelling of our recent experimental results showed that this is a very important effect when driving HHG with long-wavelength pulses, partly because the refractive index contribution from the plasma increases for longer wavelengths (see eq. (2.2.38)).

The highest ionisation fraction in the simulations shown here occurs at a CEP value of $\pi/2$ and reaches as high as 9 %. As a result, the peak intensity in the target will likely not reach the value assumed in this simple model, reducing the peak photon

energy. The plasma defocusing effect should be somewhat weaker than in previous experiments, since only the main peak of the pulse contributes significantly. However, the fact that the gas target does not have hard boundaries and so some gas will be present for an extended distance before the target means that defocusing will still occur in practice.

An interesting feature of the CEP-dependent harmonic spectra in fig. 6.7 is that while the highest photon energy occurs at a CEP value of $\pi/2$ as suggested by the field profiles in fig. 6.6, the highest overall flux in the water window is obtained at a value of ca. $0.75\,\pi$, indicated by the dashed black line in fig. 6.7. From these simulations it is not immediately clear why this is so, however it is likely related to the detailed interplay between the CEP-dependent shape of the carrier and pulse envelope.

The absence of any discrete harmonics in the spectra already indicates that the SXR radiation is generated as an isolated pulse. To establish that this pulse is indeed of attosecond duration and to obtain the bandwidth over which a single pulse can be isolated, it is useful to calculate the spectrogram of the SFA dipole. The spectrogram is created using a Gaussian gate and is therefore calculated as

$$S(\omega, t) = \mathcal{F}_{t'}\left[D(t')\,\mathrm{e}^{-[(t'-t)/\Delta t']^2}\right], \tag{6.3.5}$$

where $\Delta t'$ is the $1/\mathrm{e}$ half width of the gate, chosen here as 100 as.

The spectrogram of the SFA dipole with an input CEP of 0.75π is shown in fig. 6.8. The fact that energies above 200 eV are only created for the strongest peak of the pulse indicates that they are emitted as an isolated attosecond burst. Although the transform-limited pulse duration of this spectrum is below 20 as, the real pulse duration even in this simulation is much longer, around 200 as. The cause for this can be seen in the spectrogram: different photon energies are emitted at different times due to the dependence of the photon energy on the excursion time. This inherent dispersion is known as *attochirp*. It can be compensated to an extent by passing the harmonic beam through thin metal films, which are transmissive as well as anomalously dispersive for certain regions of the SXR photon energy range [159].

In an experiment, the limited bandwidth of high-harmonic phase-matching as well as

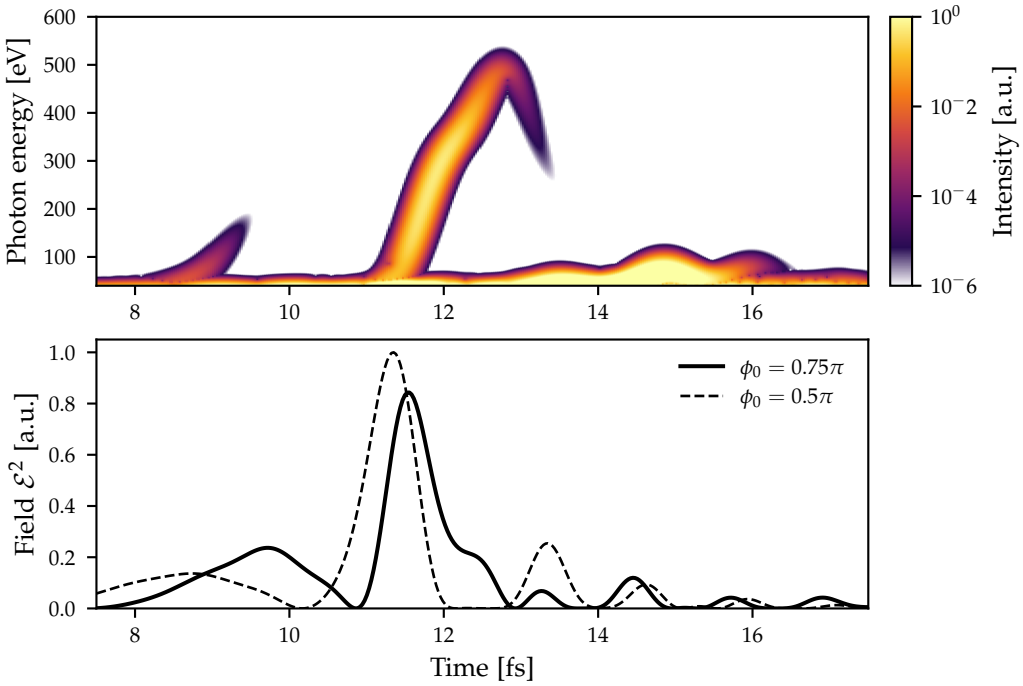

Figure 6.8.: Top: spectrogram of the dipole $D(t)$ as obtained by the SFA for an input CEP value of 0.75π. The window function is a Gaussian with a $1/e$ half width of 100 as.
Bottom: squared electric field $\mathcal{E}(t)^2$ of the driving pulse. The dashed line shows the field profile for an input CEP of 0.5π

the dispersion of the medium would likely further limit the pulse duration. However, a sub-femtosecond pulse should be easily obtainable. The source proposed here is thus capable in principle of producing both the pump and probe pulses for a transient absorption experiment with an unprecedented time resolution of around 3 fs using only a single pulse compression stage.

7. Conclusions and Outlook

Progress in ultrafast science has depended on the development of new light sources for the entirety of the field's relatively short history. As the focus shifts from atoms to molecules and from gas-phase experiments to solid and liquid targets, the availability and characterisation of extremely short pulses at new wavelengths is one of the dominant issues. In many areas of light source development, the state of the art is being overturned by the application of novel micro-structured waveguides and their unique capabilities.

In this thesis I have presented work towards harnessing these capabilities to generate ultrashort pulses across the deep ultraviolet for time-resolved experiments, with a particular focus on applying it to soft X-ray transient absorption spectroscopy. In addition, I have demonstrated that the phenomenon of soliton self-compression, which forms the basis of this pulse generation mechanism, can be used for both parts of the experiment simultaneously in a single fibre, greatly reducing complexity and improving both the reliability and performance of the apparatus.

In chapter 4 I described the design, modelling and implementation of an experimental apparatus to use soliton self-compression and dispersive wave emission as a source of tunable ultrashort pulses in the deep ultraviolet. To increase the conversion efficiency, the driving field is pre-compressed by a factor of 2 in a hollow-capillary pulse compressor. The design of the apparatus places particular emphasis on fulfilling the stability requirements for long data acquisition times and the need to protect the photonic crystal fibre from damage. While not record-breaking, the energy conversion efficiency is already sufficient for many experiments and can likely be improved further. Delivery of the generated pulses to an experiment while avoiding pulse distortions

and stretching poses a significant challenge, and one that has remained somewhat unexamined until now. To address this issue, the simplest and most effective solution is to connect the light source directly to a vacuum chamber, which results in a pressure gradient in the waveguide. This poses additional complications both in a numerical model and in the experiment, however the utility of the source remains unchanged.

To verify that dispersive wave emission can indeed be used as a source of ultrashort pulses in a real experiment, in chapter 5 I presented the characterisation of deep ultraviolet pulses using ptychographic methods under conditions which closely match those of a future time-resolved experiment. The measurements show that direct-to-vacuum coupling of the generation apparatus provides essentially dispersion-free transport of the pulse from generation to target, so that very short pulses can be delivered. This remains true when the wavelength of the pulse is changed, indicating that the full tuning capabilities of the dispersive wave can be utilised without significant additional complications.

Unfortunately, the fact that a sufficiently thin nonlinear medium was not available precludes definitive statements about the pulse shape for the moment, so that a systematic investigation of the optimal parameters for the shortest possible pulses will have to wait. However, the data so far strongly suggests that durations significantly below 5 fs are achievable in this system while retaining tunability. Once the appropriate nonlinear medium is available, a higher-quality measurement will be straightforward.

In the numerical model for the generation of ultraviolet pulses, a particular scenario appeared in which both the driving pulse and the dispersive wave exit the fibre as high-quality few-femtosecond pulses. In chapter 6, I explored this phenomenon further with a view to apply it to the generation of soft X-ray attosecond pulses. This would harness both the self-compression and secondary pulse generation capabilities. From the simulations, the great potential of this kind of source is easily apparent. The creation of high-energy sub-cycle pulses alone makes the self-compression technique an attractive option for harmonic generation, especially because most of the pulse energy is contained in a single peak and can contribute to X-ray emission. The calculated

single-atom harmonic spectrum extends across the water window and beyond even with relatively loose focusing. In combination with the simultaneous creation of a bright second pulse in the ultraviolet, the proposed source is uniquely capable of driving a time-resolved soft X-ray experiment using existing laser sources.

To make this proposal a reality, more work is required on the experimental implementation. For instance, the phase-matching conditions for harmonic generation have not been addressed thus far. This is partly due to the fact that phase-matching of soft X-ray harmonics is poorly understood, and so far no consensus has emerged on routes towards source optimisation. Therefore, only sophisticated and resource-intensive numerical modelling or experiments can give more insight into how the brightness of high-harmonic soft X-ray sources can be increased. Given the extensive expertise and previous work on this topic in the Laser Consortium, this laboratory is very well-placed to implement and test the proposed source and use it in experiments.

While this work focused on the deep ultraviolet and short-wavelength infrared parts of the electromagnetic spectrum, all techniques presented here can be applied to a number of combinations of driving and dispersive wave frequency. Especially for the dispersive wave, this potential extends in both directions. The generation of ultrashort pulses in the vacuum ultraviolet is more challenging due to detection issues, but there are no fundamental restrictions that prevent it. On the other hand, dispersive wave generation across nearly all of the visible spectral range has already been demonstrated.

There are two main avenues for future work based on the results acquired and challenges encountered during the experimental work for this thesis. Initially, the first of these consists simply of repeating the pulse characterisation measurements in the deep ultraviolet with an appropriate nonlinear medium to establish the true pulse duration achieved with the source presented here. The next step will be to explore in more detail how the various experimental parameters influence the pulse duration in the dispersive wave, and consequently how the duration can be improved further towards the sub-femtosecond regime. If the bandwidth can be improved sufficiently, this could provide a route towards true attosecond-resolution experiments.

Depending on the dynamics being investigated and the relevant absorption resonances, two dispersive wave sources may be used instead of a probe based on high-harmonic generation, increasing the total pulse energy available for the experiment by several orders of magnitude.

The second avenue for future work is to apply the source that is already available to experiments with few-femtosecond resolution, such as a soft X-ray transient absorption study. Work towards this goal is ongoing in our laboratory, and the availability of tunable pump pulses for such an experiment opens up a wide variety of possible samples. Depending on the molecular system, this may require more pump pulse energy than has been achieved so far. Improving on the pulse energy of around 250 nJ reported in chapter 4 can be done in two ways, by increasing the conversion efficiency or by using more energy to drive the process. The first will consist mostly of an exploration of the experimental parameter space, however this will have to be done in conjunction with pulse characterisation measurements to establish whether efficiency and bandwidth have to be traded off in a similar way to more conventional frequency conversion schemes. The second is mostly dependent on the availability of applicable waveguides with larger core diameters, since a large increase in intensity would require even shorter lengths of fibre than shown in section 5.4, which may not be practical.

The work in this thesis establishes soliton self-compression dispersive wave emission as a versatile, reliable and practical source of ultrashort pulses for time-resolved experiments. The exploration of the breadth of possible applications has only just begun.

A. Detailed derivations

Temporal phase shift due to GVD

A unchirped Gaussian pulse in time, centred at ω_0,

$$E(t) = E_0 \, \mathrm{e}^{-\frac{1}{2}\left(\frac{t}{\tau}\right)^2} \mathrm{e}^{-i\omega_0 t} \,, \tag{A.1}$$

has a complex frequency spectrum given by

$$E(\omega) = \mathcal{F}_t[E(t)] = \sqrt{2\pi}\, E_0 \tau \, \mathrm{e}^{-\frac{1}{2}\tau^2(\omega-\omega_0)^2} \,. \tag{A.2}$$

After propagation through a dispersive medium of length L with a GVD at ω_0 given by β_2, this is modified:

$$E(\omega, L) = \sqrt{2\pi}\, E_0 \tau \, \mathrm{e}^{-\frac{1}{2}\tau^2(\omega-\omega_0)^2} \mathrm{e}^{i\frac{1}{2}\beta_2 L(\omega-\omega_0)^2} \,, \tag{A.3}$$

where we have ignored the phase and group delays β_0 and β_1, since we are only interested in the shape of the pulse. The form of the field in eq. (A.3) is a Gaussian with a complex variance of $(\tau^2 + i\beta_2)^{-1}$. The field in the time domain is therefore given by

$$E(t, L) = E_0' \, \mathrm{e}^{\frac{1}{2}(\tau^2 + i\beta_2 L)^{-1} t^2} \mathrm{e}^{-i\omega_0 t} = E_0' \exp\left[\frac{1}{2}\frac{\tau^2 - i\beta_2 L}{\tau^4 + \beta_2^2 L^2} t^2 - i\omega_0 t\right] . \tag{A.4}$$

Note that the peak field E_0' is also different to the unchirped case, however its value is irrelevant to the present discussion. The temporal phase shift due to GVD is then given by

$$\phi_L(t) = \frac{\beta_2 t^2 L}{\tau^4 + \beta_2^2 L^2} \,. \tag{A.5}$$

The single-mode forward Maxwell equation

The nonlinear polarisation *induced by* the field in the fundamental mode is given by (see eq. (2.2.5))

$$P_{\text{NL}}(r,\theta,z,t) = \frac{3}{4}\varepsilon_0\chi^{(3)}|E(r,\theta,z,t)|^2E(r,\theta,z,t) \tag{A.6}$$

$$= \frac{3}{4}\varepsilon_0\chi^{(3)}|E(z,t)|^2E(z,t)\,|e_0(r,\theta)|^2e_0(r,\theta)\,, \tag{A.7}$$

where we have neglected the THG term for simplicity[1], and $e_0(r,\theta)$ is the transverse field distribution of the fundamental mode. Since $|e_0(r,\theta)|^2e_0(r,\theta) \neq e_0(r,\theta)$, to find the part of this which radiates in the fundamental mode, we project it as in eq. (3.2.12):

$$P_{\text{NL}}^{(0)}(r,\theta,z,t) = \hat{e}_0(r,\theta)\int_S P_{\text{NL}}(r,\theta,z,t)\hat{e}_0^*(r,\theta)\,\mathrm{d}\mathbf{S}\,, \tag{A.8}$$

where the integration region S is the fibre core and $\mathrm{d}\mathbf{S} = r\mathrm{d}r\mathrm{d}\theta$. Note that this is expressed in the time domain since the nonlinear polarisation then takes a simple form, and we have retained the transverse spatial dimension by adding the factor $\hat{e}_0(r,\theta)$. To carry out the projection, we first need to normalise $e_0(r,\theta)$ to find $\hat{e}_0(r,\theta)$:

$$\hat{e}_0(r,\theta) = \left[\int_S |e_0(r,\theta)|^2\,\mathrm{d}\mathbf{S}\right]^{-\frac{1}{2}} e_0(r,\theta)\,. \tag{A.9}$$

Using this to express $\hat{e}_0(r,\theta)$ in terms of $e_0(r,\theta)$ in eqs. (A.7) and (A.8) yields

$$P_{\text{NL}}^{(0)}(r,\theta,z,t) = \frac{3}{4}\varepsilon_0\chi^{(3)}|E(z,t)|^2E(z,t)\,\frac{\int_S |e_0(r,\theta)|^2e_0(r,\theta)e_0^*(r,\theta)\,\mathrm{d}\mathbf{S}}{\int_S |e_0(r,\theta)|^2\,\mathrm{d}\mathbf{S}}\,e_0(r,\theta) \tag{A.10}$$

$$= \frac{3}{4}\varepsilon_0\chi^{(3)}|E(z,t)|^2E(z,t)\,\frac{\int_S |e_0(r,\theta)|^4\,\mathrm{d}\mathbf{S}}{\int_S |e_0(r,\theta)|^2\,\mathrm{d}\mathbf{S}}\,e_0(r,\theta) \tag{A.11}$$

$$= \frac{3}{4}\varepsilon_0\chi^{(3)}|E(z,t)|^2E(z,t)\,\Gamma^{-1}\,e_0(r,\theta)\,, \tag{A.12}$$

which defines the normalisation constant Γ:

$$\Gamma = \frac{\int_S |e_0(r,\theta)|^2\,\mathrm{d}\mathbf{S}}{\int_S |e_0(r,\theta)|^4\,\mathrm{d}\mathbf{S}}\,. \tag{A.13}$$

[1] Since $e(r,\theta)$ can be taken as real-valued, including the THG term in this analysis would lead to the same result.

If we now define a re-scaled field $\tilde{E}$ by

$$E \equiv \sqrt{\Gamma}\tilde{E}\,, \tag{A.14}$$

and substitute for E in the FME eq. (3.2.7) for the fundamental mode, assuming that the nonlinear polarisation is solely due to the Kerr effect, we arrive at

$$\sqrt{\Gamma}\partial_z\tilde{E}^{(0)}(r,\theta,z,\omega) = i\beta\sqrt{\Gamma}\tilde{E}^{(0)}(r,\theta,z,\omega) + \Gamma^{\frac{3}{2}}\Gamma^{-1}\frac{i\omega^2}{2\varepsilon_0 c^2\beta}P_{\mathrm{NL}}^{(0)}(r,\theta,z,\omega)\,, \tag{A.15}$$

where the factor Γ^{-1} in the nonlinear term comes from the projection, and the factor $\Gamma^{\frac{3}{2}}$ from the third power of $\tilde{E}$. Since all of the factors of Γ and the spatial distribution $e_0(r,\theta)$ cancel, we recover the FME in the same form:

$$\partial_z\tilde{E}^{(0)}(z,\omega) = i\beta\tilde{E}^{(0)}(z,\omega) + \frac{i\omega^2}{2\varepsilon_0 c^2\beta}P_{\mathrm{NL}}^{(0)}(z,\omega)\,, \tag{A.16}$$

and transforming coordinates into the moving frame leads to the single-mode FME eq. (3.2.14).

By integrating the square of the field E over the transverse surface, we calculate the total time-dependent power $W(t)$ in the fibre. We can also express this in terms of the re-scaled field $\tilde{E}$:

$$W(t) = \frac{c\varepsilon_0}{2}\int_S |E(r,\theta,z,t)|^2 \mathrm{d}\mathbf{S} \tag{A.17}$$

$$= \frac{c\varepsilon_0}{2}\int_S |E(z,t)|^2 |e_0(r,\theta)|^2 \mathrm{d}\mathbf{S} \tag{A.18}$$

$$= \frac{c\varepsilon_0}{2}\,\Gamma\int_S \left|\tilde{E}(z,t)\right|^2 |e_0(r,\theta)|^2 \mathrm{d}\mathbf{S} \tag{A.19}$$

$$= \frac{c\varepsilon_0}{2}\,\frac{\left(\int_S |e_0(r,\theta)|^2\,\mathrm{d}\mathbf{S}\right)^2}{\int_S |e_0(r,\theta)|^4\,\mathrm{d}\mathbf{S}}\left|\tilde{E}(z,t)\right|^2 \tag{A.20}$$

$$= \frac{c\varepsilon_0}{2}A_{\mathrm{eff}}\left|\tilde{E}(z,t)\right|^2, \tag{A.21}$$

which defines the effective area A_{eff}:

$$A_{\mathrm{eff}} = \frac{\left(\int_S |e_0(r,\theta)|^2\,\mathrm{d}\mathbf{S}\right)^2}{\int_S |e_0(r,\theta)|^4\,\mathrm{d}\mathbf{S}}\,. \tag{A.22}$$

The pre-conditioned forward Maxwell equation

The FME is a first-order ODE, the generic form of which is commonly written as

$$\frac{\partial y}{\partial t} = f(y,t)\,. \tag{A.23}$$

In the FME, t is the propagation coordinate z, $y(t)$ is the field $E(\omega,z)$ and $f(y,t)$ contains the linear propagation as well as the nonlinear effects. To better incorporate the (analytically solvable) linear part of the equation, we can write the generic equation instead as

$$\frac{\partial y}{\partial t} = \mathcal{L}(t)y(t) + f(y,t)\,, \tag{A.24}$$

where $\mathcal{L}$ is the linear part (in the FME this incorporates dispersion through $\beta(\omega,z)$, the loss through $\alpha(\omega,z)$, and the moving frame) and $f(y,t)$ now only includes the nonlinear part. We then define the pre-conditioned function $\bar{y}(t)$:

$$\bar{y}(t) = \exp\left\{-\int_0^t \mathcal{L}(t')\mathrm{d}t'\right\} y(t) = \mathrm{e}^{-L(t)}y(t)\,, \tag{A.25}$$

where we have also defined the linear propagator $L(t)$ for brevity:

$$L(t) \equiv \int_0^t \mathcal{L}(t')\mathrm{d}t'\,. \tag{A.26}$$

The left-hand side of the differential equation can thus be expressed as

$$\frac{\partial y}{\partial t} = \mathrm{e}^{L(t)}\left[\frac{\partial \bar{y}}{\partial t} + \frac{\partial L}{\partial t}\bar{y}(t)\right] \tag{A.27}$$

$$= \mathrm{e}^{L(t)}\left[\frac{\partial \bar{y}}{\partial t} + \mathcal{L}(t)\bar{y}(t)\right]. \tag{A.28}$$

Also substituting for $y(t)$ using $\bar{y}(t)$ on the right-hand side, the whole equation then becomes

$$\mathrm{e}^{L(t)}\left[\frac{\partial \bar{y}}{\partial t} + \mathcal{L}(t)\bar{y}(t)\right] = \mathrm{e}^{L(t)}\mathcal{L}(t)\bar{y}(t) + f(y,t)\,, \tag{A.29}$$

and the two terms proportional to $\mathcal{L}(t)$ cancel, leaving us with

$$\frac{\partial \bar{y}}{\partial t} = \mathrm{e}^{-L(t)}f(y,t)\,. \tag{A.30}$$

By further defining the pre-conditioned function $\bar{f}(\bar{y},t)$ as

$$\bar{f}(\bar{y},t) = \mathrm{e}^{-L(t)}f\left(\mathrm{e}^{L(t)}\bar{y},t\right), \tag{A.31}$$

we recover a generic first-order ODE which now includes the pre-conditioner:

$$\frac{\partial \bar{y}}{\partial t} = \bar{f}(\bar{y}, t)\,. \tag{A.32}$$

Applying a numerical integrator such as the Dormand-Prince method now automatically includes the linear propagation with the highest possible accuracy, speeding up computation considerably.

The ptychographic iterative engine for DFG

Here I will derive the form of the update function in the ePIE and rPIE for an XFROG produced with a DFG process. In this process, the signal field ψ is given by

$$\psi(t) = E(t)G^*(t-\tau)\,. \tag{A.33}$$

The fundamental structure of the algorithm is to calculate ψ from the current guesses for $E(t)$ and $G(t)$, replace its magnitude with the corresponding slice of the measured XFROG trace, and then update the guesses using this updated field $\psi'(t)$. To obtain the form of this update, we can write the updated test pulse $E'(t)$ as

$$E'(t) = \frac{\psi'(t)}{G^*(t-\tau)} = \frac{G(t-\tau)}{|G(t-\tau)|^2}\psi'(t)\,, \tag{A.34}$$

that is, we simply apply eq. (A.33). However, this update will only be accurate where the gate pulse field $G(t-\tau)$ is strong; at times where $|G(t-\tau)|^2$ is weak, the update will be very noisy. Therefore, it is prudent to add a weighting function to the update. Where the gate field is weak, the update simply retains the previous guess for $E(t)$:

$$E'(t) = [1 - w(t)]E(t) + w(t)\frac{G(t-\tau)}{|G(t-\tau)|^2}\psi'(t)\,, \tag{A.35}$$

where $w(t)$ is near unity where the gate field is strong, and near zero where $G(t-\tau)$ is weak. Since we can express $E(t)$ in terms of $\psi(t)$ by using eq. (A.33),

$$E(t) = \frac{\psi(t)}{G^*(t-\tau)} = \frac{G(t-\tau)}{|G(t-\tau)|^2}\psi(t)\,, \tag{A.36}$$

this update can also be expressed as

$$E'(t) = E(t) + w(t)\frac{G(t-\tau)}{|G(t-\tau)|^2}\left[\psi'(t) - \psi(t)\right]. \tag{A.37}$$

A very simple weighting function is to use the normalised intensity of the gate field:

$$w(t) = \alpha \frac{|G(t-\tau)|^2}{|G(t-\tau)|^2_{\max}}. \tag{A.38}$$

This leads to the update function for DFG in the ePIE:

$$E'(t) = E(t) + \alpha \frac{G(t-\tau)}{|G(t-\tau)|^2_{\max}} \left[\psi'(t) - \psi(t)\right]. \tag{A.39}$$

The update function in the rPIE is obtained by using a different weight function which includes the regularisation component in the denominator:

$$w_{\mathrm{rPIE}}(t) = \frac{|G(t-\tau)|^2}{(1-\alpha)|G(t-\tau)|^2 + \alpha |G(t-\tau)|^2_{\max}}. \tag{A.40}$$

Using the form of the DFG field eq. (A.33) to obtain an updated guess for the gate pulse instead yields

$$G'^*(t-\tau) = \frac{\psi'(t)}{E(t)} = \frac{E^*(t)}{|E(t)|^2}\psi'(t). \tag{A.41}$$

Following the same procedure as for the test pulse, this update can be expressed as

$$G'^*(t-\tau) = G^*(t-\tau) + w(t)\frac{E^*(t)}{|E(t)|^2}\left[\psi'(t) - \psi(t)\right]. \tag{A.42}$$

In this case, the weighting function takes into account the normalised intensity of the test pulse instead:

$$w(t) = \beta \frac{|E(t)|^2}{|E(t)|^2_{\max}}, \tag{A.43}$$

which leads to the ePIE update function for the gate pulse:

$$G'(t-\tau) = G(t-\tau) + \beta \frac{E(t)}{|E(t)|^2_{\max}} \left[\psi'(t) - \psi(t)\right]^*, \tag{A.44}$$

which can again be improved by adding a regularisation term in the same manner as before.

Bibliography

[1] A. H. Zewail. *Femtochemistry: Atomic-Scale Dynamics of the Chemical Bond*. The Journal of Physical Chemistry A, 104(24):pp. 5660–5694, **2000**. ISSN 1089-5639. doi:10.1021/jp001460h [Cited on pp. 2 and 4.]

[2] M. A. M. Fox. *Quantum optics : an introduction*. Oxford University Press, **2006**. ISBN 0198566735 [Cited on p. 3.]

[3] M. Barbatti, et al. *Relaxation mechanisms of UV-photoexcited DNA and RNA nucleobases*. Proceedings of the National Academy of Sciences, 107(50):pp. 21453–21458, **2010**. ISSN 0027-8424. doi:10.1073/pnas.1014982107 [Cited on p. 4.]

[4] S. Matsika, et al. *Ultrafast Relaxation Dynamics of Uracil Probed via Strong Field Dissociative Ionization*. Journal of Physical Chemistry A, 117(48):pp. 12796–12801, **2013**. ISSN 1089-5639. doi:10.1021/jp408073d [Cited on p. 4.]

[5] V. Averbukh, Kolorenc, and Premysl. *Collective Interatomic Decay of Multiple Vacancies in Clusters*. Physical Review Letters, 103(18), **2009**. ISSN 00319007. doi:10.1103/PhysRevLett.103.183001 [Cited on p. 4.]

[6] R. Berera, R. van Grondelle, and J. T. M. Kennis. *Ultrafast transient absorption spectroscopy: principles and application to photosynthetic systems.* Photosynthesis research, 101(2-3):pp. 105–18, **2009**. ISSN 1573-5079. doi:10.1007/s11120-009-9454-y [Cited on pp. 5 and 48.]

[7] J. C. Travers, et al. *Ultrafast nonlinear optics in gas-filled hollow-core photonic crystal fibers [Invited]*. Journal of the Optical Society of America B, 28(12):pp. A11–A26,

2011. ISSN 0740-3224. doi:10.1364/JOSAB.28.000A11 [Cited on pp. 6, 66, 71, 94, and 96.]

[8] P. S. J. Russell, et al. *Hollow-core photonic crystal fibres for gas-based nonlinear optics*, **2014**. doi:10.1038/nphoton.2013.312 [Cited on p. 6.]

[9] N. Y. Joly, et al. *Bright spatially coherent wavelength-tunable deep-UV laser source using an Ar-filled photonic crystal fiber*. Physical Review Letters, 106(20), **2011**. ISSN 00319007. doi:10.1103/PhysRevLett.106.203901 [Cited on p. 6.]

[10] K. F. Mak, et al. *Tunable vacuum-UV to visible ultrafast pulse source based on gas-filled Kagome-PCF*. Optics express, 21(9):pp. 10942–53, **2013**. ISSN 1094-4087. doi:10.1364/OE.21.010942 [Cited on pp. 6, 7, 92, 105, and 131.]

[11] A. Ermolov, et al. *Characterization of few-fs deep-UV dispersive waves by ultra-broadband transient-grating XFROG*. Optics Letters, 41(23):p. 5535, **2016**. ISSN 0146-9592. doi:10.1364/OL.41.005535 [Cited on pp. 6, 92, 131, and 148.]

[12] G. R. Fowles. *Introduction to modern optics*. Holt, Rinehart and Winston, **1975**. ISBN 0030894042 [Cited on pp. 10, 22, 31, 32, and 67.]

[13] R. W. Boyd. *Nonlinear optics*. Academic Press, **2003**. ISBN 0121216829 [Cited on pp. 10, 19, 20, 29, 31, 33, and 94.]

[14] A. Monmayrant, S. Weber, and B. Chatel. *A newcomer's guide to ultrashort pulse shaping and characterization*. J. Phys. B: At. Mol. Opt. Phys. J. Phys. B: At. Mol. Opt. Phys, 43(43), **2010**. doi:10.1088/0953-4075/43/10/103001 [Cited on p. 17.]

[15] I. Walmsley, L. Waxer, and C. Dorrer. *The role of dispersion in ultrafast optics*. Review of Scientific Instruments, 72(1), **2001**. doi:10.1063/1.1330575 [Cited on pp. 17 and 141.]

[16] M. Th Hassan, et al. *Optical attosecond pulses and tracking the nonlinear response of bound electrons*. Nature, 530, **2016**. doi:10.1038/nature16528 [Cited on pp. 19 and 39.]

[17] V. G. Dmitriev, G. G. Gurzadyan, and D. N. Nikogosyan. *Handbook of Nonlinear Optical Crystals*. Springer, **1997**. ISBN 3540683925 [Cited on pp. 20, 131, and 133.]

[18] R. A. Lewis. *A review of terahertz sources*. Journal of Physics D: Applied Physics, 47(37):p. 374001, **2014**. ISSN 0022-3727. doi:10.1088/0022-3727/47/37/374001 [Cited on p. 21.]

[19] P. A. Franken, et al. *Generation of Optical Harmonics*. Physical Review Letters, 7(4):pp. 118–119, **1961**. ISSN 0031-9007. doi:10.1103/PhysRevLett.7.118 [Cited on p. 21.]

[20] Y. B. Band. *Light and matter : electromagnetism, optics, spectroscopy and lasers*. John Wiley, **2006**. ISBN 0471899313 [Cited on pp. 31, 32, and 36.]

[21] R. L. Sutherland, D. G. McLean, and S. Kirkpatrick. *Handbook of nonlinear optics*. Marcel Dekker, **2003**. ISBN 0824742435 [Cited on pp. 33 and 133.]

[22] W. Wood, C. Siders, and M. Downer. *Femtosecond growth dynamics of an underdense ionization front measured by spectral blueshifting*. IEEE Transactions on Plasma Science, 21(1):pp. 20–33, **1993**. ISSN 00933813. doi:10.1109/27.221098 [Cited on p. 36.]

[23] C. Vozzi, et al. *Optimal spectral broadening in hollow-fiber compressor systems*. Applied Physics B: Lasers and Optics, 80(3):pp. 285–289, **2005**. ISSN 09462171. doi:10.1007/s00340-004-1721-1 [Cited on pp. 37 and 163.]

[24] A. Couairon and A. Mysyrowicz. *Femtosecond filamentation in transparent media*. Physics Reports, 441:pp. 47–189, **2007**. doi:10.1016/j.physrep.2006.12.005 [Cited on p. 37.]

[25] L. V. Keldysh. *Ionization in the field of a strong electromagnetic wave*. Soviet Physics JETP, 20(5):pp. 1307–1314, **1965**. ISSN 0038-5646 [Cited on pp. 37 and 38.]

[26] M. V. Ammosov, N. B. Delone, and V. P. Krainov. *Tunnel Ionization Of Complex*

Atoms And Atomic Ions In Electromagnetic Field. Soviet Physics JETP, 64(6):pp. 1191–1194, **1986**. doi:10.1117/12.938695 [Cited on pp. 37 and 38.]

[27] A. M. Perelomov, V. S. Popov, and M. V. Terent 'ev. *Ionization of atoms in an alternating electric field*. Soviet Physics JETP, 23(50):pp. 1393–1409, **1966** [Cited on pp. 37 and 38.]

[28] G. L. Yudin and M. Y. Ivanov. *Nonadiabatic tunnel ionization: Looking inside a laser cycle*. Physical Review A, 64(1):p. 013409, **2001**. ISSN 1050-2947. doi:10.1103/PhysRevA.64.013409 [Cited on pp. 37 and 38.]

[29] P. B. Corkum. *Plasma perspective on strong field multiphoton ionization*. Physical Review Letters, 71(13):pp. 1994–1997, **1993**. ISSN 0031-9007. doi:10.1103/PhysRevLett.71.1994 [Cited on pp. 39, 40, and 171.]

[30] M. Lewenstein, P. Salières, and A. L'Huillier. *Phase of the atomic polarization in high-order harmonic generation*. Physical Review A, 52(6):pp. 4747–4754, **1995**. ISSN 10502947. doi:10.1103/PhysRevA.52.4747 [Cited on pp. 40 and 172.]

[31] M. Lewenstein, et al. *Theory of high-harmonic generation by low-frequency laser fields*. Physical Review A, 49(3):pp. 2117–2132, **1994**. ISSN 10502947. doi:10.1103/PhysRevA.49.2117 [Cited on pp. 40, 44, 170, and 172.]

[32] D. Popmintchev, et al. *Ultraviolet surprise: Efficient soft x-ray high-harmonic generation in multiply ionized plasmas.* Science, 350(6265):pp. 1225–31, **2015**. ISSN 1095-9203. doi:10.1126/science.aac9755 [Cited on p. 40.]

[33] P. M. Paul, et al. *Observation of a train of attosecond pulses from high harmonic generation*. Science, 292(5522):pp. 1689–1692, **2001**. ISSN 00368075. doi:10.1126/science.1059413 [Cited on p. 41.]

[34] X. Feng, et al. *Generation of isolated attosecond pulses with 20 to 28 femtosecond lasers*. Physical Review Letters, 103(18), **2009**. ISSN 00319007. doi:10.1103/PhysRevLett.103.183901 [Cited on p. 42.]

[35] G. Sansone, et al. *Isolated single-cycle attosecond pulses*. Science, 314(5798):pp. 443–446, **2006**. ISSN 00368075. doi:10.1126/science.1132838 [Cited on p. 42.]

[36] H. Mashiko, et al. *Double optical gating of high-order harmonic generation with carrier-envelope phase stabilized lasers*. Physical Review Letters, 100(10), **2008**. ISSN 00319007. doi:10.1103/PhysRevLett.100.103906 [Cited on p. 42.]

[37] M. J. Abel, et al. *Isolated attosecond pulses from ionization gating of high-harmonic emission*. Chemical Physics, 366(1-3):pp. 9–14, **2009**. ISSN 0301-0104. doi:10.1016/J.CHEMPHYS.2009.09.016 [Cited on p. 42.]

[38] M. Hentschel, et al. *Attosecond metrology*. Nature, 414(6863):pp. 509–513, **2001**. ISSN 0028-0836. doi:10.1038/35107000 [Cited on p. 42.]

[39] F. Krausz and M. Ivanov. *Attosecond physics*. Reviews of Modern Physics, 81(1):pp. 163–234, **2009**. ISSN 00346861. doi:10.1103/RevModPhys.81.163 [Cited on pp. 42, 43, and 172.]

[40] F. Calegari, et al. *Ultrafast electron dynamics in phenylalanine initiated by attosecond pulses*. Science, 346(6207):pp. 336–339, **2014**. ISSN 10959203. doi:10.1126/science.1254061 [Cited on p. 42.]

[41] E. Goulielmakis, et al. *Real-time observation of valence electron motion*. Nature, 466(7307):pp. 739–743, **2010**. ISSN 0028-0836. doi:10.1038/nature09212 [Cited on p. 42.]

[42] B. Shan and Z. Chang. *Dramatic extension of the high-order harmonic cutoff by using a long-wavelength driving field*. Physical Review A, 65(1):p. 4, **2002**. ISSN 10941622. doi:10.1103/PhysRevA.65.011804 [Cited on pp. 42 and 43.]

[43] M. B. Gaarde, J. L. Tate, and K. J. Schafer. *Macroscopic aspects of attosecond pulse generation*. Journal of Physics B: Atomic, Molecular and Optical Physics, 41(13):p. 132001, **2008**. ISSN 0953-4075. doi:10.1088/0953-4075/41/13/132001 [Cited on pp. 42 and 43.]

[44] A. Paul, et al. *Phase-Matching Techniques for Coherent Soft X-Ray Generation*. IEEE Journal of Quantum Electronics, 42(1):pp. 14–26, **2006**. ISSN 0018-9197. doi:10.1109/JQE.2005.859914 [Cited on p. 42.]

[45] B. Henke, E. Gullikson, and J. Davis. *X-Ray Interactions: Photoabsorption, Scattering, Transmission, and Reflection at E = 50-30,000 eV, Z = 1-92*. Atomic Data and Nuclear Data Tables, 54(2):pp. 181–342, **1993**. ISSN 0092-640X. doi:10.1006/ADND.1993.1013 [Cited on p. 44.]

[46] T. Popmintchev, et al. *Bright coherent ultrahigh harmonics in the kev x-ray regime from mid-infrared femtosecond lasers*. Science, 336(6086):pp. 1287–1291, **2012**. ISSN 10959203. doi:10.1126/science.1218497 [Cited on p. 43.]

[47] S. L. Cousin, et al. *High-flux table-top soft x-ray source driven by sub-2-cycle, CEP stable, 1.85-µm 1-kHz pulses for carbon K-edge spectroscopy*. Optics Letters, 39(18):p. 5383, **2014**. ISSN 0146-9592. doi:10.1364/OL.39.005383 [Cited on p. 43.]

[48] D. R. Austin and J. Biegert. *Strong-field approximation for the wavelength scaling of high-harmonic generation*. Physical Review A, 86, **2012**. doi:10.1103/PhysRevA.86.023813 [Cited on p. 44.]

[49] A. S. Johnson. *Long and Short Wavelength Optical Sources for Attosecond Science*. Ph.D. thesis, Imperial College London, **2017** [Cited on pp. 45, 47, 48, and 170.]

[50] J. Yano and V. K. Yachandra. *X-ray absorption spectroscopy.* Photosynthesis research, 102(2-3):pp. 241–54, **2009**. ISSN 1573-5079. doi:10.1007/s11120-009-9473-8 [Cited on p. 45.]

[51] C. Bressler, R. Abela, and M. Chergui. *Exploiting EXAFS and XANES for time-resolved molecular structures in liquids*. Zeitschrift für Kristallographie, 223(4-5/2008):pp. 307–321, **2008**. ISSN 0044-2968. doi:10.1524/zkri.2008.0030 [Cited on p. 45.]

[52] M. Saes, et al. *Observing Photochemical Transients by Ultrafast X-Ray Absorption Spectroscopy*. Physical Review Letters, 90(4):p. 4, **2003**. ISSN 10797114. doi:10.1103/PhysRevLett.90.047403 [Cited on p. 45.]

[53] A. S. Johnson, et al. *Measurement of sulfur L2,3 and carbon K edge XANES in a polythiophene film using a high harmonic supercontinuum*. Structural Dynamics, 3(6), **2016**. ISSN 23297778. doi:10.1063/1.4964821 [Cited on pp. 46, 47, 48, 158, and 175.]

[54] T. Etzkorn, et al. *Gas-phase absorption cross sections of 24 monocyclic aromatic hydrocarbons in the UV and IR spectral ranges*. Atmospheric Environment, 33:pp. 525–540, **1999** [Cited on p. 48.]

[55] K. Sunanda, et al. *Electronic states of carbon disulphide in the 5.5-11.8 eV region by VUV photo absorption spectroscopy*. Journal of Quantitative Spectroscopy and Radiative Transfer, 151:pp. 76–87, **2015**. ISSN 0022-4073. doi:10.1016/J.JQSRT.2014.08.020 [Cited on p. 48.]

[56] A. Takahashi, et al. *New femtosecond streak camera with temporal resolution of 180 fs*. In R. P. Trebino and I. A. Walmsley (eds.), *Proceedings Volume 2116 OE/LASE '94 23-29 January 1994*, volume 2116, pp. 275–284. International Society for Optics and Photonics, **1994**. doi:10.1117/12.175863 [Cited on p. 49.]

[57] T. Gaumnitz, et al. *Streaking of 43-attosecond soft-X-ray pulses generated by a passively CEP-stable mid-infrared driver*. Optics Express, 25(22):p. 27506, **2017**. ISSN 1094-4087. doi:10.1364/OE.25.027506 [Cited on p. 49.]

[58] M. Miranda, et al. *Simultaneous compression and characterization of ultrashort laser pulses using chirped mirrors and glass wedges*. Optics Express, 20(1):p. 688, **2012**. ISSN 1094-4087. doi:10.1364/OE.20.000688 [Cited on pp. 49 and 55.]

[59] R. Priemer. *Introductory Signal Processing*. WORLD SCIENTIFIC, **1990**. ISBN 978-9971-5-0919-4. doi:10.1142/0864 [Cited on pp. 49, 101, and 123.]

[60] R. Trebino. *Frequency-Resolved Optical Gating: The Measurement of Ultrashort Laser Pulses*. Springer US, **2000**. ISBN 146151181X [Cited on pp. 50, 51, 53, 54, 55, 124, and 131.]

[61] D. Kane and R. Trebino. *Characterization of arbitrary femtosecond pulses using frequency-resolved optical gating*. IEEE Journal of Quantum Electronics, 29(2):pp. 571–579, **1993**. ISSN 00189197. doi:10.1109/3.199311 [Cited on pp. 51 and 52.]

[62] M. V. Klibanov. *On the recovery of a 2-D function from the modulus of its Fourier transform*. Journal of Mathematical Analysis and Applications, 323(2):pp. 818–843, **2006**. ISSN 0022-247X. doi:10.1016/J.JMAA.2005.10.079 [Cited on p. 52.]

[63] R. Trebino, et al. *Measuring ultrashort laser pulses in the time-frequency domain using frequency-resolved optical gating*. Review of Scientific Instruments, 68(9):pp. 3277–3295, **1997**. ISSN 0034-6748. doi:10.1063/1.1148286 [Cited on pp. 52 and 131.]

[64] G. Taft, et al. *Measurement of 10-fs laser pulses*. IEEE Journal of Selected Topics in Quantum Electronics, 2(3):pp. 575–585, **1996**. doi:10.1109/2944.571757 [Cited on pp. 52, 54, and 150.]

[65] A. Baltuska, M. S. Pshenichnikov, and D. A. Wiersma. *Second-Harmonic Generation Frequency-Resolved Optical Gating in the Single-Cycle Regime*. IEEE Journal of Quantum Electronics, 35(4), **1999**. doi:10.1109/3.753651 [Cited on pp. 52 and 131.]

[66] K. W. DeLong, et al. *Pulse retrieval in frequency-resolved optical gating based on the method of generalized projections*. Optics Letters, 19(24):p. 2152, **1994**. ISSN 0146-9592. doi:10.1364/OL.19.002152 [Cited on pp. 55 and 56.]

[67] H. Stark. *Image recovery*. Academic Press, **1987**. ISBN 9780323145978 [Cited on p. 55.]

[68] D. J. Kane, et al. *Simultaneous measurement of two ultrashort laser pulses from a*

single spectrogram in a single shot. Journal of the Optical Society of America B, 14(4):p. 935, **1997**. ISSN 0740-3224. doi:10.1364/JOSAB.14.000935 [Cited on p. 56.]

[69] H. Anton. *Elementary linear algebra*. Wiley, **2010**. ISBN 9780470458211 [Cited on p. 56.]

[70] D. Kane. *Real-time measurement of ultrashort laser pulses using principal component generalized projections*. IEEE Journal of Selected Topics in Quantum Electronics, 4(2):pp. 278–284, **1998**. ISSN 1077260X. doi:10.1109/2944.686733 [Cited on p. 56.]

[71] G. P. G. P. Agrawal. *Nonlinear fiber optics*. Academic Press, **1989**. ISBN 0120451409 [Cited on pp. 59, 74, 75, and 83.]

[72] M. Born and E. Wolf. *Principles of optics : electromagnetic theory of propagation, interference and diffraction of light*. Cambridge University Press, **1999**. ISBN 0521642221 [Cited on pp. 59, 115, and 134.]

[73] E. A. J. Marcatili and R. A. Schmeltzer. *Hollow Metallic and Dielectric Waveguides for Long Distance Optical Transmission and Lasers*. Bell System Technical Journal, 43(4):pp. 1783–1809, **1964**. ISSN 00058580. doi:10.1002/j.1538-7305.1964.tb04108.x [Cited on pp. 60, 61, 62, 63, and 65.]

[74] E.-G. Neumann. *Single-Mode Fibers*. Springer-Verlag, **1988**. ISBN 978-3-540-48173-7 [Cited on p. 64.]

[75] L. Grüner-Nielsen, et al. *Dispersion Compensating Fibers*. Optical Fiber Technology, 6(2):pp. 164–180, **2000**. ISSN 1068-5200. doi:10.1006/OFTE.1999.0324 [Cited on p. 64.]

[76] M. J. Weber. *Handbook of optical materials*. CRC Press, **2003**. ISBN 0849335124 [Cited on p. 65.]

[77] J. C. Knight, et al. *All-silica single-mode optical fiber with photonic crystal cladding*. Optics Letters, 21(19):p. 1547, **1996**. ISSN 0146-9592. doi:10.1364/OL.21.001547 [Cited on p. 65.]

[78] J. C. Knight, et al. *Photonic band gap guidance in optical fibers*. Science, 282(5393):pp. 1476–8, **1998**. ISSN 1095-9203. doi:10.1126/SCIENCE.282.5393.1476 [Cited on p. 65.]

[79] D. Atkin, et al. *Full 2-D photonic bandgaps in silica/air structures*. Electronics Letters, 31(22):pp. 1941–1943, **1995**. ISSN 0013-5194. doi:10.1049/el:19951306 [Cited on p. 66.]

[80] J. C. Knight. *Photonic crystal fibres*. Nature, 424(6950):pp. 847–851, **2003**. ISSN 0028-0836. doi:10.1038/nature01940 [Cited on p. 66.]

[81] Y. Chen, et al. *Detailed study of macrobending effects in a wide transmission bandwidth hollow-core photonic bandgap fiber*. In *SPIE Photonics Europe 2016*, p. 98860X. **2016**. ISBN 9781510601314. ISSN 1996756X. doi:10.1117/12.2225861 [Cited on p. 66.]

[82] P. S. Light, et al. *Double photonic bandgap hollow-core photonic crystal fiber*. Optics Express, 17(18):p. 16238, **2009**. ISSN 1094-4087. doi:10.1364/OE.17.016238 [Cited on p. 66.]

[83] S.-J. Im, A. Husakou, and J. Herrmann. *Guiding properties and dispersion control of kagome lattice hollow-core photonic crystal fibers*. Optics Express, 17(15):p. 13050, **2009**. ISSN 1094-4087. doi:10.1364/OE.17.013050 [Cited on pp. 67, 68, and 69.]

[84] J.-L. Archambault, et al. *Loss calculations for antiresonant waveguides*. Journal of Lightwave Technology, 11(3):pp. 416–423, **1993**. ISSN 07338724. doi:10.1109/50.219574 [Cited on p. 68.]

[85] F. Yu, M. Xu, and J. C. Knight. *Experimental study of low-loss single-mode performance in anti-resonant hollow-core fibers*. Optics Express, 24(12):p. 12969, **2016**. ISSN 1094-4087. doi:10.1364/OE.24.012969 [Cited on p. 68.]

[86] R. M. Carter, et al. *Measurement of resonant bend loss in anti-resonant hollow core optical fiber*. Optics Express, 25(17):p. 20612, **2017**. ISSN 1094-4087. doi:10.1364/OE.25.020612 [Cited on pp. 69 and 96.]

[87] F. Köttig, et al. *Generation of microjoule pulses in the deep ultraviolet at megahertz repetition rates*. Optica, 4(10):p. 1272, **2017**. ISSN 2334-2536. doi:10.1364/OPTICA.4.001272 [Cited on pp. 69, 92, and 127.]

[88] F. Yu and J. C. Knight. *Negative Curvature Hollow-Core Optical Fiber*. IEEE Journal of Selected Topics and Quantum Electronics, 22(2), **2016**. doi:10.1109/JSTQE.2015.2473140 [Cited on p. 69.]

[89] F. Benabid, et al. *Stimulated Raman Scattering in Hydrogen-Filled Hollow-Core Photonic Crystal Fiber*. Science, 298(5592):pp. 399–402, **2002**. ISSN 00368075. doi:10.1126/science.1076408 [Cited on p. 69.]

[90] C. Wei, et al. *Negative curvature fibers*. Advances in Optics and Photonics, 9(3):p. 504, **2017**. ISSN 1943-8206. doi:10.1364/AOP.9.000504 [Cited on p. 69.]

[91] M. A. Finger, et al. *Accuracy of the capillary approximation for gas-filled kagomé-style photonic crystal fibers*. Optics letters, 39(4):pp. 821–4, **2014**. ISSN 1539-4794. doi:10.1364/OL.39.000821 [Cited on p. 69.]

[92] W. Chang, et al. *Influence of ionization on ultrafast gas-based nonlinear fiber optics*. Optics Express, 19(21):p. 21018, **2011**. ISSN 1094-4087. doi:10.1364/OE.19.021018 [Cited on p. 71.]

[93] P. Kinsler. *Optical pulse propagation with minimal approximations*. Physical Review A - Atomic, Molecular, and Optical Physics, 81(1), **2010**. ISSN 10502947. doi:10.1103/PhysRevA.81.013819 [Cited on pp. 71, 72, and 75.]

[94] a. Couairon, et al. *Practitioner's guide to laser pulse propagation models and simulation*. The European Physical Journal Special Topics, 199(1):pp. 5–76, **2011**. ISSN 1951-6355. doi:10.1140/epjst/e2011-01503-3 [Cited on pp. 71, 72, and 75.]

[95] M. Nisoli, S. De Silvestri, and O. Svelto. *Generation of high energy 10 fs pulses by a new pulse compression technique*. Citation: Appl. Phys. Lett, 68:p. 2793, **1996**. doi:10.1063/1.116609 [Cited on p. 78.]

[96] C. S. Gardner, et al. *Method for solving the Korteweg-deVries equation*. Physical Review Letters, 19(19):pp. 1095–1097, **1967**. ISSN 00319007. doi:10.1103/PhysRevLett.19.1095 [Cited on p. 80.]

[97] V. E. Zakharovand and A. B. Shabat. *Exact Theory of Two-Dimensional Self-Focusing and One- Dimensional Self-Modulation of Waves in Nonlinear Media*. Journal of Experimental and Theoretical Physics, 34(61):pp. 118–134, **1972**. ISSN 1063-7761. doi:10.1007/BF01075696 [Cited on p. 80.]

[98] F. Mitschke, C. Mahnke, and A. Hause. *Soliton Content of Fiber-Optic Light Pulses*. Applied Sciences, 7(6):p. 635, **2017**. ISSN 2076-3417. doi:10.3390/app7060635 [Cited on pp. 80, 82, 83, and 89.]

[99] N. Akhmediev and M. Karlsson. *Cherenkov radiation emitted by solitons in optical fibers*. Physical Review A, 51(3):pp. 2602–2607, **1995**. ISSN 10502947. doi:10.1103/PhysRevA.51.2602 [Cited on pp. 84 and 86.]

[100] J. N. Elgin, T. Brabec, and S. M. J. Kelly. *A perturbative theory of soliton propagation in the presence of third order dispersion*. Optics Communications, 114:pp. 321–328, **1995**. doi:10.1016/0030-4018(94)00602-Q [Cited on p. 86.]

[101] D. R. Austin, et al. *Dispersive wave blue-shift in supercontinuum generation*. Optics express, 14(25):pp. 11997–12007, **2006**. ISSN 1094-4087. doi:10.1364/OE.14.011997 [Cited on p. 87.]

[102] A. V. Husakou and J. Herrmann. *Supercontinuum Generation of Higher-Order Solitons by Fission in Photonic Crystal Fibers*. Physical Review Letters, 87(20):p. 203901, **2001**. ISSN 0031-9007. doi:10.1103/PhysRevLett.87.203901 [Cited on p. 89.]

[103] J. M. Dudley, G. Genty, and S. Coen. *Supercontinuum generation in photonic crystal fiber*. Reviews of Modern Physics, 78(4):pp. 1135–1184, **2006**. ISSN 00346861. doi:10.1103/RevModPhys.78.1135 [Cited on p. 89.]

[104] F. Reiter, et al. *Generation of sub-3 fs pulses in the deep ultraviolet*. Optics Letters, 35(13):p. 2248, **2010**. ISSN 0146-9592. doi:10.1364/OL.35.002248 [Cited on p. 91.]

[105] C. G. Durfee, et al. *Intense 8-fs pulse generation in the deep ultraviolet*. Optics Letters, 24(10):p. 697, **1999**. ISSN 0146-9592. doi:10.1364/OL.24.000697 [Cited on p. 91.]

[106] T. Fuji, T. Horio, and T. Suzuki. *Generation of 12 fs deep-ultraviolet pulses by four-wave mixing through filamentation in neon gas*. Optics Letters, 32(17):p. 2481, **2007**. ISSN 0146-9592. doi:10.1364/OL.32.002481 [Cited on p. 91.]

[107] J. Ringling, et al. *Tunable femtosecond pulses in the near vacuum ultraviolet generated by frequency conversion of amplified Ti:sapphire laser pulses*. Optics Letters, 18(23):p. 2035, **1993**. ISSN 0146-9592. doi:10.1364/OL.18.002035 [Cited on p. 91.]

[108] F. Rotermund and V. Petrov. *Generation of the fourth harmonic of a femtosecond Ti:sapphire laser*. Optics Letters, 23(13):p. 1040, **1998**. ISSN 0146-9592. doi:10.1364/OL.23.001040 [Cited on p. 91.]

[109] G. Cerullo and S. De Silvestri. *Ultrafast optical parametric amplifiers*. Review of Scientific Instruments, 74(1):pp. 1–18, **2003**. ISSN 0034-6748. doi:10.1063/1.1523642 [Cited on p. 91.]

[110] A. Cerqueira Sodre, et al. *Highly efficient generation of broadband cascaded four-wave mixing products*. Optics Express, 16(4):p. 2816, **2008**. ISSN 1094-4087. doi:10.1364/OE.16.002816 [Cited on p. 93.]

[111] M. Erkintalo, et al. *Cascaded phase matching and nonlinear symmetry breaking in fiber frequency combs*. Physical Review Letters, 109(22), **2012**. ISSN 00319007. doi:10.1103/PhysRevLett.109.223904 [Cited on p. 93.]

[112] F. Tani, et al. *Effect of anti-crossings with cladding resonances on ultrafast nonlinear dynamics in gas-filled photonic crystal fibers*. Photonics Research, 6(2):p. 84, **2018**. ISSN 2327-9125. doi:10.1364/PRJ.6.000084 [Cited on p. 96.]

[113] S. J. Farlow. *Partial differential equations for scientists and engineers*. Dover Publications, **1993**. ISBN 048667620X [Cited on pp. 96 and 97.]

[114] W. H. Press. *Numerical recipes : the art of scientific computing*. Cambridge University Press, **2007**. ISBN 0521880688 [Cited on pp. 96, 97, 136, and 138.]

[115] J. Dormand and P. Prince. *A family of embedded Runge-Kutta formulae*. Journal of Computational and Applied Mathematics, 6(1):pp. 19–26, **1980**. ISSN 0377-0427. doi:10.1016/0771-050X(80)90013-3 [Cited on p. 97.]

[116] J. C. Tremblay and T. Carrington. *Using preconditioned adaptive step size Runge-Kutta methods for solving the time- dependent Schrödinger equation*. The Journal of Chemical Physics, 1211(10):pp. 11535–174108, **2004**. doi:10.1063/1.1814103 [Cited on p. 98.]

[117] D. R. Austin and J. Biegert. *Attosecond pulse shaping using partial phase matching*. New Journal of Physics, 16(11):p. 113011, **2014**. ISSN 1367-2630. doi:10.1088/1367-2630/16/11/113011 [Cited on p. 98.]

[118] E. Hairer and G. Wanner. *Solving Ordinary Differential Equations II*. Springer-Verlag, **2000**. ISBN 9783540604525 [Cited on p. 99.]

[119] J. Dormand and P. Prince. *Runge-Kutta triples*. Computers & Mathematics with Applications, 12(9):pp. 1007–1017, **1986**. ISSN 0898-1221. doi:10.1016/0898-1221(86)90025-8 [Cited on p. 99.]

[120] D. M. Bishop and J. Pipin. *Improved dynamic hyperpolarizabilities and field-gradient polarizabilities for helium*. The Journal of Chemical Physics, 91(3549), **1989**. doi:10.1063/1.456886 [Cited on p. 99.]

[121] D. P. Shelton and J. E. Rice. *Measurements and Calculations of the Hyperpolarizabilities of Atoms and Small Molecules in the Gas Phase*. Chemical Reviews, 94(1):pp. 3–29, **1994**. ISSN 15206890. doi:10.1021/cr00025a001 [Cited on pp. 99 and 109.]

[122] M. Geissler, et al. *Light Propagation in Field-Ionizing Media: Extreme Nonlinear Optics*. Physical Review Letters, 83(15):pp. 2930–2933, **1999**. ISSN 0031-9007. doi:10.1103/PhysRevLett.83.2930 [Cited on pp. 99 and 101.]

[123] F. Tani, J. C. Travers, and P. St.J. Russell. *Multimode ultrafast nonlinear optics in optical waveguides: numerical modeling and experiments in kagomé photonic-crystal fiber*. Journal of the Optical Society of America B, 31(2):p. 311, **2014**. ISSN 0740-3224. doi:10.1364/JOSAB.31.000311 [Cited on p. 100.]

[124] A. Suda, et al. *Generation of sub-10-fs, 5-mJ-optical pulses using a hollow fiber with a pressure gradient*. Appl. Phys. Lett, 86:p. 111116, **2005**. doi:10.1063/1.1883706 [Cited on p. 103.]

[125] A. Bideau-Mehu, et al. *Measurement of refractive indices of neon, argon, krypton and xenon in the 253.7140.4 nm wavelength range. Dispersion relations and estimated oscillator strengths of the resonance lines*. Journal of Quantitative Spectroscopy and Radiative Transfer, 25(5):pp. 395–402, **1981**. ISSN 0022-4073. doi:10.1016/0022-4073(81)90057-1 [Cited on p. 109.]

[126] F. Frank, et al. *Invited Review Article: Technology for Attosecond Science*. Review of Scientific Instruments, 83(7):p. 071101, **2012**. ISSN 0034-6748. doi:10.1063/1.4731658 [Cited on p. 112.]

[127] C. Knapp and G. Carter. *The generalized correlation method for estimation of time delay*. IEEE Transactions on Acoustics, Speech, and Signal Processing, 24(4):pp. 320–327, **1976**. ISSN 0096-3518. doi:10.1109/TASSP.1976.1162830 [Cited on p. 113.]

[128] C. Scott and Craig. *Introduction to Optics and Optical Imaging*. **1998** [Cited on p. 121.]

[129] Eksma Optics. *Ultrathin nonlinear crystals* [Cited on p. 131.]

[130] R. Schalck and Society of Photo-optical Instrumentation Engineers. *The*

proper care of optics : cleaning, handling, storage, and shipping. ISBN 0819494577 [Cited on p. 132.]

[131] S. Linden, H. Giessen, and J. Kuhl. *XFROG - A new method for amplitude and phase characterization of weak ultrashort pulses*. Physica Status Solidi (B) Basic Research, 206(1):pp. 119–124, **1998**. ISSN 03701972. doi:10.1002/(SICI)1521-3951(199803)206:1¡119::AID-PSSB119¿3.0.CO;2-X [Cited on pp. 133 and 134.]

[132] S. Linden, J. Kuhl, and H. Giessen. *Amplitude and phase characterization of weak blue ultrashort pulses by downconversion*. Optics Letters, 24(8):pp. 569–571, **1999**. ISSN 0146-9592. doi:10.1364/OL.24.000569 [Cited on pp. 133 and 134.]

[133] B. Seifert, H. Stolz, and M. Tasche. *Nontrivial ambiguities for blind frequency-resolved optical gating and the problem of uniqueness*. Journal of the Optical Society of America B, 21(5):p. 1089, **2004**. ISSN 0740-3224. doi:10.1364/JOSAB.21.001089 [Cited on p. 134.]

[134] J. R. Fienup. *Phase retrieval algorithms: a comparison*. Applied Optics, 21(15):p. 2758, **1982**. ISSN 0003-6935. doi:10.1364/AO.21.002758 [Cited on p. 134.]

[135] J. M. Rodenburg and H. M. L. Faulkner. *A phase retrieval algorithm for shifting illumination*. Applied Physics Letters, 85(108), **2004**. doi:10.1063/1.1823034 [Cited on p. 134.]

[136] J. Rodenburg, A. Hurst, and A. Cullis. *Transmission microscopy without lenses for objects of unlimited size*. Ultramicroscopy, 107(2-3):pp. 227–231, **2007**. ISSN 0304-3991. doi:10.1016/J.ULTRAMIC.2006.07.007 [Cited on p. 134.]

[137] P. D. Nellist, B. C. McCallum, and J. M. Rodenburg. *Resolution beyond the 'information limit' in transmission electron microscopy*. Nature, 374(6523):pp. 630–632, **1995**. ISSN 0028-0836. doi:10.1038/374630a0 [Cited on p. 134.]

[138] P. Thibault, et al. *Probe retrieval in ptychographic coherent diffractive*

imaging. Ultramicroscopy, 109(4):pp. 338–343, **2009**. ISSN 03043991. doi:10.1016/j.ultramic.2008.12.011 [Cited on p. 135.]

[139] A. M. Maiden and J. M. Rodenburg. *An improved ptychographical phase retrieval algorithm for diffractive imaging*. Ultramicroscopy, 109(10):pp. 1256–1262, **2009**. ISSN 0304-3991. doi:10.1016/J.ULTRAMIC.2009.05.012 [Cited on p. 135.]

[140] D. Spangenberg, et al. *Time-domain ptychography*. Physical Review A - Atomic, Molecular, and Optical Physics, 91(2), **2015**. ISSN 10941622. doi:10.1103/PhysRevA.91.021803 [Cited on p. 135.]

[141] A. M. Heidt, et al. *Measurement of complex supercontinuum light pulses using time domain ptychography*. Optics Letters, pp. 1–5, **2016**. ISSN 0146-9592. doi:10.1364/OL.41.004903 [Cited on p. 135.]

[142] P. Sidorenko, et al. *Ptychographic reconstruction algorithm for frequency-resolved optical gating: super-resolution and supreme robustness*. Optica, 3(12):p. 1320, **2016**. ISSN 2334-2536. doi:10.1364/OPTICA.3.001320 [Cited on pp. 135, 136, 137, and 139.]

[143] D. Spangenberg, et al. *Ptychographic ultrafast pulse reconstruction*. Optics Letters, 40(6):p. 1002, **2015**. ISSN 0146-9592. doi:10.1364/OL.40.001002 [Cited on p. 135.]

[144] A. Maiden, D. Johnson, and P. Li. *Further improvements to the ptychographical iterative engine*. Optica, 4(7):p. 736, **2017**. ISSN 2334-2536. doi:10.1364/OPTICA.4.000736 [Cited on p. 135.]

[145] D.-M. Spangenberg, et al. *All-optical implementation of a time-domain ptychographic pulse reconstruction setup*. Applied Optics, 55(21):p. 5771, **2016**. ISSN 0003-6935. doi:10.1364/AO.55.005771 [Cited on pp. 135 and 139.]

[146] T. Witting, et al. *Time-domain ptychography of over-octave-spanning laser pulses in the single-cycle regime*. Opt. Lett., 41(18):pp. 4218–4221, **2016**. ISSN 0146-9592. doi:10.1364/OL.41.004218 [Cited on p. 135.]

[147] J. Hyyti, et al. *Interferometric time-domain ptychography for ultrafast pulse characterization*. Optics Letters, 42(11):p. 2185, **2017**. ISSN 0146-9592. doi:10.1364/OL.42.002185 [Cited on p. 135.]

[148] R. L. Fork, O. E. Martinez, and J. P. Gordon. *Negative dispersion using pairs of prisms*. Optics Letters, 9(5):p. 150, **1984**. ISSN 0146-9592. doi:10.1364/OL.9.000150 [Cited on p. 141.]

[149] A. Ermolov, et al. *Low loss hollow optical-waveguide connection from atmospheric pressure to ultra-high vacuum*. Applied Physics Letters, 103(26), **2013**. ISSN 00036951. doi:10.1063/1.4860947 [Cited on p. 142.]

[150] P. O 'shea, et al. *Increased-bandwidth in ultrashort-pulse measurement using an angle-dithered nonlinear-optical crystal*. Optics, 7(10):pp. 342–349, **2000**. doi:10.1364/OE.7.000342 [Cited on p. 153.]

[151] E. R. Simpson. *Attosecond transient absorption spectroscopy in atomic species*. Ph.D. thesis, Imperial College London, **2016** [Cited on p. 157.]

[152] T. Balciunas, et al. *A strong-field driver in the single-cycle regime based on self-compression in a kagome fibre*. Nature Communications, 6, **2015**. doi:10.1038/ncomms7117 [Cited on p. 158.]

[153] F. Tani, et al. *Continuously wavelength-tunable high harmonic generation via soliton dynamics*. Optics Letters, 42:p. 1768, **2017**. ISSN 0146-9592. doi:10.1364/OL.42.001768 [Cited on p. 158.]

[154] D. R. Austin, et al. *Spatio-temporal characterization of intense few-cycle 2 μm pulses*. Optics Express, 24(21):p. 24786, **2016**. ISSN 1094-4087. doi:10.1364/OE.24.024786 [Cited on pp. 163 and 175.]

[155] A. Baltuška, T. Fuji, and T. Kobayashi. *Controlling the Carrier-Envelope Phase of Ultrashort Light Pulses with Optical Parametric Amplifiers*. Physical Review Letters,

88(13):p. 4, **2002**. ISSN 10797114. doi:10.1103/PhysRevLett.88.133901 [Cited on p. 169.]

[156] P. Balcou, et al. *Generalized phase-matching conditions for high harmonics: The role of field-gradient forces*. Physical Review A, 55(4):pp. 3204–3210, **1997**. ISSN 10941622. doi:10.1103/PhysRevA.55.3204 [Cited on p. 172.]

[157] Z. Chang. *Fundamentals of attosecond optics*. CRC Press, **2011**. ISBN 9781420089370 [Cited on p. 173.]

[158] A. Gordon and F. X. Kärtner. *Quantitative modeling of single atom high harmonic generation*. Physical Review Letters, 95(22), **2005**. ISSN 00319007. doi:10.1103/PhysRevLett.95.223901 [Cited on p. 173.]

[159] R. López-Martens, et al. *Amplitude and Phase Control of Attosecond Light Pulses*. Physical Review Letters, 94(3):p. 033001, **2005**. ISSN 0031-9007. doi:10.1103/PhysRevLett.94.033001 [Cited on p. 176.]

Republishing permissions

Figures 3.2a and 3.2b

From: pubscopyright <copyright@osa.org>
Sent: 28 February 2018 20:35
To: Brahms, Christian; pubscopyright
Subject: RE: Request to reproduce figure in PhD thesis

Dear Christian Brahms,

Thank you for contacting The Optical Society.

For the use of figure 1 from John C. Travers, Wonkeun Chang, Johannes Nold, Nicolas Y. Joly, and Philip St. J. Russell, "Ultrafast nonlinear optics in gas-filled hollow-core photonic crystal fibers [Invited]," J. Opt. Soc. Am. B 28, A11-A26 (2011):

> OSA considers your requested use of its copyrighted material to be Fair Use under United States Copyright Law. It is requested that a complete citation of the original material be included in any publication.

While your publisher should be able to provide additional guidance, OSA prefers the below citation formats:

> For citations in figure captions:
>
> [Reprinted/Adapted] with permission from ref [x], [Publisher]. (with full citation in reference list)
>
> For images without captions:
>
> Journal Vol. #, first page (year published) An example: J. Opt. Soc. Am. B 28, A11 (2011)

OSA considers this email to be sufficient authorization for the use of the requested material.

Let me know if you have any questions.

Kind Regards,

Rebecca Robinson

Rebecca Robinson
February 28, 2018
Authorized Agent, The Optical Society